

Chapter 4

Configuring Dropping Behavior with RED and WRED

This chapter provides information for configuring dropping behavior using RED and WRED on the E-series router.

QoS topics are discussed in the following sections:

- Dropping Behavior Overview on page 25
- RED and WRED Overview on page 26
- Configuring RED on page 27
- RED Configuration Examples on page 28
- Configuring WRED on page 30
- WRED Configuration Examples on page 32
- Monitoring RED and WRED on page 35

Dropping Behavior Overview

Drop profiles control the dropping behavior of a set of egress queues. They define the range within the queue where random early detection (RED) operates, the maximum percentage of packets to drop, and sensitivity to bursts of packets. Weighted random early detection (WRED) is an extension to RED that enables you to assign different RED drop profiles to each color of traffic.

The purpose of RED and WRED is to signal end-to-end protocols, such as TCP, that the router is becoming congested along a particular egress path. The intent is to trigger TCP congestion avoidance in a random set of TCP flows before congestion becomes severe and causes tail dropping on a large number of flows. Tail dropping can lead to TCP slow-starts, and tail dropping on a large number of flows results in global synchronization.

By default, tail dropping occurs when the length of a queue exceeds a threshold. Drop profiles allow you to employ active queue management by specifying RED and WRED parameters to be applied to an egress queue.

Congestion of an egress queue occurs when the rate of traffic destined for the queue exceeds the rate of traffic draining from the queue; the queue fills to its limit, and any further traffic destined to it must be discarded until there is room in the queue. RED and WRED monitor average queue length over time to detect incipient congestion.

You can combine drop profiles and queue profiles within a queue rule of a QoS profile to specify up to 256 unique queuing behaviors within the router. You can then associate these queuing behaviors in any combination with any of the egress queues.

Related Topics

- [Queuing and Buffer Management Overview on page 17](#)


RED and WRED Overview

The scheduler maintains an average queue length for each queue configured for RED. When a packet is enqueued, the current queue length is weighted into the average queue length based on the average-length exponent in the drop profile.

- Small exponent values weight the current queue length heavily, so the average queue length is more responsive to transient bursts.
- Large exponent values weight the current queue length lightly, so the average queue length is less responsive to bursts.

When the average queue length exceeds the minimum threshold, RED begins randomly dropping packets. While the average queue length increases toward the maximum threshold, RED drops packets with increasing frequency, up to the maximum drop probability. When the average queue length exceeds the maximum drop threshold, all packets are dropped. Figure 2 shows this behavior.

Figure 2: Packets Dropped as Queue Length Increases



WRED is an extension of RED that allows you to assign different RED drop thresholds to each color of traffic. The router assigns a color to each packet. Committed means green, conformed means yellow, and exceeded means red. When the queue fills above the exceeded threshold, the router drops red packets, but still queues yellow and green packets. When the queue fills above the conformed drop threshold, the router queues only green packets.

Configuring RED

Each line module supports a default drop profile and 15 configurable drop profiles. You can configure the default drop profile on all E-series line modules except for the ES2 10G LM.

To configure RED:

1. Create a drop profile and enter Drop Profile Configuration mode.

```
host1(config)#drop-profile internetDropProfile
host1(config-drop-profile)#
```

You can configure up to 16 drop profiles.

2. Set the average-length exponent, which specifies the exponent used to weight the average queue length over time, controlling WRED responsiveness.

```
host1(config-drop-profile)#average-length-exponent 9
```

- Specifying an average-length exponent enables the RED average queue length computation.
- A higher value smooths out the average and slows WRED reaction to congestion and decongestion, accommodating short bursts without dropping. Too large a value can smooth the average to the point that WRED does not react at all.
- A lower value speeds up WRED reaction. Too low a value can cause overreaction to short bursts, dropping packets unnecessarily.

3. (Optional) Set the minimum and maximum threshold for committed traffic.

```
host1(config-drop-profile)#committed-threshold percent 30 90 4
```

4. (Optional) Set the minimum and maximum threshold for conformed traffic.

```
host1(config-drop-profile)#conformed-threshold percent 25 90 5
```

5. (Optional) Set the minimum and maximum threshold for exceeded traffic.

```
host1(config-drop-profile)#exceeded-threshold percent 20 90 6
```

The thresholds specify a linear relationship between average queue length and drop probability.

