



WRED-Explicit Congestion Notification

Last Updated: December 2, 2011

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.

This document describes the WRED--Explicit Congestion Notification feature in Cisco IOS Release 12.2(8)T.

- [Finding Feature Information, page 1](#)
- [Prerequisites for WRED-Explicit Congestion Notification, page 1](#)
- [Information About WRED-Explicit Congestion Notification, page 2](#)
- [How to Configure WRED-Explicit Congestion Notification, page 4](#)
- [Configuration Examples for WRED-Explicit Congestion Notification, page 7](#)
- [Additional References, page 8](#)
- [Feature Information for WRED-Explicit Congestion Notification, page 10](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the Feature Information Table at the end of this document.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for WRED-Explicit Congestion Notification



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

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.

Information About WRED-Explicit Congestion Notification

- [WRED-Explicit Congestion Notification Feature Overview, page 2](#)
- [ECN Extends WRED Functionality, page 2](#)
- [Benefits of WRED-Explicit Congestion Notification, page 4](#)

WRED-Explicit Congestion Notification 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.

- [How WRED Works, page 2](#)

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. The table below 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, page 3](#)

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 minimum 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.

Benefits of WRED-Explicit Congestion Notification

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*.

How to Configure WRED-Explicit Congestion Notification

- [Configuring Explicit Congestion Notification, page 4](#)
- [Verifying the Explicit Congestion Notification Configuration, page 6](#)

Configuring Explicit Congestion Notification

To configure ECN, complete the following steps.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **policy-map** *policy-map-name*
4. **class** { *class-name* | **class-default** }
5. **bandwidth** { *bandwidth-kbps* | **percent percent** }
6. **random-detect**
7. **random-detect ecn**
8. **end**

DETAILED STEPS

Command or Action	Purpose
<p>Step 1 enable</p> <p>Example:</p> <pre>Router> enable</pre>	<p>Enables privileged EXEC mode.</p> <ul style="list-style-type: none"> Enter your password if prompted.
<p>Step 2 configure terminal</p> <p>Example:</p> <pre>Router# configure terminal</pre>	<p>Enters global configuration mode.</p>
<p>Step 3 policy-map <i>policy-map-name</i></p> <p>Example:</p> <pre>Router(config)# policy-map policy1</pre>	<p>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.</p> <ul style="list-style-type: none"> Enter the name of the policy map.
<p>Step 4 class {<i>class-name</i> class-default}</p> <p>Example:</p> <pre>Router(config-pmap)# class class-default</pre>	<p>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. Enters policy-map-class configuration mode.</p> <ul style="list-style-type: none"> Enter the name of the class or enter the class-default keyword.
<p>Step 5 bandwidth {<i>bandwidth-kbps</i> percent <i>percent</i>}</p> <p>Example:</p> <pre>Router(config-pmap-c)# bandwidth percent 35</pre>	<p>Specifies or modifies the bandwidth (either in kbps or a percentage) allocated for a class belonging to a policy map.</p> <ul style="list-style-type: none"> Enter the bandwidth in kilobytes per second or enter the bandwidth percentage.
<p>Step 6 random-detect</p> <p>Example:</p> <pre>Router(config-pmap-c)# random-detect</pre>	<p>Enables WRED or distributed WRED (dWRED).</p>

Command or Action	Purpose
Step 7 <code>random-detect ecn</code> Example: <pre>Router(config-pmap-c)# random-detect ecn</pre>	Enables ECN.
Step 8 <code>end</code> Example: <pre>Router(config-pmap-c)# end</pre>	(Optional) Exits policy-map class configuration mode.

Verifying the Explicit Congestion Notification Configuration

To verify the ECN configuration, complete the following steps.

SUMMARY STEPS

1. `enable`
2. `show policy-map`
3. `show policy-map interface`
4. `end`

DETAILED STEPS

Command or Action	Purpose
Step 1 <code>enable</code> Example: <pre>Router> enable</pre>	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2 <code>show policy-map</code> Example: <pre>Router# show policy-map</pre>	If ECN is enabled, displays ECN marking information for a specified policy map.

Command or Action	Purpose
Step 3 show policy-map interface Example: Router# show policy-map interface	If ECN is enabled, displays ECN marking information for a specified interface.
Step 4 end Example: Router# end	(Optional) Exits privileged EXEC mode.

Configuration Examples for WRED-Explicit Congestion Notification

- [Example Enabling ECN, page 7](#)
- [Example Verifying the ECN Configuration, page 7](#)

Example Enabling ECN

The following example enables ECN in the policy map called poll1:

```
Router(config)# policy-map poll1
Router(config-pmap)# class class-default
Router(config-pmap-c)# bandwidth per 70
Router(config-pmap-c)# random-detect
Router(config-pmap-c)# random-detect ecn
```

Example Verifying the ECN Configuration

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 poll1
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:precl (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

```

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Commands List, All Releases
QoS commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco IOS Quality of Service Solutions Command Reference</i>

Related Topic	Document Title
MQC	"Applying QoS Features Using the MQC" module
Congestion avoidance concepts	"Congestion Avoidance Overview" module

Standards

Standard	Title
No new or modified standards are supported, and support for existing standards has not been modified.	--

MIBs

MIB	MIBs Link
No new or modified MIBs are supported, and support for existing MIBs has not been modified.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
RFC 2309	<i>Internet Performance Recommendation</i>
RFC 2884	<i>Performance Evaluation of Explicit Congestion Notification (ECN) in IP Networks</i>
RFC 3168	<i>The Addition of Explicit Congestion Notification (ECN) to IP</i>

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for WRED-Explicit Congestion Notification

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 2 Feature Information for WRED-Explicit Congestion Notification

Feature Name	Software Releases	Feature Configuration Information
WRED--Explicit Congestion Notification	12.2(8)T	<p>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.</p> <p>The following commands were introduced or modified: random-detect ecn, show policy-map, show policy-map interface.</p>

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks.

Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2011 Cisco Systems, Inc. All rights reserved.