

opsec
Internet-Draft
Intended status: Best Current Practice
Expires: January 4, 2018

F. Gont
UTN-FRH / SI6 Networks
R. Hunter
Globis Consulting BV
J. Massar
Massar Networking
W. Liu
Huawei Technologies
July 3, 2017

Defeating Attacks which employ Forged ICMPv4/ICMPv6 Error Messages
draft-gont-opsec-icmp-ingress-filtering-03.txt

Abstract

Over the years, a number of attack vectors that employ forged ICMPv4/ICMPv6 error messages have been disclosed and exploited in the wild. The aforementioned attack vectors do not require that the source address of the packets be forged, but do require that the addresses of the IPv4/IPv6 packet embedded in the ICMPv4/ICMPv6 payload be forged. This document discusses a simple, effective, and straightforward method for using ingress traffic filtering to mitigate attacks that use forged addresses in the IPv4/IPv6 packet embedded in an ICMPv4/ICMPv6 payload.

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 January 4, 2018.

Copyright Notice

Copyright (c) 2017 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	2
2. Terminology	3
3. Applicability Statement	3
4. Overview	3
4.1. Generation of ICMP Error Messages in Legitimate Scenarios	4
4.2. Attack Scenario	5
5. ICMPv4/ICMPv6 Network Ingress Filtering	7
6. IANA Considerations	7
7. Security Considerations	7
8. Acknowledgements	8
9. References	8
9.1. Normative References	8
9.2. Informative References	9
Authors' Addresses	9

1. Introduction

Over the years, a number of attack vectors that employ forged ICMPv4/ICMPv6 error messages have been disclosed and exploited in the wild. The effects of these attack vectors have ranged from Denial of Service (DoS) to performance degradation [US-CERT] [RFC5927] [I-D.gont-v6ops-ipv6-ehs-packet-drops].

The aforementioned attack vectors do not require that the Source Address of the ICMP [RFC0792] or ICMPv6 [RFC4443] attack packets to be forged, but do require that the Destination Address of the IPv4 [RFC0791] (in the case of ICMPv4) or IPv6 (in the case of ICMPv6) packet embedded in the ICMPv4/ICMPv6 payload be forged. Thus, performing ingress filtering (ala BCP38 [RFC2827]) on the Destination Address of the embedded IPv4/IPv6 packet results in a simple, effective, and straightforward mitigation for any attack vectors based on ICMPv4/ICMPv6 error messages.

Section 4 provides an overview of how ICMP/ICMPv6 error messages are generated, and how packets are crafted to perform attacks based on

ICMPv4/ICMPv6 error messages. Section 5 specifies network ingress filtering based on the ICMP/ICMPv6 payload.

2. Terminology

Throughout this document the term "IP" is employed to refer to both the IPv4 [RFC0791] and IPv6 [RFC2460] protocols. That is, the term "IP" is employed when we do not mean to make a distinction between both versions of the protocol. In a similar vein, the term "ICMP" is employed to refer to both the ICMPv4 [RFC0792] and ICMPv6 [RFC4443] protocols. That is, the term "ICMP" is employed when we do not mean to make a distinction between both versions of the protocol.

For obvious reasons, ICMPv4 will only be employed in conjunction with IPv4, and ICMPv6 will always be employed in conjunction with IPv6. That is, the phrase "the IP packet embedded in the ICMP payload" means "the IPv4 packet embedded in the ICMPv4 payload" payload or "the IPv6 packet embedded in the ICMPv6 payload" (but NOT e.g. "the IPv4 packet embedded in the ICMPv6 payload").

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 [RFC2119].

3. Applicability Statement

The filtering policy specified in this document could be enforced at the border firewall of a non-multihomed network or at a CPE router, such that users of that network are prevented from performing ICMP-based attacks against other parties.

The filtering policy specified in this document SHOULD NOT be enforced in multihoming scenarios, or other scenarios where this policy could lead to false positives and therefore incorrect packet drops.

