

Network Working Group
Internet-Draft
Expires: January 31, 2005

F. Gont
UTN/FRH
August 2, 2004

Increasing the payload of ICMP error messages
draft-gont-icmp-payload-00.txt

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of section 3 of RFC 3667. By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she become aware will be disclosed, in accordance with RFC 3668. This document may not be modified, and derivative works of it may not be created, except to publish it as an RFC and to translate it into languages other than English.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 31, 2005.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

The original ICMP specification states that when a packet elicits an ICMP error message, the IP header plus the next 64 bits of the original datagram must be returned in the payload of the ICMP error message. This imposes a constraint on the design of transport-layer protocols, which are forced to include all the relevant information needed to identify an instance of communication in the first 64 bits

of their protocol header. It also limits the amount of data from the original packet available to the transport-layer when acting on the ICMP error message. Including only the first 64 bits of the original datagram's payload may also not be enough to demultiplex ICMP error messages if IP is being used to tunnel some other network-layer protocol. This document proposes to increase the amount of data of the original datagram to be included in the payload of ICMP error messages.

1. Introduction

The Internet Control Message Protocol (ICMP) [1] is used in the Internet Architecture to perform the fault isolation function, that is, the group of actions that hosts and routers take to determine that there is some network failure [4].

The original ICMP specification [1] states that, whenever a packet elicits an ICMP error message, the internet header plus the first 64 bits of the original datagram's data must be included in the payload of the ICMP error message. These data are used by the receiving host to match the error message to the instance of communication that elicited it.

This limit on the amount information returned in the payload of ICMP error messages has two drawbacks:

- o It imposes a constraint on the design of transport-layer protocols, which are forced to include all the relevant information needed to identify a communication instance in the first 64 bits of their protocol header.
- o It limits the amount of data the transport-protocol has available to perform, for example, security checks on the returned datagram.
- o If IP [5] is being used for tunneling purposes, including just the first 8 bytes of the payload of the original datagram may not be enough information to demultiplex the ICMP error message.

As discussed in [1] and [6], in order to allow ICMP error messages to be demultiplexed, transport protocols are forced to include in the first 64 bits of their headers all the information needed to identify a communication instance. Thus, this limit somehow constrains the design of transport protocols.

There are a number of scenarios in which a larger amount of data from the original datagram may be needed, or, at least, desirable. For example, additional data from the original datagram could be used to perform security checks on the received ICMP error message [7].

Also, in case IP is being used to tunnel some other protocol, the first 64 bits of the original datagrams's payload may not provide enough information to the demultiplex the ICMP error message.

Even when the Host Requirements RFC [2] states that more than 8 octets of the original datagram's payload MAY be included in the payload of an ICMP error message, it does not require any specific amount of data, and thus does not remove the constraints discussed above.

This document proposes a modification to the original ICMP specification to increase the amount of data of the original packet to be included in the payload of ICMP error messages.

2. Conventions

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 [3].

3. Specification

When a host or router sends an ICMP error message, it MUST include in the payload of the ICMP error message as many bytes of the original datagram as possible. However, the resulting IP datagram MUST NOT be greater than 576 bytes.

It must be noted that 576 is the minimum reassembly buffer size [2].

4. Security Considerations

This document proposes a minor modification to the original ICMP specification [1], to increase the amount of data of the original packet to be included in the payload of ICMP error messages. This modification does not raise any new security implications.

5. Acknowledgements

The author would like to thank Guillermo Gont and Michael Kerrisk for providing many valuable comments.

6. References

6.1 Normative References

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

- [2] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, October 1989.
- [3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

6.2 Informative References

- [4] Clark, D., "Fault isolation and recovery", RFC 816, July 1982.
- [5] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [6] Clark, D., "Name, addresses, ports, and routes", RFC 814, July 1982.
- [7] Gont, F., "ICMP attacks against TCP", (work in progress) draft-gont-tcpm-icmp-attacks-00.txt, 2004.

Author's Address

Fernando Gont
Universidad Tecnologica Nacional
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Phone: +54 11 4650 8472
EMail: fernando@gont.com.ar

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

