Network Working Group
Internet-Draft
Intended status: BCP
Expires: August 1, 2014

M. Shore
No Mountain Software
C. Pignataro
Cisco Systems, Inc.
January 28, 2014

An Acceptable Use Policy for New ICMP Types and Codes
draft-shore-icmp-aup-10

Abstract

In this document we provide a basic description of ICMP's role in the IP stack and some guidelines for future use.

This document is motivated by concerns about lack of clarity concerning when to add new Internet Control Message Protocol (ICMP) types and/or codes. These concerns have highlighted a need to describe policies for when adding new features to ICMP is desirable and when it is not.

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 1, 2014.

Copyright Notice

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

There has been some recent concern expressed about a lack of clarity around when to add new message types and codes to ICMP (including ICMPv4 [RFC0792] and ICMPv6 [RFC4443]). We lay out a description of when (and when not) to move functionality into ICMP.

This document is the result of discussions among ICMP experts within the OPS area's IP Diagnostics Technical Interest Group [1] and concerns expressed by the OPS area leadership.

Note that this document does not supercede the IANA Allocation Guidelines for Values in the Internet Protocol and Related Headers, RFC 2780 [RFC2780], which specifies best practices and processes for the allocation of values in the IANA registries but does not describe the policies to be applied in the standards process.

2. Acceptable use policy

In this document we describe an acceptable use policy for new ICMP message types and codes, and provide some background behind the policy.

In summary, any future message types added to ICMP should be limited to two broad categories:

1. to inform a datagram's originator that a forwarding plane anomaly has been encountered downstream. The datagram originator must be able to determine whether or not the datagram was discarded by examining the ICMP message
2. to discover and convey dynamic information about a node (other than information usually carried in routing protocols), to discover and convey network-specific parameters, and to discover on-link routers and hosts.

Normally, other uses such as implementing a general-purpose routing or network management protocol are not advisable. However, ICMP does have a role to play in conveying dynamic information about a network, which would belong in category 2 above.

2.1. Classification of existing message types

This section provides a rough breakdown of existing message types according to the taxonomy described in Section 2 at the time of publication.

IPv4 forwarding plane anomaly reporting:

3: Destination unreachable
4: Source quench (deprecated)
6: Alternate host address (deprecated)
11: Time exceeded
12: Parameter problem
31: Datagram conversion error (deprecated)
41: ICMP messages utilized by experimental mobility protocols, such as Seamoby

IPv4 router or host discovery:

0: Echo reply
5: Redirect
8: Echo
9: Router advertisement
10: Router solicitation
13: Timestamp
14: Timestamp reply
15: Information request (deprecated)
16: Information reply (deprecated)
17: Address mask request (deprecated)
18: Address mask reply (deprecated)
30: Traceroute (deprecated)
32: Mobile host redirect (deprecated)
33: IPv6 Where-Are-You (deprecated)
34: IPv6 I-Am-Here (deprecated)
35: Mobile registration request (deprecated)
36: Mobile registration reply (deprecated)

37: Domain name request (deprecated)
38: Domain name reply (deprecated)
39: SKIP (deprecated)
40: Photuris
41: ICMP messages utilized by experimental mobility protocols, such as Seamoby

Please note that some ICMP message types were formally deprecated by [RFC6918].

IPv6 forwarding plane anomaly reporting:

1: Destination unreachable
2: Packet too big
3: Time exceeded

4: Parameter problem
150: ICMP messages utilized by experimental mobility protocols, such as Seamoby

IPv6 router or host discovery:

128: Echo request
129: Echo reply
130: Multicast listener query
131: Multicast listener report
132: Multicast listener done
133: Router solicitation
134: Router advertisement
135: Neighbor solicitation
136: Neighbor advertisement
137: Redirect message
138: Router renumbering
139: ICMP node information query
140: ICMP node information response
141: Inverse neighbor discovery solicitation message
142: Inverse neighbor discovery advertisement message
143: Version 2 multicast listener report

144: Home agent address discovery request message
145: Home agent address discovery reply message
146: Mobile prefix solicitation
147: Mobile prefix advertisement
148: Certification path solicitation message
149: Certification path advertisement message
150: ICMP messages utilized by experimental mobility protocols, such as Seamoby
151: Multicast router advertisement
152: Multicast router solicitation
153: Multicast router termination
154: FMIPv6 messages
155: RPL control message

2.1.1. ICMP Use as a Routing Protocol

As mentioned in Section 2, using ICMP as a general-purpose routing or network management protocol is not advisable.

ICMP has a role in the Internet as an integral part of the IP layer, and not as a transport protocol for other layers including routing information. From a more pragmatic perspective, some of the key characteristics of ICMP do not support using it as a routing protocol. Those include that ICMP is frequently filtered, is not authenticated, and is easily spoofed.

2.1.2. A few notes on RPL

RPL, the IPv6 Routing protocol for low-power and lossy networks (see [RFC6550]) uses ICMP as a transport. It is something of an outlier among the existing ICMP message types, as the expansion of its acronym appears to be an actual routing protocol.

This should be considered anomalous and is not a model for future ICMP message types. That is, ICMP is not intended as a transport for other protocols and should not be used in that way in future specifications. In particular, while it is adequate to use ICMP as a discovery protocol, this does not extend to full routing capabilities.