You can express thresholds as either percentages of maximum queue size by including the keyword **percent**, or as absolute byte values by omitting the keyword.

Related Topics

- Configuring WRED on page 30
- Monitoring RED and WRED on page 35
- **average-length-exponent** command
- **committed-threshold** command
- **conformed-threshold** command
- **drop-profile** command
- **exceeded-threshold** command

RED Configuration Examples

This section describes how to configure the RED average queue length computation, configure RED for colored traffic, and configure RED so that packets are dropped without regard to color.

Configuring Average Queue Length

To enable calculation of average queue length, create a drop profile with a nonzero average-length exponent, reference the drop profile within a QoS profile, and attach the QoS profile to an interface. The following drop profile enables the average queue length calculation, but does not initiate RED dropping behavior:

```
host1(config)#drop-profile averageOnly
host1(config-drop-profile)#average-length-exponent 10
```

Configuring Thresholds

You can specify different dropping behavior for committed (green), conformed (yellow), and exceeded (red) packets by specifying a minimum queue threshold, maximum queue threshold, and maximum drop probability for each color of traffic.

By default, conformed threshold and exceeded threshold take the same values as the committed threshold. Therefore, if you specify only a committed threshold, conformed and exceeded traffic is treated like committed traffic. Similarly, if you specify a conformed threshold without an exceeded threshold, exceeded traffic is treated like committed traffic.

The following drop profiles result in identical behavior:

```
host1(config)#drop-profile colorblind1
host1(config-drop-profile)#committed-threshold percent 30 90 5
host1(config-drop-profile)#exit
```


```
host1(config)#drop-profile colorblind2
host1(config-drop-profile)#committed-threshold percent 30 90 5
host1(config-drop-profile)#conformed-threshold percent 30 90 5
host1(config-drop-profile)#exit
```

```
host1(config)#drop-profile colorblind3
host1(config-drop-profile)#committed-threshold percent 30 90 5
host1(config-drop-profile)#conformed-threshold percent 30 90 5
host1(config-drop-profile)#exceeded-threshold percent 30 90 5
```

Configuring Color-Blind RED

You can configure RED so that packets are dropped without regard to color. To do so, you combine a drop profile that has a committed threshold configured with a queue profile that specifies the same queue length for committed, conformed, and exceeded packets, as shown in Figure 3 on page 29.

Figure 3: Color-Blind RED Drop Profile with Colorless Queue Profile




In the following example, the drop profile and queue profile combine to specify the following:

- When the average queue length is between 30 percent full (30 KB) and 90 percent full (90 KB), up to 5 percent of the packets are randomly dropped regardless of their color.
- When the average queue length is greater than 90 percent, all packets are dropped regardless of color.

```
host1(config)#drop-profile nocolor
host1(config-drop-profile)#committed-threshold percent 30 90 5
host1(config-drop-profile)#exit
host1(config)#queue-profile colorless
host1(config-queue)#committed-length 100000 100000
host1(config-queue)#conformed-fraction 100
host1(config-queue)#exceeded-fraction 100
```

To achieve the same drop treatment for each color, you can specify color-blind RED in combination with a color-sensitive queue profile, as shown in Figure 4.

Figure 4: Color-Blind RED Drop Profile with Color-Sensitive Queue Profile

In the following example, the drop profile and queue profile combine to specify the following:

- When the average queue length is between 30 percent full (30 KB) and 90 percent full (90 KB), up to 5 percent of the packets are dropped randomly. In this case, the maximum queue length is 100 KB for green packets, 50 KB for yellow packets, and 25 KB for red packets. Therefore, the router randomly drops:
 - Red packets when the average queue length is between 7.5 KB and 22.5 KB
 - Yellow packets when the average queue length is between 15 KB and 45 KB
 - Green packets when the average queue length is between 30 KB and 90 KB
- When the average queue length is greater than 90 percent of the maximum queue length, all packets are dropped. Therefore, the router drops:
 - Red packets when the average queue length is greater than 22.5 KB
 - Yellow packets when the average queue length is greater than 45 KB
 - Green packets when the average queue length is greater than 90 KB

```
host1(config)#drop-profile colorblindRed
host1(config-drop-profile)#committed-threshold percent 30 90 5
host1(config-drop-profile)#exit
host1(config)#queue-profile colorSensitive
host1(config-queue)#committed-length 100000 100000
```

Configuring WRED

The main difference between RED and WRED is that WRED deals with different colored packets. The router assigns a color to each packet. Committed means green, conformed means yellow, and exceeded means red.

