Chapter 10

Configuring RED Drop Profiles

You can configure two parameters to control congestion at the output stage. The first parameter defines the delay-buffer bandwidth, which provides packet buffer space to absorb burst traffic up to the specified duration of delay. Once the specified delay buffer becomes full, packets with 100 percent drop probability are dropped from the head of the buffer. For more information, see "Configuring the Scheduler Buffer Size" on page 122.

The second parameter defines the drop probabilities across the range of delay-buffer occupancy, supporting the random early detection (RED) process. When the number of packets queued is greater than the ability of the routing platform to empty a queue, the queue requires a method for determining which packets to drop from the network. To address this, the JUNOS software provides the option of enabling RED on individual queues.

Depending on the drop probabilities, RED might drop many packets long before the buffer becomes full, or it might drop only a few packets even if the buffer is almost full.

A *drop profile* is a mechanism of RED that defines parameters that allow packets to be dropped from the network. Drop profiles define the meanings of the loss priorities.

When you configure drop profiles, there are two important values: the queue fullness and the drop probability. The *queue fullness* represents a percentage of the memory used to store packets in relation to the total amount that has been allocated for that specific queue. Similarly, the *drop probability* is a percentage value that correlates to the likelihood that an individual packet is dropped from the network. These two variables are combined in a graph-like format, as shown in Figure 9 on page 112.

Figure 9 shows both a segmented and an interpolated graph. Although the formation of these graph lines is different, the application of the profile is the same. When a packet reaches the head of the queue, a random number between 0 and 100 is calculated by the routing platform. This random number is plotted against the drop profile using the current queue fullness of that particular queue. When the random number falls above the graph line, the packet is transmitted onto the physical media. When the number falls below graph the line, the packet is dropped from the network.

Figure 9: Segmented and Interpolated Drop Profiles

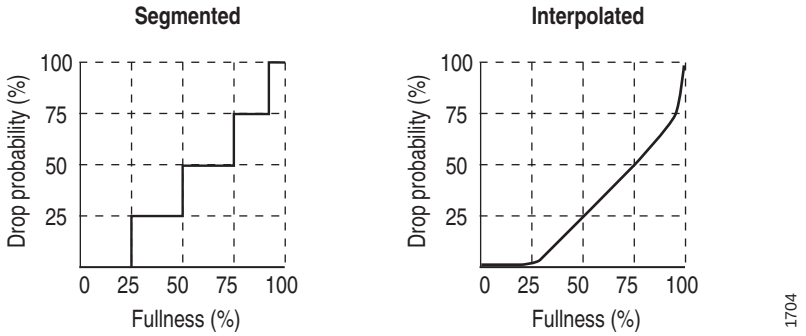

By defining multiple fill levels and drop probabilities, you create a segmented drop profile. The line segments are defined in terms of the following graphical model: in the first quadrant, the x axis represents the fill level, and the y axis represents the drop probability. The initial line segment spans from the origin (0,0) to the point (<l1>, <p1>); a second line runs from (<l1>, <p1>) to (<l2>, <p2>) and so forth, until a final line segment connects (100, 100). The software automatically constructs a drop profile containing 64 fill levels at drop probabilities that approximate the calculated line segments.

NOTE: If you configure the interpolate statement, you can specify more than 64 pairs, but the system generates only 64 discrete entries.

You specify drop probabilities in the drop profile section of the class-of-service (CoS) configuration hierarchy and reference them in each scheduler configuration. For each scheduler, you can configure multiple separate drop profiles, one for each combination of loss priority (low, medium-low, medium-high, or high) and IP transport protocol (TCP or non-TCP).

You can configure a maximum of 32 different drop profiles.

To configure RED drop profiles, include the following statements at the [edit class-of-service] hierarchy level of the configuration:

```
class-of-service {
    drop-profiles {
        profile-name {
            fill-level percentage drop-probability percentage;
            interpolate {
                drop-probability [ values ];
                fill-level [ values ];
            }
        }
    }
}
```

This chapter discusses the following topics:

Default Drop Profile

By default, if you configure no drop profiles, RED is still in effect and functions as the primary mechanism for managing congestion. In the default RED drop profile, when the fill-level is 0 percent, the drop probability is 0 percent. When the fill-level is 100 percent, the drop probability is 100 percent.

As a backup method for managing congestion, tail dropping takes effect when congestion of small packets occurs. On M320 and T-series platforms, the software supports *tail-RED*, which means that when tail dropping occurs, the software uses RED to execute intelligent tail drops. On other platforms, the software executes tail drops unconditionally.

Configuring RED Drop Profiles

You enable RED by applying a drop profile to a scheduler. When RED is operational on an interface, the queue no longer drops packets from the tail of the queue. Rather, packets are dropped after they reach the head of the queue.

