INTAREA
Internet-Draft
Intended status: Standards Track
Expires: December 29, 2016

R. Bonica
R. Thomas
Juniper Networks

June 27, 2016

Extended Ping (eping)
draft-bonica-intarea-eping-01

Abstract

This document describes a new diagnostic tool called Extended Ping (eping). Network operators execute eping to determine whether a remote interface is active. In this respect, eping is similar to ping. Eping differs from ping in that it does not require network reachability between itself and remote interface whose status is being queried.

Eping relies on two new ICMP messages, called Extended Echo and Extended Echo Reply. Both ICMP messages are defined herein.

Requirements Language

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].

Status of This Memo

Table of Contents

1. Problem Statement

Network operators use ping [RFC2151] to determine whether a remote interface is alive. Ping sends an ICMP [RFC0792] [RFC4443] Echo message to the interface being probed and waits for an ICMP Echo Reply. If ping receives the expected ICMP Echo Reply, it reports that interface is alive.

In order for the Echo message to reach the probed interface, the probed interface must be addressed appropriately. IP addresses are scoped as follows:

- o Global [RFC4291]

- o Private [RFC1918] [RFC4193]

o Link-local [RFC3927] [RFC4291]

Global addresses are the most widely scoped. A globally addressed interface can be reached from any node on the Internet. By contrast, link-local addresses are the least widely scoped. An interface whose only address is link-local can be reached from on-link interfaces only.

Network operators seek to decrease their dependence on widely-scoped interface addressing. For example:

o The operator of an IPv4 network currently assigns global addresses to all interfaces. In order to conserve scarce IPv4 address space, this operator seeks to renumber selected interfaces with private addresses.

o The operator of an IPv4 network currently assigns private addresses to all interfaces. In order to achieve operational efficiencies, this operator seeks to leave selected interfaces unnumbered.

o The operator of an IPv6 network currently assigns global addresses to all interfaces. In order to achieve operational efficiencies, this operator seeks to allow selected interfaces to be automatically configured with link-local addresses.

When a network operator renumbers an interface, replacing a more widely-scoped address with a less widely-scope address, the operator also reduces the number of nodes from which ping can probe the interface. Furthermore, when a network operator removes all addresses from an interface, leaving it unnumbered, the operator makes that interface totally inaccessible to ping. Therefore, many network operators who rely on ping remain dependant upon widely-scoped interface addressing.

This document describes a new diagnostic tool called Extended Ping (eping). Network operators use eping to determine whether a remote interface is active. In this respect, eping is similar to ping. Eping differs from ping in that it does not require reachability between the probing node and the probed interface. Or, said another way, eping does not require reachability between the node upon which it executes and the interface whose status is being queried.

Eping relies on two new ICMP messages, called Extended Echo and Extended Echo Reply. The Extended Echo message makes a semantic distinction between the destination interface and the probed interface. The destination interface is the interface to which the Extended Echo message is delivered. It must be reachable from the

probing node. The probed interface is the interface whose status is being queried. It does not need to be reachable from the probing node. However, the destination and probed interfaces must be local to one another (i.e., the same node must support both interfaces).

Because the Extended Echo message makes a distinction between the destination and probed interfaces, eping can probe every interface on a node if it can reach any node on the node. In many cases, this allows network operators to decrease their dependence on widely-scoped interface addressing.

This document is divided into sections, with Section 2 describing the Extended Echo message and Section 3 describing the Extended Echo Reply message. Section 4 describes how the probed node processes the Extended Echo message and Section 5 describes the eping application.

2. ICMP Extended Echo

The ICMP Extended Echo message is applicable to both ICMPv4 and ICMPv6. Like any ICMP message, the ICMP Extended Echo message is encapsulated in an IP header. The ICMPv4 version of the Extended Echo message is encapsulated in an IPv4 header, while the ICMPv6 version is encapsulated in an IPv6 header.