4. Overview

Attack vectors based on ICMP error messages have been known for a long time, and have been described in detail in [RFC5927]. The following subsections provide an overview of how ICMP error messages are generated in legitimate scenarios, and how an attacker would forge an ICMP error message in order to perform an attack based on ICMP error messages.

4.1. Generation of ICMP Error Messages in Legitimate Scenarios

The following figure illustrates a very simple network scenario in which two hosts (H1 and H2) are connected to each other by means of the router R1:

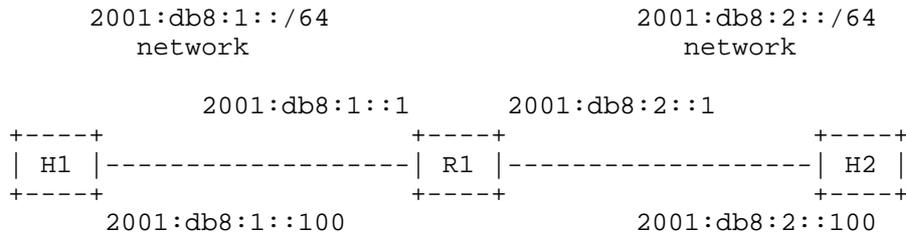


Figure 1: Sample Scenario for ICMP/ICMPv6 Error Generation

The aforementioned figure illustrates the IPv6 addresses assigned to each of the involved network interfaces. For simplicity sake, this figure employs only IPv6 addresses, but the same logic applies to the IPv4 case.

Let us assume that H1 sends a packet towards H2, and that R1 encounters an error condition while processing such a packet. Typically, the error condition will be reported to H1 by means of an ICMPv6 error message. The error message will have the following structure:

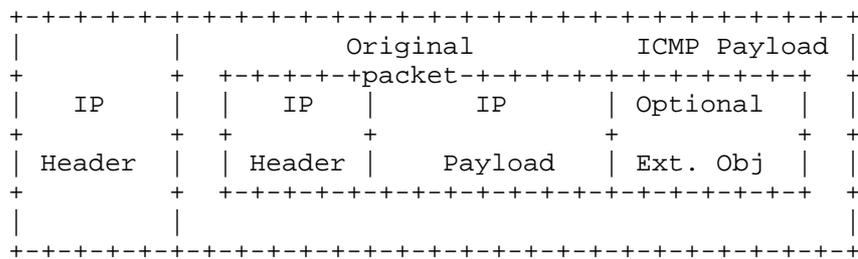


Figure 2: Structure of ICMPv4/ICMPv6 Error Messages

NOTES:

For completeness-sake, the figure above depicts the structure of ICMP error messages including ICMP extension objects (see [RFC4884]). Use of such extension objects does not affect the discussion in this document.

In the IPv6 case, the "IP header" corresponds to the entire IPv6 header chain. Additionally, in the IPv4 scenarios in which

Network Address Translation (NAT) is in place, the NAT device could fail to translate the IPv4 addresses of the embedded packet.

where the ICMPv6 error message embeds the whole (or part of) the original packet that elicited the error message.

In our scenario, the relevant header fields would have the following values:

- o Source Address: 2001:db8:1::1
- o Destination Address: 2001:db8:1::100
- o Source Address (embedded packet): 2001:db8:1::100
- o Destination Address (embedded packet): 2001:db8:2::100

It should be clear that the Source Address of the packet could be virtually any address (since it corresponds to the IP address of a router reporting the error), while the Destination Address of the packet will be that of the target/destination of the ICMP error message. On the other hand, the IP addresses of the embedded packet will be those of the packet that elicited the ICMP error message.

The embedded IP packet is typically employed by the receiving system to demultiplex the ICMP error message.