Each line module supports a default drop profile and 15 configurable drop profiles.

WRED is not supported on the ES2 10G Uplink LM. On the ES2 10G LM, you must configure WRED in one of the 15 configurable drop profiles; you cannot configure its default drop profile.

To configure WRED:

1. Create a drop profile and enter Drop Profile Configuration mode.

```
host1(config)#drop-profile internetDropProfile
host1(config-drop-profile)#
```

You can configure up to 16 drop profiles.

2. Set the average-length exponent, which specifies the exponent used to weight the average queue length over time, controlling WRED responsiveness.

```
host1(config-drop-profile)#average-length-exponent 9
```

- Specifying an average-length exponent enables the RED average queue length computation.
- A higher value smooths out the average and slows WRED reaction to congestion and decongestion, accommodating short bursts without dropping. Too large a value can smooth the average to the point that WRED does not react at all.
- A lower value speeds up WRED reaction. Too low a value can cause overreaction to short bursts, dropping packets unnecessarily.

3. (Optional) Set the minimum and maximum threshold for committed traffic.

```
host1(config-drop-profile)#committed-threshold percent 30 90 4
```

4. (Optional) Set the minimum and maximum threshold for conformed traffic.

```
host1(config-drop-profile)#conformed-threshold percent 25 90 5
```

5. (Optional) Set the minimum and maximum threshold for exceeded traffic.

```
host1(config-drop-profile)#exceeded-threshold percent 20 90 6
```

The thresholds specify a linear relationship between average queue length and drop probability.

You can express thresholds as either percentages of maximum queue size by including the keyword **percent**, or as absolute byte values by omitting the keyword.

Related Topics

- Configuring RED on page 27
- Monitoring RED and WRED on page 35
- **average-length-exponent** command
- **committed-threshold** command
- **conformed-threshold** command

- **drop-profile** command
- **exceeded-threshold** command

WRED Configuration Examples

This section explains how to configure different treatment of colored packets, different drop behavior for each queue, RED and dynamic queue thresholds, and average queue lengths for WRED.


Configuring Different Treatment of Colored Packets for WRED

Figure 5 on page 32 shows a WRED drop profile that yields progressively more aggressive drop treatment for each color. Exceeded traffic is dropped over a wider range and with greater maximum drop probability than conformed or committed traffic. Conformed traffic is dropped over a wider range and with greater maximum drop probability than committed traffic.

The commands to configure this example are:

