

Network Working Group
Internet-Draft
Intended status: BCP
Expires: April 24, 2014

M. Shore
No Mountain Software
C. Pignataro
Cisco Systems, Inc.
October 21, 2013

An Acceptable Use Policy for New ICMP Types and Codes draft-shore-icmp-aup-05

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 April 24, 2014.

Copyright Notice

Copyright (c) 2013 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.	Acceptable use policy	3
2.1.	Classification of existing message types	3
2.1.1.	A few notes on RPL	5
2.2.	Extending ICMP	6
3.	ICMP's role in the internet	6
4.	Management vs. control	7
5.	Security considerations	8
6.	IANA considerations	8
7.	Acknowledgments	8
8.	Informative references	8
	Authors' Addresses	9

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 attempt to 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 2880](#). [[RFC2780](#)] That document 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 a proposed acceptable use policy for new ICMP message types and codes, and provide some background behind the proposed policy.

In summary, we propose that 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 on-link routers and hosts, and network-specific parameters

While ICMP's role is not to be a general-purpose network management protocol, it does have a role to play in conveying dynamic information about a network.

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)
- 5: Redirect
- 6: Alternate host address (deprecated)
- 11: Time exceeded
- 12: Parameter problem
- 31: Datagram conversion error (deprecated)
- 32: Mobile host redirect (deprecated)
- 41: ICMP messages utilized by experimental mobility protocols,
such as Seamoby

IPv4 router or host discovery:

- 0: Echo reply
- 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)
- 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
- 137: Redirect message
- 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
- 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. A few notes on RPL

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

This should be considered anomalous and is not a model for future ICMP message types. Our understanding is that the working group initially defined a discovery protocol extending existing ICMPv6 Neighbor Discovery messages before moving to its own native ICMP type.

It is typically the case that routing protocols have transport requirements that are not met by ICMP. For example, there will be reliability guarantees and security guarantees that are not provided by ICMP, forcing protocol developers to design their own mechanisms. Given the availability of other IETF standard transports for routing, this reinvention should be avoided.

2.2. 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](#)].

3. ICMP's role in the internet

ICMP was originally intended to be a mechanism for routers to report error conditions back to hosts in ICMPv4 [[RFC0792](#)], and ICMPv6 [[RFC4443](#)] models 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.

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

4. Management vs. control

In this section we attempt to draw a distinction between management and control planes, acknowledging in advance that this may serve to muddle the differences even further. Ultimately the difference may not matter that much for the purpose of creating a policy for adding new types to ICMP, but because that terminology has become ubiquitous, even in IETF discussions, and because it has come up in prior discussions of ICMP policies, it seems worthwhile to take a few paragraphs to describe what they are and what they are not.

The terms "management plane" and "control plane" came into use to describe one aspect of layering in telecommunications networks. "Management plane" is described in [[I-D.ietf-opsawg-oam-overview](#)], and "control plane" is defined in [[RFC6192](#)].

It is particularly important, in the context of this discussion, to understand that "control plane" in telecomm networks almost always refers to 'signaling,' or call control and network control information. This includes "call" establishment and teardown, route establishment and teardown, requesting QoS or other parameters, and so on.

"Management," on the other hand, involves an exchange between a management application and managed entities such as network nodes, and includes "inline management" and "management" per se. Typical "inline management" functions include fault management and performance monitoring (Service Level Agreement (SLA) compliance), discovery, and typical "management" include protocols such as SNMP and NETCONF.

The correct answer to the question of where ICMP fits into the management/control/data taxonomy is that it doesn't, at least not neatly. While some of the message types are unambiguously management message, at least within the narrow confines of a management/control dichotomy (ICMP type 3, or "unreachable" messages), others are less clearly identifiable. For example, the "redirect" (ICMP type 5) message can be construed to contain control (in this case, routing) information, even though it is in some very real sense an error message.

At this time,

- o there are many, many other protocols that can be (and are) used for control traffic, whether they're routing protocols, telephony signaling protocols, QoS protocols, middlebox protocols, AAA protocols, etc.

- o the transport characteristics needed by control traffic can be incompatible with the ICMP protocol standard -- for example, they may require reliable delivery, very large payloads, or have security requirements that cannot be met.
- and because of this we propose that any future message types added to ICMP should conform to the policy proposed in [Section 2](#). ICMP should not be used as a routing or network management protocol.

5. Security considerations

This document attempts to describe 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 security considerations.

6. IANA considerations

There are no actions required by IANA.

7. 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 grateful for feedback from Joe Clarke and Wen Zhang.

8. 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.
- [RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", [RFC 6192](#), March 2011.
- [I-D.ietf-opsawg-oam-overview]
Mizrahi, T., Sprecher, N., Bellagamba, E., and Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Mechanisms", [draft-ietf-opsawg-oam-overview-09](#) (work in progress), July 2013.
- [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

