WRED — Explicit Congestion Notification

Feature History

Release	Modification
12.2(8)T	This feature was introduced.

This document describes the WRED — Explicit Congestion Notification feature in Cisco IOS Release 12.2(8)T. It includes the following sections:

Feature Overview

Currently, the congestion control and avoidance algorithms for Transmission Control Protocol (TCP) are based on the idea that packet loss is an appropriate indication of congestion on networks transmitting data using the best-effort service model. When a network uses the best-effort service model, the network delivers data if it can, without any assurance of reliability, delay bounds, or throughput. However, these algorithms and the best-effort service model are not suited to applications that are sensitive to delay or packet loss (for instance, interactive traffic including Telnet, web-browsing, and transfer of audio and video data). Weighted Random Early Detection (WRED), and by extension, Explicit Congestion Notification (ECN), helps to solve this problem.

RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*, states that with the addition of active queue management (for example, WRED) to the Internet infrastructure, routers are no longer limited to packet loss as an indication of congestion.

Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

How WRED Works

WRED makes early detection of congestion possible and provides a means for handling multiple classes of traffic. WRED can selectively discard lower priority traffic when the router begins to experience congestion and provide differentiated performance characteristics for different classes of service. It also protects against global synchronization. Global synchronization occurs as waves of congestion crest, only to be followed by periods of time during which the transmission link is not used to capacity. For these reasons, WRED is useful on any output interface or router where congestion is expected to occur.

WRED is implemented at the core routers of a network. Edge routers assign IP precedences to packets as the packets enter the network. With WRED, core routers then use these precedences to determine how to treat different types of traffic. WRED provides separate thresholds and weights for different IP precedences, enabling the network to provide different qualities of service, in regard to packet dropping, for different types of traffic. Standard traffic may be dropped more frequently than premium traffic during periods of congestion.

For more information about WRED, refer to the "Congestion Avoidance Overview" module.

ECN Extends WRED Functionality

WRED drops packets, based on the average queue length exceeding a specific threshold value, to indicate congestion. ECN is an extension to WRED in that ECN marks packets instead of dropping them when the average queue length exceeds a specific threshold value. When configured with the WRED — Explicit Congestion Notification feature, routers and end hosts would use this marking as a signal that the network is congested and slow down sending packets.

As stated in RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP,* implementing ECN requires an ECN-specific field that has two bits—the ECN-capable Transport (ECT) bit and the CE (Congestion Experienced) bit—in the IP header. The ECT bit and the CE bit can be used to make four ECN field combinations of 00 to 11. The first number is the ECT bit and the second number is the CE bit. Table 1 lists each of the ECT and CE bit combination settings in the ECN field and what the combinations indicate.

Table 1 ECN Bit Setting

ECT Bit	CE Bit	Combination Indicates
0	0	Not ECN-capable
0	1	Endpoints of the transport protocol are ECN-capable
1	0	Endpoints of the transport protocol are ECN-capable
1	1	Congestion experienced

The ECN field combination 00 indicates that a packet is not using ECN.

The ECN field combinations 01 and 10—called ECT(1) and ECT(0), respectively—are set by the data sender to indicate that the endpoints of the transport protocol are ECN-capable. Routers treat these two field combinations identically. Data senders can use either one or both of these two combinations. For more information about these two field combinations, and the implications of using one over the other, refer to RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP.*

The ECN field combination 11 indicates congestion to the endpoints. Packets arriving a full queue of a router will be dropped.

How Packets Are Treated When ECN Is Enabled

- If the number of packets in the queue is below the minimum threshold, packets are transmitted. This happens whether or not ECN is enabled, and this treatment is identical to the treatment a packet receives when WRED only is being used on the network.
- If the number of packets in the queue is between the minimum threshold and the maximum threshold, one of the following three scenarios can occur:
  - If the ECN field on the packet indicates that the endpoints are ECN-capable (that is, the ECT bit is set to 1 and the CE bit is set to 0, or the ECT bit is set to 0 and the CE bit is set to 1)—and the WRED algorithm determines that the packet should have been dropped based on the drop probability—the ECT and CE bits for the packet are changed to 1, and the packet is transmitted. This happens because ECN is enabled and the packet gets marked instead of dropped.
  - If the ECN field on the packet indicates that neither endpoint is ECN-capable (that is, the ECT bit is set to 0 and the CE bit is set to 0), the packet may be dropped based on the WRED drop probability. This is the identical treatment that a packet receives when WRED is enabled without ECN configured on the router.
  - If the ECN field on the packet indicates that the network is experiencing congestion (that is, both the ECT bit and the CE bit are set to 1), the packet is transmitted. No further marking is required.
- If the number of packets in the queue is above the maximum threshold, packets are dropped based on the drop probability. This is the identical treatment a packet receives when WRED is enabled without ECN configured on the router.

