

Transport Area Working Group (tsvwg)
Internet-Draft
Updates: [792](#), [1122](#), [1812](#)
(if approved)
Intended status: Standards Track
Expires: August 25, 2012

F. Gont
UTN-FRH / SI6 Networks
February 22, 2012

Deprecation of ICMP Source Quench messages
draft-ietf-tsvwg-source-quench-05.txt

Abstract

This document formally deprecates the use of ICMP Source Quench messages by transport protocols, formally updating [RFC 792](#), [RFC 1122](#), and [RFC 1812](#).

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 25, 2012.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	ICMP Source Quench messages	3
3.	Updating RFC 1122	4
4.	Updating RFC 1812	4
5.	Clarification for UDP, SCTP, and DCCP	5
6.	General Advice to Transport Protocols	5
7.	Recommendation Regarding RFC 1016	5
8.	Security Considerations	5
9.	IANA Considerations	6
10.	Acknowledgements	6
11.	References	7
11.1.	Normative References	7
11.2.	Informative References	7
Appendix A.	Survey of support of ICMP Source Quench in some popular TCP/IP implementations	8
Appendix B.	Changes from previous versions of the draft (to be removed by the RFC Editor before publishing this document as an RFC)	8
B.1.	Changes from draft-ietf-tsvwg-source-quench-04	8
B.2.	Changes from draft-ietf-tsvwg-source-quench-03	8
B.3.	Changes from draft-ietf-tsvwg-source-quench-02	9
B.4.	Changes from draft-ietf-tsvwg-source-quench-01	9
B.5.	Changes from draft-ietf-tsvwg-source-quench-00	9
B.6.	Changes from draft-gont-tsvwg-source-quench-01	9
B.7.	Changes from draft-gont-tsvwg-source-quench-00	9
Author's Address	9

1. Introduction

The ICMP specification [[RFC0792](#)] defined the ICMP Source Quench message (type 4, code 0), which was meant as a mechanism for congestion control. ICMP Source Quench has been known to be an ineffective (and unfair) antidote for congestion, and generation of ICMP Source Quench messages by routers has been formally deprecated by [[RFC1812](#)] since 1995. However, reaction to ICMP Source Quench messages in transport protocols has never been formally deprecated.

This document formally deprecates reaction to ICMP Source Quench messages by transport protocols such as TCP, formally updating [[RFC0792](#)], [[RFC1122](#)], and [[RFC1812](#)]. Additionally, it provides recommendation against the implementation of [[RFC1016](#)]. The rationale for these specification updates is:

- o Processing of ICMP Source Quench messages by routers has been deprecated for more than 20 years [[RFC1812](#)].
- o Virtually all popular host implementations have removed support for ICMP Source Quench messages since (at least) 2005 [[RFC5927](#)].
- o Widespread deployment of ICMP filtering makes it impossible to rely on ICMP Source Quench messages for congestion control.
- o The IETF has moved away from ICMP Source Quench messages for congestion control (note e.g. the development of ECN [[RFC3168](#)], and the fact that ICMPv6 [[RFC4443](#)] does not even specify a Source Quench message).

ICMP Source Quench messages are not normally seen in the deployed Internet and were considered rare at least as far back as 1994. [[Floyd1994](#)]

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. ICMP Source Quench messages

The ICMP specification [[RFC0792](#)] defined the ICMP Source Quench message (type 4, code 0), which was meant to provide a mechanism for congestion control. The Host Requirements RFC [[RFC1122](#)] stated in [Section 4.2.3.9](#) that hosts MUST react to ICMP Source Quench messages by slowing transmission on the connection, and further added that the RECOMMENDED procedure was to put the corresponding connection in the slow-start phase of TCP's congestion control algorithm [[RFC5681](#)].

[RFC1812] noted that research suggested that ICMP Source Quench was an ineffective (and unfair) antidote for congestion, and formally deprecated the generation of ICMP Source Quench messages by routers, stating that routers SHOULD NOT send ICMP Source Quench messages in response to congestion.

[RFC5927] discussed the use of ICMP Source Quench messages for performing "blind throughput-reduction" attacks, and noted that most TCP implementations silently ignore ICMP Source Quench messages.

We note that TCP implements its own congestion control mechanisms [RFC5681] [RFC3168], that do not depend on ICMP Source Quench messages.

It is interesting to note that ICMPv6 [RFC4443] does not specify a "Source Quench" message.

