Independent Submission
Request for Comments: 7514
Category: Experimental
ISSN: 2070-1721

M. Luckie
CAIDA

1 April 2015

Really Explicit Congestion Notification (RECN)

Abstract

This document proposes a new ICMP message that a router or host may use to advise a host to reduce the rate at which it sends, in cases where the host ignores other signals provided by packet loss and Explicit Congestion Notification (ECN).

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. This is a contribution to the RFC Series, independently of any other RFC stream. The RFC Editor has chosen to publish this document at its discretion and makes no statement about its value for implementation or deployment. Documents approved for publication by the RFC Editor are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7514.

Table of Contents

1. Introduction

The deployment of Explicit Congestion Notification (ECN) [RFC3168] remains stalled. While most operating systems support ECN, it is currently disabled by default because of fears that enabling ECN will break transport protocols. This document proposes a new ICMP message that a router or host may use to advise a host to reduce the rate at which it sends, in cases where the host ignores other signals such as packet loss and ECN. We call this message the "Really Explicit Congestion Notification" (RECN) message because it delivers a less subtle indication of congestion than packet loss and ECN.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. RECN Message Format

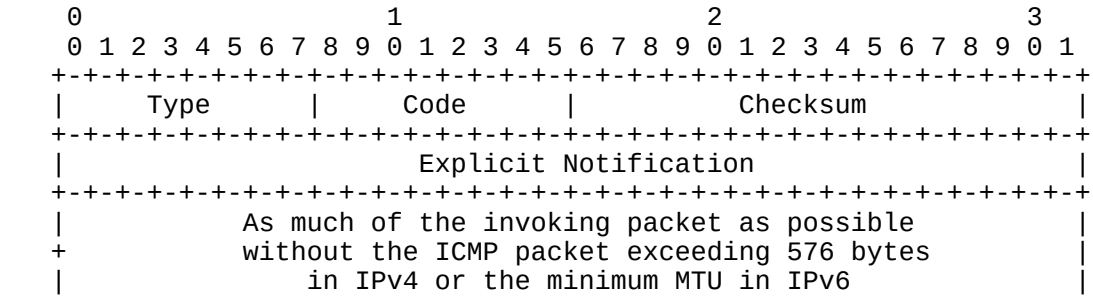

Type

IPv4: 4

IPv6: 201

Code

0

Checksum

The checksum is the 16-bit ones's complement of the one's complement sum of the ICMP message starting with the ICMP type field. When an RECN message is encapsulated in an IPv6 packet, the computation includes a "pseudo-header" of IPv6 header fields as specified in Section 8.1 of [RFC2460]. For computing the checksum, the checksum field is first set to zero.

Explicit Notification

A 4-byte value that conveys an explicit notification in the ASCII format defined in [RFC20]. This field MUST NOT be set to zero.

Description

An RECN message SHOULD be sent by a router in response to a host that is generating traffic at a rate persistently unfair to other competing flows and that has not reacted to previous packet losses or ECN marks.

The contents of an RECN message MUST be conveyed to the user responsible for the traffic. Precisely how this is accomplished will depend on the capabilities of the host. If text-to-speech capabilities are available, the contents should be converted to sound form and audibly rendered. If the system is currently muted, a pop-up message will suffice.

2.1. Advice to Implementers

As the Explicit Notification field is only 4 bytes, it is not required that the word be null terminated. A client implementation should be careful not to use more than those 4 bytes. If a router chooses a word less than 4 bytes in size, it should null-terminate that word.

A router should not necessarily send an RECN message every time it discards a packet due to congestion. Rather, a router should send these messages whenever it discards a burst of packets from a single sender. For every packet a router discards in a single burst, it should send an RECN message. A router may form short sentences composed of different 4-byte words, and a host should play these sentences back to the user. A router may escalate the content in the Explicit Notification field if it determines that a sender has not

adjusted its transmission rate in response to previous RECN messages. There is no upper bound on the intensity of the escalation, either in content or sentence length.

2.2. Relationship to ICMP Source Quench

The RECN message was inspired by the ICMP Source Quench message, which is now deprecated [RFC6633]. Because the RECN message uses a similar approach, an RECN message uses the same ICMP type when encapsulated in IPv4 as was used by the ICMP Source Quench message.

3. IANA Considerations

This is an Experimental RFC; the experiment will conclude two years after the publication of this RFC. During the experiment, implementers are free to use words of their own choosing (up to four letters) in RECN messages. If RECN becomes a Standard of any kind, a list of allowed words will be maintained in an IANA registry. There are no IANA actions required at this time.

4. Security Considerations

ICMP messages may be used in various attacks [RFC5927]. An attacker may use the RECN message to cause a host to reduce their transmission rate for no reason. To prevent such an attack, a host must ensure the quoted message corresponds to an active flow on the system, and an attacker MUST set the security flag defined in [RFC3514] to 1 when the RECN message is carried in an IPv4 packet.

5. Normative References

[RFC20] Cerf, V., "ASCII format for network interchange", STD 80, RFC 20, October 1969, <http://www.rfc-editor.org/info/rfc20>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998, <http://www.rfc-editor.org/info/rfc2460>.

[RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, September 2001, <http://www.rfc-editor.org/info/rfc3168>.

[RFC3514] Bellovin, S., "The Security Flag in the IPv4 Header", RFC 3514, April 2003, <http://www.rfc-editor.org/info/rfc3514>.

[RFC5927] Gont, F., "ICMP Attacks against TCP", RFC 5927, July 2010, <http://www.rfc-editor.org/info/rfc5927>.

[RFC6633] Gont, F., "Deprecation of ICMP Source Quench Messages", RFC 6633, May 2012, <http://www.rfc-editor.org/info/rfc6633>.

Author's Address

Matthew Luckie
CAIDA
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0505
United States

EMail: mjl@caida.org