To configure a drop profile, include the drop-profiles statement at the [edit class-of-service] hierarchy level:

```
[edit class-of-service]
drop-profiles {
    profile-name {
        fill-level percentage drop-probability percentage;
        interpolate {
            drop-probability [ values ];
            fill-level [ values ];
        }
    }
}
```

In this configuration, include either the interpolate statement and its options, or the fill-level and drop-probability *percentage* values. These two alternatives enable you to configure either each drop probability at up to 64 fill-level/drop-probability paired values, or a profile represented as a series of line segments, as shown in Figure 9 on page 112.

After you configure a drop profile, you must assign the drop profile to a drop-profile map, and assign the drop-profile map to a scheduler, as discussed in “Configuring Schedulers” on page 117.

Packet Loss Priority

Loss priority settings help determine which packets are dropped from the network during periods of congestion. The software supports multiple packet loss priority (PLP) designations: low and high. (In addition, medium-low and medium-high PLPs are supported when you configure tricolor marking, as discussed in "Configuring Tricolor Marking" on page 173.) You can set PLP by configuring a behavior aggregate or multifield classifier, as discussed in "Classifying Packets by Behavior Aggregate" on page 47 and "Classifying Packets Based on Various Packet Header Fields" on page 65.

A drop-profile map examines the loss priority setting of an outgoing packet: high, medium-high, medium-low, low, or any. In addition, some Layer 4 protocol information is examined to determine if the packet is associated with the Transmission Control Protocol (TCP), non-TCP, or any form of traffic.

Obviously, *low*, *medium-low*, *medium-high,* and *high* are relative terms, which by themselves have no meaning. Drop profiles define the meanings of the loss priorities. In the following example, the low-drop drop profile defines the meaning of low PLP as a 10 percent drop probability when the fill level is 75 percent, and a 40 percent drop probability when the fill level is 95 percent. The high-drop drop profile defines the meaning of high PLP as a 50 percent drop probability when the fill level is 25 percent, and a 90 percent drop probability when the fill level is 50 percent.

In this 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.

```
[edit class-of-service]
drop-profiles {
    low-drop {
        interpolate {
            drop-probability [ 10 40 ];
            fill-level [ 75 95 ];
        }
    }
    high-drop {
        interpolate {
            drop-probability [ 50 90 ];
            fill-level [ 25 50 ];
        }
    }
}

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;
        }
    }
}
```

For more information, see "Configuring Schedulers" on page 117.

Example: Configuring RED Drop Profiles

Create a segmented configuration and an interpolated configuration that correspond to the graphs in Figure 10. The values defined in the configuration are matched to represent the data points in the graph line. In this example, the drop probability is 25 percent when the queue is 50 percent full. The drop probability increases to 50 percent when the queue is 75 percent full.

Figure 10: Segmented and Interpolated Drop Profiles

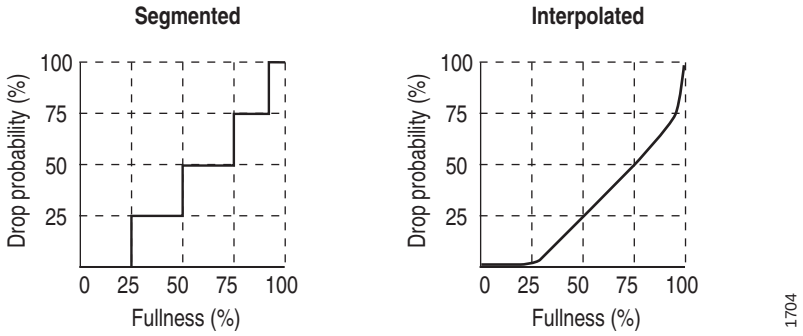

Segmented

```
class-of-service {
    drop-profiles {
        segmented-style-profile {
            fill-level 25 drop-probability 25;
            fill-level 50 drop-probability 50;
            fill-level 75 drop-probability 75;
            fill-level 95 drop-probability 100;
        }
    }
}
```

To create the profile's graph line, the software begins at the bottom-left corner, representing a 0 percent fill level and a 0 percent drop probability. This configuration draws a line directly to the right until it reaches the first defined fill level, 25 percent for this configuration. The software then continues the line vertically until the first drop probability is reached. This process is repeated for all of the defined levels and probabilities until the top-right corner of the graph is reached.

Interpolated Create a smoother graph line by configuring the profile with the interpolate statement. This allows the software to automatically generate 64 data points on the graph beginning at (0, 0) and ending at (100, 100). Along the way, the graph line intersects specific data points, which you define as follows:

```
class-of-service {
    drop-profiles
        interpolated-style-profile {
            interpolate {
                fill-level [ 50 75 ];
                drop-probability [ 25 50 ];
            }
        }
    }
}
```