Figure 1 depicts the ICMP Extended Echo message.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Type      |     Code      |          Checksum             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Identifier          |        Sequence Number        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   ICMP Extensions ........
```

Figure 1: ICMP Extened Echo Message

IP Source Address: Identifies an interface on the probing node.

IP Destination Address: Identifies the destination interface (i.e., the interface to which this message will be delivered).

Type: Extended Echo (TBD. Value to be assigned by IANA.)

Code: 0

Checksum: For ICMPv4, see RFC 792. For ICMPv6, see RFC 4443.

Identifier: An identifier to aid in matching Extended Echo Replies to this Extended Echo Request. May be zero.

Sequence Number: A sequence number to aid in matching Extended Echo Replies to this Extended Echo Request. May be zero.

If the destination interface is different from the probed interface, the Extended Echo message MUST include ICMP Extensions [RFC4884]. ICMP Extensions MUST include the Interface Identification Object [RFC5837].

If the Extended Echo message does not include the Interface Identification Object, the destination and probed interfaces are understood to be the same.

2.1. Interface Identification Object

The Interface Identification Object identifies the probed interface. It includes an ICMP Object Header (RFC 4884) and object payload,

The ICMP Object Header contains Class-Num and C-Type fields. The Class-Num field MUST be set to Interface Identification Class (2). The C-Type contains an Interface Role and several C-Type flags. The Interface Role MUST be 3 (Next-hop). At least one of the following C-Type flags MUST be set:

- o IPAddress
- o ifIndex
- o name

The MTU flag MUST NOT be set.

If the IPAddress flag is set, the object payload MUST contain an Interface IP Address Sub-Object. If the name flag is set, the object payload MUST contain an Interface IP Name Sub-Object. If the ifIndex flag is set, the object payload MUST contain a 32-bit ifIndex.

If the probed interface is identified by address, its address family does not need to be the same as that of the destination address. For example, the probed interface can be identified by its Ethernet address while the destination address is identified by an IPv4 address.

By default, implementations SHOULD NOT support probing by ifName or ifIndex. See Section 7 for details.

3. ICMP Extended Echo Reply

The ICMP Extended Echo Reply message is applicable to both ICMPv4 and ICMPv6. Like any ICMP message, the ICMP Extended Echo Reply message is encapsulated in an IP header. The ICMPv4 version of the Extended Echo Reply message is encapsulated in an IPv4 header, while the ICMPv6 version is encapsulated in an IPv6 header.

Figure 2 depicts the ICMP Extended Echo Reply message.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Type      |     Code      |          Checksum             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Identifier          |        Sequence Number        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   ICMP Extensions ........
```

Figure 2: ICMP Extened Echo Reply Message

IP source address: Identifies the interface to which the corresponding ICMP Extended Echo message was sent

IP destination address: Identifies the interface from which the corresponding ICMP Extended Echo message was sent

Type: Extended Echo Reply (TBD. Value to be assigned by IANA.)

Code: Indicates operational status of probed interface. Defined values are:

- o Inactive (value to be assigned by IANA)
- o IPv4_Active (value to be assigned by IANA)
- o IPv6_Active (value to be assigned by IANA)
- o IPv4_and_IPv6 Active (value to be assigned by IANA)
- o Interface_does_not_exist (value to be assigned by IANA)
- o Malformed_query (value to be assigned by IANA)

o Query_not_supported (value to be assigned by IANA)

Checksum: For ICMPv4, see RFC 792. For ICMPv6, see RFC 4443.

Identifier: An identifier to aid in matching Extended Echo Replies to this Extended Echo Request. May be zero.

Sequence Number: A sequence number to aid in matching Extended Echo Replies to this Extended Echo Request. May be zero.

ICMP Extensions: By default, the ICMP Extended Echo Reply message MUST NOT include ICMP Extensions. However, the responding node MAY be configured to provide additional information regarding the probed interface using the Interface Identification Object.