2.2. Applications using ICMP

Some applications make use of ICMP error notifications, or even deliberately create anomalous conditions in order to elicit ICMP messages, to then use those ICMP messages to generate feedback to the higher layer. Some of these applications include most widespread examples such as TRACEROUTE and Path MTU Discovery (PMTUD). These uses are considered acceptable as they do not add new message types to ICMP.

2.3. Extending ICMP

ICMP multi-part messages are specified in [RFC4884] by defining an extension mechanism for selected ICMP messages. This mechanism addresses a fundamental problem in ICMP extensibility. An ICMP multi-part message carries all of the information that ICMP messages carried previously, as well as additional information that applications may require.

Some currently defined ICMP extensions include ICMP extensions for Multiprotocol Label Switching [RFC4950] and ICMP extensions for interface and next-hop identification [RFC5837].

Extensions to ICMP should follow [RFC4884].

2.4. ICMPv4 vs. ICMPv6

Because ICMPv6 is used for IPv6 Neighbor Discovery, deployed IPv6 routers, IPv6-capable security gateways, and IPv6-capable firewalls normally support administrator configuration of how specific ICMPv6 message types are handled. By contrast, deployed IPv4 routers, IPv4-capable security gateways, and IPv4-capable firewalls are less likely to allow an administrator to configure how specific ICMPv4 message types are handled. So, at present, ICMPv6 messages usually have a

higher probability of travelling end-to-end than ICMPv4 messages.

3. ICMP's role in the internet

ICMP was originally intended to be a mechanism for gateways or destination hosts to report error conditions back to source hosts in ICMPv4 [RFC0792], and ICMPv6 [RFC4443] is modeled after it. The word

"control" in the protocol name did not describe ICMP's function (i.e. it did not "control" the internet), but rather that it was used to communicate about the control functions in the internet. For example, even though ICMP included a redirect message type that affects routing behavior in the context of a LAN segment, it was and is not used as a generic routing protocol.

ICMP is defined to be an integral part of IP, and must be implemented by every IP module. This is true for ICMPv4 as an integral part of IPv4 (see the Introduction of [RFC0792]), and for ICMPv6 as an integral part of IPv6 (see Section 2 of [RFC4443]). When first defined, ICMP messages were thought of as IP messages that didn't carry any higher layer data. It could be conjectured that the term 'control' was used given that ICMP messages were not 'data' messages.

Most likely because of the presence of the word "control" in the protocol name, ICMP is often understood to be a control protocol, borrowing some terminology from circuit networks and the PSTN. That is probably not correct - it might be more correct to describe it as being closer to a management plane protocol, given the data plane/ control plane/ management plane taxonomy often used in describing telephony protocols. However, layering in IP networks is not very clean and there's often some intermingling of function that can tend to lead to confusion about where to place new functions.

In the following section we provide some background on the differences between control and management traffic.

4. Security considerations

This document describes a high-level policy for adding ICMP types and codes. While special attention must be paid to the security implications of any particular new ICMP type or code, this recommendation presents no new security considerations.

From a security perspective, ICMP plays a part in the Photuris [RFC2521] protocol. But more generally, ICMP is not a secure protocol, and does not include features to be used to discover network security parameters or to report on network security

anomalies in the forwarding plane.

Additionally, new ICMP functionality (e.g., ICMP extensions, or new ICMP types or codes) needs to consider potential ways of how ICMP can be abused (e.g., Smurf IP DoS [CA-1998-01]).

5. IANA considerations

There are no actions required by IANA.

6. Acknowledgments

This document was originally proposed by, and received substantial review and suggestions from, Ron Bonica. Discussions with Pascal Thubert helped clarify the history of RPL's use of ICMP. We are very grateful for the review, feedback, and comments from Ran Atkinson, Tim Chown, Joe Clarke, Ray Hunter, Hilarie Orman, Eric Rosen, JINMEI Tatuya, and Wen Zhang, which resulted in a much improved document.

7. Informative references

[RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.

[RFC2780] Bradner, S. and V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", BCP 37, RFC 2780, March 2000.

[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.

[RFC6550] Winter, T., Thubert, P., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, March 2012.

[RFC6918] Gont, F. and C. Pignataro, "Formally Deprecating Some ICMPv4 Message Types", RFC 6918, April 2013.

[RFC4884] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, "Extended ICMP to Support Multi-Part Messages", RFC 4884, April 2007.

[RFC4950] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, "ICMP

Extensions for Multiprotocol Label Switching", RFC 4950, August 2007.

[RFC5837] Atlas, A., Bonica, R., Pignataro, C., Shen, N., and JR. Rivers, "Extending ICMP for Interface and Next-Hop Identification", RFC 5837, April 2010.

[RFC2521] Karn, P. and W. Simpson, "ICMP Security Failures Messages", RFC 2521, March 1999.

[CA-1998-01] CERT, "Smurf IP Denial-of-Service Attacks", CERT Advisory CA-1998-01, January 1998, <http://www.cert.org/advisories/CA-1998-01.html>.

[1] <https://svn.tools.ietf.org/area/ops/trac/wiki/TIG_DIAGNOSTICS>

Authors' Addresses

Melinda Shore
No Mountain Software
PO Box 16271
Two Rivers, AK 99716
US

Phone: +1 907 322 9522
Email: melinda.shore@nomountain.net

Carlos Pignataro
Cisco Systems, Inc.
7200-12 Kit Creek Road
Research Triangle Park, NC 27709
US

Email: cpignata@cisco.com