3. Updating [RFC 1122](#)

This document hereby updates [Section 3.2.2.3 of \[RFC1122\]](#) as follows:

A host MUST NOT send ICMP Source Quench messages.

If a Source Quench message is received, the IP layer MAY silently discard it.

[Section 4.2.3.9 of \[RFC1122\]](#) is updated as follows:

TCP MUST silently discard any received ICMP Source Quench messages.

The consensus of the TSV WG was that there are no valid reasons for a host to generate or react to an ICMP Source Quench message in the current Internet. The recommendation that a sender "MUST NOT" send an ICMP Source Quench message is because there is no known valid reason for a host to generate this message. The only known impact of a sender ignoring this requirement is that it may necessarily consume network and endpoint resources. Discarding ICMP Source Quench messages at the internet-layer (rather than at the transport layer) is a performance optimization that is permitted by this update.

4. Updating [RFC 1812](#)

This document hereby updates [Section 4.3.3.3 of \[RFC1812\]](#) as follows:

A router MUST ignore any ICMP Source Quench messages it receives.

The consensus of the TSV WG was that there are no valid reasons for a router to react to ICMP Source Quench messages in the current Internet.

5. Clarification for UDP, SCTP, and DCCP

UDP did not explicitly specify support for ICMP Source Quench messages. Hereby we clarify that UDP end-points MUST silently discard received ICMP Source Quench messages.

It is understood that SCTP and DCCP did not specify support for processing received ICMP Source Quench messages. Hereby we clarify that DCCP and SCTP end-points MUST silently discard received ICMP Source Quench messages.

6. General Advice to Transport Protocols

If a Source Quench message is received by any other transport-protocol instance, it MUST be silently ignored.

The TSV WG is not aware of any use that requires processing of these messages, and therefore expects other transports to follow the recommendations in [Section 3](#). Note that for IETF-specified transports, this document formally deprecates reaction to ICMP Source Quench messages, and that generation of ICMP Source Quench messages has been deprecated for both hosts and routers. Therefore, future applications can not expect to receive these messages.

7. Recommendation Regarding [RFC 1016](#)

[RFC 1016](#) [[RFC1016](#)] described an experimental approach to ICMP Source Quench message handling in hosts that was being thought about in 1987. The IETF notes that [RFC 1016](#) has never been on the IETF standards-track, but for clarity and avoidance of doubt, the note that the approach described in [RFC 1016](#) [[RFC1016](#)] MUST NOT be implemented.

8. Security Considerations

ICMP Source Quench messages could be leveraged for performing blind throughput-reduction attacks against TCP and similar protocols. This attack vector, along with possible countermeasures, has been

discussed in great detail in [[RFC5927](#)] and [[CPNI-TCP](#)]. Silently ignoring ICMP Source Quench messages, as specified in this document, eliminates the aforementioned attack vector.

For current TCP implementations, receipt of an ICMP Source Quench message should not result in security issues because, as noted in [[RFC5927](#)] and [[CPNI-TCP](#)], virtually all current versions of popular TCP implementations already silently ignore ICMP Source Quench messages. This is also the case for SCTP and DCCP implementations.

Hosts, security gateways, and firewalls MUST silently discard received ICMP Source Quench packets and SHOULD log such drops as a security fault with at least minimal details (IP Source Address, IP Destination Address, ICMP message type, and date/time the packet was seen).

We note that security devices such as the Snort Network Intrusion Detection System (NIDS) has logged ICMP Source Quench messages as such for more than ten years. [[Anderson2002](#)].

9. IANA Considerations

IANA is requested to mark ICMP type 4 (Source Quench) as "Deprecated" in the ICMP Parameters registry [[ICMPPARREG](#)] with a reference to this document.

10. Acknowledgements

The author of this document would like to thank Ran Atkinson, who contributed text that was incorporated into this document and also provided valuable feedback on earlier versions of this document.

The author of this document would like to thank (in alphabetical order) Fred Baker, David Black, Scott Bradner, James Carlson, Antonio De Simone, Wesley Eddy, Gorrry Fairhurst, Alfred Hoenes, Mahesh Jethanandani, Kathleen Moriarty, Carlos Pignataro, James Polk, Anantha Ramaiah, Randall Stewart, Dan Wing, and Andrew Yourtchenko, for providing valuable feedback on earlier versions of this document.

This document has benefited from discussions within the TCPM Working Group while working on [[RFC5927](#)].