For More Information

For more information about implementing ECN and about the changes required at the routers and end hosts, refer to the following RFCs:

- RFC 2309, *Internet Performance Recommendations*
- RFC 2884, *Performance Evaluation of Explicit Congestion Notification* (*ECN) in IP Networks*
- RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*

Benefits

Improved Method for Congestion Avoidance

This feature provides an improved method for congestion avoidance by allowing the network to mark packets for transmission later, rather than dropping them from the queue. Marking the packets for transmission later accommodates applications that are sensitive to delay or packet loss and provides improved throughput and application performance.

Enhanced Queue Management

Currently, dropped packets indicate that a queue is full and the network is experiencing congestion. When a network experiences congestion, this feature allows networks to mark the IP header of a packet with a CE bit. This marking, in turn, triggers the appropriate congestion avoidance mechanism and allows the network to better manage the data queues. With this feature, ECN-capable routers and end hosts can respond to congestion before a queue overflows and packets are dropped, providing enhanced queue management.

For more information on the benefits associated with ECN, refer to RFC 2309, *Internet Performance Recommendations*.

Related Documents

- "Congestion Avoidance Overview" module
- "Applying QoS Features Using the MQC" module
- *Cisco IOS Quality of Service Solutions Command Reference*
- RFC 2309, *Internet Performance Recommendations*
- RFC 2884, *Performance Evaluation of Explicit Congestion Notification (ECN) in IP Networks*
- RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*

Supported Platforms

- Cisco 805
- Cisco 806
- Cisco 820
- Cisco 828
- Cisco 1400 series
- Cisco 1600 series
- Cisco 1751
- Cisco 2420
- Cisco 3631
- Cisco 3725
- Cisco 3745
- Cisco 7100 series
- Cisco 7200 series
- Cisco 7500
- Cisco 7700
- Cisco CVA120 series
- Cisco MC3810
- Cisco uBR7200 series
- URM (Universal Route Module)

Determining Platform Support Through Cisco Feature Navigator

Cisco IOS software is packaged in feature sets that support specific platforms. To get updated information regarding platform support for this feature, access Cisco Feature Navigator. Cisco Feature Navigator dynamically updates the list of supported platforms as new platform support is added for the feature.

Cisco Feature Navigator is a web-based tool that enables you to quickly determine which Cisco IOS software images support a specific set of features and which features are supported in a specific Cisco IOS image. You can search by feature or release. Under the release section, you can compare releases side by side to display both the features unique to each software release and the features in common.

To access Cisco Feature Navigator, you must have an account on Cisco.com. If you have forgotten or lost your account information, send a blank e-mail to cco-locksmith@cisco.com. An automatic check will verify that your e-mail address is registered with Cisco.com. If the check is successful, account details with a new random password will be e-mailed to you. Qualified users can establish an account on Cisco.com by following the directions at http://www.cisco.com/register.

Cisco Feature Navigator is updated regularly when major Cisco IOS software releases and technology releases occur. For the most current information, go to the Cisco Feature Navigator home page at the following URL:

http://www.cisco.com/go/fn

Supported Standards, MIBs, and RFCs

Standards

No new or modified standards are supported by this feature.

MIBs

No new or modified MIBs are supported by this feature.

To obtain lists of supported MIBs by platform and Cisco IOS release, and to download MIB modules, go to the Cisco MIB website on Cisco.com at the following URL:

http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

- RFC 2309, *Internet Performance Recommendations*
- RFC 2884, *Performance Evaluation of Explicit Congestion Notification (ECN) in IP Networks*
- RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*

Prerequisites

ECN must be configured through the Modular Quality of Service Command-Line Interface (MQC). For more information about the MQC, see the "Applying QoS Features Using the MQC" module.

Configuration Tasks

See the following sections for configuration tasks for the WRED — Explicit Congestion Notification feature. Each task in the list is identified as either required or optional.