4.2. Attack Scenario

The following figure illustrates a very simple attack scenario in which an attacker (H3) tries to perform an attack against H1, while H1 is communicating with H2:

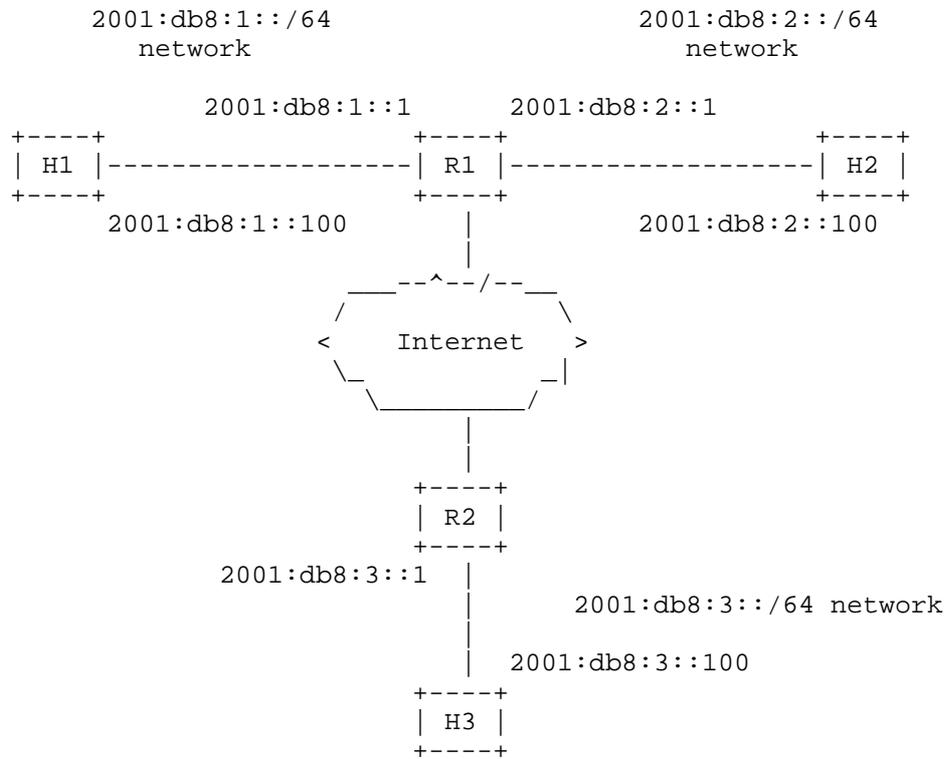


Figure 3: Hypothetical Attack Scenario

In our scenario, the attack packet sent by the attacker would have the same structure as that of Figure 2, with the following values:

- o Source Address: 2001:db8:3::100 (or forged address)
- o Destination Address: 2001:db8:1::100
- o Source Address (embedded packet): 2001:db8:1::100
- o Destination Address (embedded packet): 2001:db8:2::100

The Source Address of the packet is rather irrelevant and need not be forged. The Destination Address of the packet will be that of the attack target (H1 in our case). The Source Address of the embedded packet will be that of the attack target (H1 in our case). Finally, the Destination Address of the embedded packet will be that of the peer with which the attack target is communicating (H2 in our case).

If router R2 were to inspect the payload of the ICMP attack packet, it would conclude that the attack packet cannot be possibly valid, since packets destined to 2001:db8:2::100 would never be forwarded to the network from which the error message is originating. In a similar vein, if R1 were to examine the payload of the aforementioned ICMP error message, it would also conclude that the ICMP error message cannot be possibly valid, for the same reason stated before. Thus, filtering ICMP messages based on the ICMP payload could be employed as a countermeasure for attacks based on ICMP error messages.

5. ICMPv4/ICMPv6 Network Ingress Filtering

A node (e.g. firewall) meaning to enforce the filtering policy specified in this document SHOULD check:

```
IF    embedded packet's Destination Address is from within my network
THEN  forward as appropriate
```

```
IF    embedded packet's Destination Address is anything else
THEN  drop packet
```

NOTE: The destination match is due to a learned route (which assumes some minimal level of path or routing symmetry which firewalls tend to require anyway); or an access list.