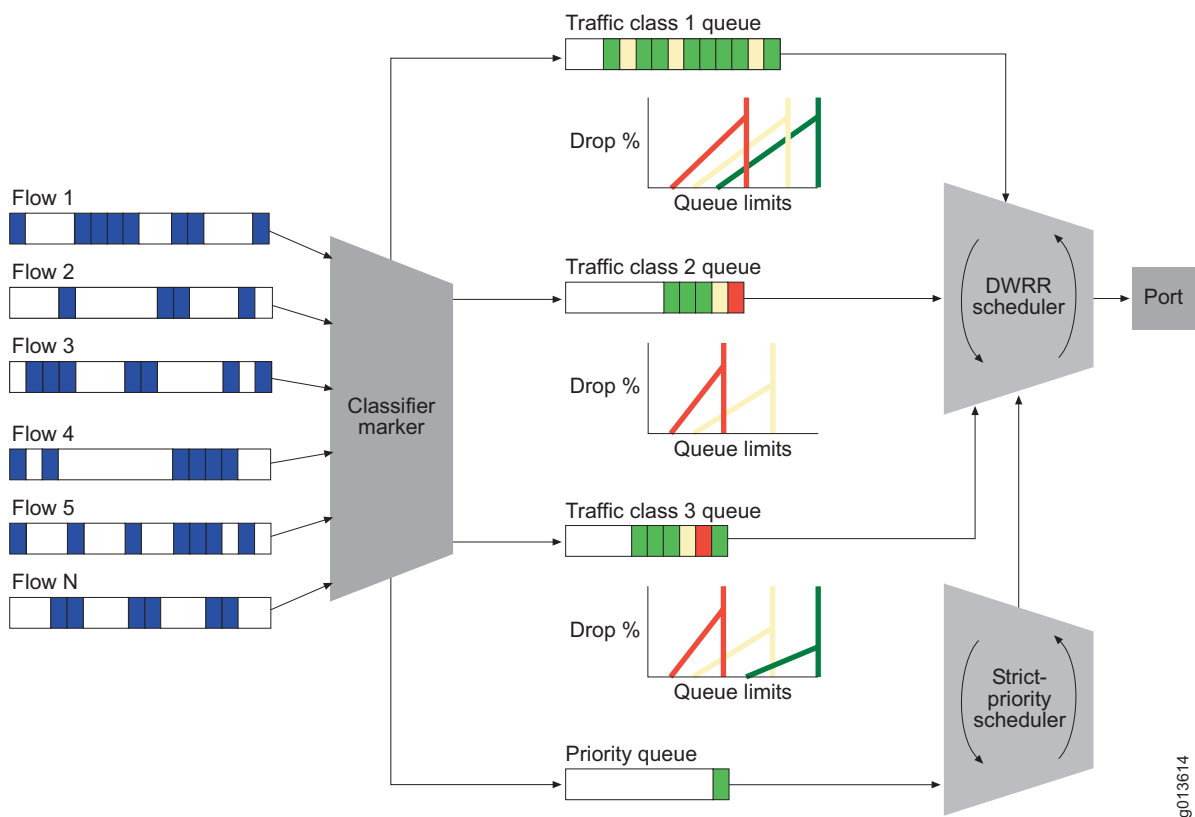
```
host1(config)#drop-profile wredColored
host1(config-drop-profile)#committed-threshold percent 30 90 3
host1(config-drop-profile)#conformed-threshold percent 25 90 5
host1(config-drop-profile)#exceeded-threshold percent 20 90 10
```

Figure 5: Different Treatment of Colored Packets



Defining Different Drop Behavior for Each Traffic Class

You can define different dropping behaviors for each traffic class in the router. By doing so, you can assign less aggressive drop profiles to higher-priority queues and more aggressive drop profiles to lower-priority queues. Figure 6 shows an example that classifies packets into one of four traffic classes. Each traffic class has a different queueing behavior, drop treatment, and scheduler treatment.

Figure 6: Defining Different Drop Behavior for Each Queue

Configuring WRED and Dynamic Queue Thresholds

RED typically operates on fixed-size queues, and you can configure the router to use fixed-size queues. However, by default, the router employs dynamic queue thresholds to provide a good balance between sharing the egress buffer memory between queues and protecting an individual queue's claim on its fair share of the egress memory. Fixed-size queues become problematic as the number of configured queues scales into the thousands, because allocating disjointed partitions of buffer memory to each queue means the allocations become quite small, and most likely not all queues are simultaneously active.

In general, you use queues as follows:


- Fixed-size queues on core routers and core-facing interfaces where the number of queues is relatively small (tens or hundreds, but not thousands).
- Dynamic queues on edge-facing interfaces where the number of queues is relatively large (thousands).

As shown in Figure 7 on page 35, queue lengths extend to oversubscribe memory when aggregate memory utilization is low, and contract to strictly partition memory when memory utilization is high. Dynamic thresholding enforces fairness when free buffers are scarce and promotes sharing when buffers are plentiful. Dynamic queue thresholds are discussed in *Queuing and Buffer Management Overview* on page 17. Figure 7 on page 35 illustrates WRED behavior with dynamic queue thresholding.

To configure WRED to run on queues whose limits dynamically expand and contract, use the **percent** keyword when you configure thresholds in a drop profile. For example:

```
host1(config)#drop-profile internetDropProfile
host1(config-drop-profile)#average-length-exponent 9
host1(config-drop-profile)#committed-threshold percent 30 90 4
host1(config-drop-profile)#conformed-threshold percent 25 90 5
host1(config-drop-profile)#exceeded-threshold percent 20 90 6
```

Figure 7: WRED and Dynamic Queue Thresholding



Monitoring RED and WRED

To monitor drop profiles, see:

- Monitoring Drop Profiles for RED and WRED on page 317