- Configuring Explicit Congestion Notification (required)
- Verifying the Explicit Congestion Notification Configuration (optional)

Configuring Explicit Congestion Notification

To configure ECN, use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# policy-map` *policy-map-name*	Creates or modifies a policy map that can be attached to one or more interfaces to specify a service policy. Enters QoS policy-map configuration mode.
Step 2	`Router(config-pmap)# class class-default`	Specifies the name of the class whose policy you want to create or change or specifies the default class (commonly known as the class-default class) before you configure its policy.
Step 3	`Router(config-pmap)# bandwidth` {*bandwidth-kbps* | `percent` *percent*}	Specifies or modifies the bandwidth (either in kbps or a percentage) allocated for a class belonging to a policy map. Enters policy-map class configuration mode.
Step 4	`Router(config-pmap-c)# random-detect`	Enables WRED or distributed WRED (dWRED).
Step 5	`Router(config-pmap-c)# random-detect ecn`	Enables ECN.

Verifying the Explicit Congestion Notification Configuration

To verify the ECN configuration, use the following commands in EXEC or privileged EXEC mode, as needed:

Command	Purpose
`Router# show policy-map`	If ECN is enabled, displays ECN marking information for a specified policy map.
`Router# show policy-map interface`	If ECN is enabled, displays ECN marking information for a specified interface.

Configuration Examples

This section provides the following configuration examples:

- Enabling ECN Example
- Verifying the ECN Configuration Example

Enabling ECN Example

The following example enables ECN in the policy map called pol1:

```
Router(config)# policy-map pol1
 Router(config-pmap)# class class-default
 Router(config-pmap)# bandwidth per 70
 Router(config-pmap-c)# random-detect
 Router(config-pmap-c)# random-detect ecn
```

Verifying the ECN Configuration Example

The following is sample output from the **show policy-map** command. The words "explicit congestion notification" (along with the ECN marking information) included in the output indicate that ECN has been enabled.

```
Router# show policy-map

   Policy Map pol1
     Class class-default
       Weighted Fair Queueing
             Bandwidth 70 (%)
             exponential weight 9
             explicit congestion notification
             class    min-threshold    max-threshold    mark-probability
             ------------------------------------------------------------
             ------------------------------------------------------------

             0        -                -                1/10
             1        -                -                1/10
             2        -                -                1/10
             3        -                -                1/10
             4        -                -                1/10
             5        -                -                1/10
             6        -                -                1/10
             7        -                -                1/10
             rsvp     -                -                1/10
```

The following is sample output from the **show policy-map interface** command. The words "explicit congestion notification" included in the output indicate that ECN has been enabled.

```
Router# show policy-map interface Serial4/1

 Serial4/1

  Service-policy output:policy_ecn
        Class-map:prec1 (match-all)
          1000 packets, 125000 bytes
          30 second offered rate 14000 bps, drop rate 5000 bps
          Match:ip precedence 1
          Weighted Fair Queueing
            Output Queue:Conversation 42
            Bandwidth 20 (%)
            Bandwidth 100 (kbps)
            (pkts matched/bytes matched) 989/123625
         (depth/total drops/no-buffer drops) 0/455/0
             exponential weight:9
             explicit congestion notification
             mean queue depth:0

     class   Transmitted  Random drop  Tail drop   Minimum     Maximum     Mark
             pkts/bytes   pkts/bytes    pkts/bytes threshold   threshold   probability
       0       0/0          0/0          0/0          20          40        1/10
       1     545/68125      0/0          0/0          22          40        1/10
       2       0/0          0/0          0/0          24          40        1/10
       3       0/0          0/0          0/0          26          40        1/10
       4       0/0          0/0          0/0          28          40        1/10
       5       0/0          0/0          0/0          30          40        1/10
       6       0/0          0/0          0/0          32          40        1/10
       7       0/0          0/0          0/0          34          40        1/10
     rsvp      0/0          0/0          0/0          36          40        1/10
```

```
class   ECN Mark
        pkts/bytes
  0      0/0
  1     43/5375
  2      0/0
  3      0/0
  4      0/0
  5      0/0
  6      0/0
  7      0/0
rsvp     0/0
```

Command Reference

The following commands are introduced or modified in the feature or features documented in this module. For information about these commands, see the *Cisco IOS Quality of Service Solutions Command Reference* at http://www.cisco.com/en/US/docs/ios/qos/command/reference/qos_book.html. For information about all Cisco IOS commands, go to the Command Lookup Tool at http://tools.cisco.com/Support/CLILookup or to the *Cisco IOS Master Commands List*.

- **random-detect ecn**
- **show policy-map**
- **show policy-map interface**