We note, however, that the techniques described in [RFC3704] should be evaluated when the aforementioned network ingress filtering is to be implemented in more complex network scenarios, such as that of a multihomed networks. In multihomed scenarios, this filtering policy tends to be undesirable since it is likely to lead to false positives.

Finally, we note that packet drops SHOULD be logged, since this then provides a basis for monitoring any suspicious activity.

6. IANA Considerations

This document has no actions for IANA.

7. Security Considerations

This document provides advice on performing network ingress filtering on ICMPv4 and ICMPv6 error messages, such that attacks based on such messages can be mitigated by means of network packet filtering. Implementation of this filtering technique may depend on the ability of the filtering device to inspect the payload of ICMP messages.

We note that a given platform may or may not be able to filter ICMP error messages based on the ICMP payload. Thus, the aforementioned filter SHOULD only be performed where applicable. Additionally, enforcing the aforementioned filtering method might impact the performance of the filtering device (see e.g., [I-D.gont-v6ops-ipv6-ehs-packet-drops] and [Zack-FW-Benchmark] for a discussion of the IPv6 case). This should be considered before enabling the aforementioned filtering method.

8. Acknowledgements

The authors of this document would like to thank (in alphabetical order) Ron Bonica, Igor Gashinsky, Joel Jaeggli, Merike Kaeo, Jen Linkova, Vic Liu, Carlos Pignataro, and Eric Vyncke, for providing valuable comments on earlier versions of this document.

9. References

9.1. Normative References

- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<http://www.rfc-editor.org/info/rfc791>>.
- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <<http://www.rfc-editor.org/info/rfc792>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, DOI 10.17487/RFC4443, March 2006, <<http://www.rfc-editor.org/info/rfc4443>>.
- [RFC4884] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, "Extended ICMP to Support Multi-Part Messages", RFC 4884, DOI 10.17487/RFC4884, April 2007, <<http://www.rfc-editor.org/info/rfc4884>>.

9.2. Informative References

- [I-D.gont-v6ops-ipv6-ehs-packet-drops]
Gont, F., Hilliard, N., Doering, G., (Will), S., and W. Kumari, "Operational Implications of IPv6 Packets with Extension Headers", draft-gont-v6ops-ipv6-ehs-packet-drops-03 (work in progress), March 2016.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <<http://www.rfc-editor.org/info/rfc2827>>.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <<http://www.rfc-editor.org/info/rfc3704>>.
- [RFC5927] Gont, F., "ICMP Attacks against TCP", RFC 5927, DOI 10.17487/RFC5927, July 2010, <<http://www.rfc-editor.org/info/rfc5927>>.
- [US-CERT] US-CERT, "US-CERT Vulnerability Note VU#222750: TCP/IP Implementations do not adequately validate ICMP error messages", <http://www.kb.cert.org/vuls/id/222750>, 2005.
- [Zack-FW-Benchmark]
Zack, E., "Firewall Security Assessment and Benchmarking IPv6 Firewall Load Tests", IPv6 Hackers Meeting #1, Berlin, Germany. June 30, 2013, <<http://www.ipv6hackers.org/meetings/ipv6-hackers-1/zack-ipv6hackers1-firewall-security-assessment-and-benchmarking.pdf>>.

Authors' Addresses

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: <https://www.si6networks.com>

Ray Hunter
Globis Consulting BV
Weegschaalstraat 3
Eindhoven 5632CW
NL

Email: v6ops@globis.net

Jeroen Massar
Massar Networking
Swiss Post Box 101811
Zuercherstrasse 161
Zuerich CH-8010
CH

Email: jeroen@massar.ch
URI: <http://jeroen.massar.ch>

Will(Shucheng) Liu
Huawei Technologies
Bantian, Longgang District
Shenzhen 518129
P.R. China

Email: liushucheng@huawei.com