11. References

11.1. Normative References

- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, [RFC 792](#), September 1981.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.
- [RFC1122] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, [RFC 1122](#), October 1989.
- [RFC1812] Baker, F., "Requirements for IP Version 4 Routers", [RFC 1812](#), June 1995.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", [RFC 5681](#), September 2009.

11.2. Informative References

- [Anderson2002]
Anderson, D., Fong, M., and A. Valdes, "Heterogeneous Sensor Correlation: A Case Study of Live Traffic Analysis", Proceedings of the 3rd Annual IEEE Information Assurance Workshop New York, NY, USA, 2002.
- [CPNI-TCP]
CPNI, "Security Assessment of the Transmission Control Protocol (TCP)", 2009, <<http://www.cpni.gov.uk/Docs/tn-03-09-security-assessment-TCP.pdf>>.
- [Floyd1994]
Floyd, S., "TCP and Explicit Congestion Notification", ACM CCR Volume 24, Issue 5, 1994.
- [FreeBSD] The FreeBSD Project, "<http://www.freebsd.org>".
- [ICMPPARREG]
Internet Control Message Protocol (ICMP) Parameters, "<http://www.iana.org/assignments/icmp-parameters>".
- [Linux] The Linux Project, "<http://www.kernel.org>".
- [NetBSD] The NetBSD Project, "<http://www.netbsd.org>".
- [OpenBSD] The OpenBSD Project, "<http://www.openbsd.org>".

[OpenSolaris]

OpenSolaris, "<http://www.opensolaris.org>".

[RFC1016] Prue, W. and J. Postel, "Something a host could do with source quench: The Source Quench Introduced Delay (SQuID)", [RFC 1016](#), July 1987.

[RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", [RFC 3168](#), September 2001.

[RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC 4443](#), March 2006.

[RFC5927] Gont, F., "ICMP Attacks against TCP", [RFC 5927](#), July 2010.

[Appendix A](#). Survey of support of ICMP Source Quench in some popular TCP/IP implementations

A large number of implementations completely ignore ICMP Source Quench messages meant for TCP connections. This behavior has been implemented in, at least, Linux [[Linux](#)] since 2004, and in FreeBSD [[FreeBSD](#)], NetBSD [[NetBSD](#)], OpenBSD [[OpenBSD](#)], and Solaris 10 since 2005. Additionally, OpenSolaris [[OpenSolaris](#)] has always shipped with support for ICMP Source Quench messages disabled.

[Appendix B](#). Changes from previous versions of the draft (to be removed by the RFC Editor before publishing this document as an RFC)

[B.1](#). Changes from [draft-ietf-tsvwg-source-quench-04](#)

- o Removes request to move [RFC 1016](#) to "Historic" status.
- o Updates the Security Considerations section.

[B.2](#). Changes from [draft-ietf-tsvwg-source-quench-03](#)

- o Added 'Obsoletes' metadata, and moved the reference to [[RFC1016](#)] from the 'Normative References' to the 'Informative References'.

B.3. Changes from [draft-ietf-tsvwg-source-quench-02](#)

- o Clarifies the requirements language.

B.4. Changes from [draft-ietf-tsvwg-source-quench-01](#)

- o Changes deprecation of ICMP SQ from "SHOULD NOT" to "MUST NOT" in response of feedback from Scott Bradner and the TSV WG.

B.5. Changes from [draft-ietf-tsvwg-source-quench-00](#)

- o Discusses the motivation for deprecating ICMP Source Quench messages (as suggested by Anantha Ramaiah).
- o Incorporates IANA considerations such that ICMP Source Quench messages are deprecated in the corresponding registry.

B.6. Changes from [draft-gont-tsvwg-source-quench-01](#)

- o Addresses nits and editorial changes suggested by Gorrry Fairhurst.
- o Added the status of Solaris and OpenSolaris to [Appendix A](#).
- o Document resubmitted as [draft-ietf](#).

B.7. Changes from [draft-gont-tsvwg-source-quench-00](#)

- o This revision reflects the recent discussion about ICMP Source Quench messages on the tsvwg mailing-list. A detailed list of the changes is available at:
<http://www.ietf.org/mail-archive/web/tsvwg/current/msg10407.html>

Author's Address

Fernando Gont
UTN-FRH / SI6 Networks
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Phone: +54 11 4650 8472
Email: fgont@si6networks.com
URI: <http://www.si6networks.com>