4. ICMP Extended Echo and Extended Echo Reply Processing

When a node receives an ICMPv4 Extended Echo, it MUST format an ICMP Extended Echo Reply as follows:

o Don't Fragment flag (DF) is 1

o More Fragments flag is 0

o Fragment Offset is 0

o TTL is 255

o Protocol is ICMP

When a node receives an ICMPv6 Extended Echo, it MUST format an ICMPv6 Extended Echo Reply as follows:

o Hop Limit is 255

o Next Header is ICMPv6

o Flow Label is 0

In either case, the responding node MUST:

o Copy the source address from the Extended Echo message to the destination address of the Extended Echo Reply

o Copy the destination address from the Extended Echo message to the source address of the Extended Echo Reply

o Set the DiffServ codepoint to CS0 [RFC4594]

o Set the ICMP Type to Extended Echo Reply

o Copy the Identifier from the Extended Echo message to the Extended Echo Reply

o Copy the sequence number from the Extended Echo message to the Extended Echo Reply

o Set the code appropriately

o Append ICMP Extensions as required

o Set the checksum appropriately

The following rules govern how the Code should be set:

o If the query is malformed, set the Code to Malformed_query

o If the query type is not supported, set the Code to Query_not_supported

o Determine which interface is being probed. The probed interface matches all of the sub-TLVs in the incoming Interface Identification Object.

o If the interface does not exist, set the Code to Interface_does_not_exist

o If the destination interface is in one security domain and the probed interface is in another security domain, set the Code to Interface_does_not_exist. Virtual Private Networks are examples of security domains.

o Set the code to Inactive, IPv4_active, IPv6_active or IPv4_and_IPv6_aqctive as appropriate.

5. The Eping Application

The eping application accepts input parameters, sets a counter and enters a loop to be exited when the counter is equal to zero. On each iteration of the loop, eping emits an ICMP Extended Echo, decrements the counter, sets a timer, waits for the timer to expire. If an expected ICMP Extended Echo Reply arrives while eping is waiting for the timer to expire, eping relays information returned by that message to its user. However, on each iteration of the loop, eping waits for the timer to expire, regardless of whether an Extended Echo Reply message arrives.

Eping accepts the following parameters:

- o Count
- o Wait
- o Source Interface Address
- o Hop Count
- o Destination Interface Address
- o Probed Interface Identifier

Count is a positive integer whose default value is 3. Count determines the number of times that eping iterates through the above-mentioned loop.

Wait is a positive integer whose minimum and default values are 1. Wait determines the duration of the above-mentioned timer, measured in seconds.

Source Interface Address specifies the source address of ICMP Extended Echo.

The destination Interface Address identifies the interface to which the ICMP Extended Echo message is sent. It can be an IPv4 address or an IPv6 address. If it is an IPv4 address, eping emits an ICMPv4 message. If it is an IPv6 address, eping emits an ICMPv6 message.

The probed interface is the interface whose status is being queried. If the probed interface identifier is not specified, the eping application invokes the traditional ping application and terminates. If the probed interface identifier is specified, it can be any combination of the following:

- o an interface name
- o an address from any address family (e.g., IPv4, IPv6, MAC)
- o an ifIndex

The probed interface identifier can have any scope. For example, the probed interface identifier can be:

- o an IPv6 address, whose scope is global
- o an IPv6 address, whose scope is link-local

- o an interface name, whose scope is node-local
- o an ifIndex, whose scope is node-local

If the probed interface identifier is an address, it does not need to be of the same address family as the destination interface address. For example, eping accepts an IPv4 destination interface address and an IPv6 probed interface identifier.

6. IANA Considerations

This document requests the following actions from IANA:

- o Add an entry to the "ICMP Type Number" registry, representing the Extended Echo. This entry has one code (0).
- o Add an entry to the "ICMPv6 Type Number" registry, representing the Extended Echo. This entry has one code (0).
- o Add an entry to the "ICMP Type Number" registry, representing the Extended Echo Reply. This entry has the following codes: Inactive, IPv4_active, IPv6_acive, IPv4_and_IPv6_active, Interface_does_not_exist, and Query_not_supported.
- o Add an entry to the "ICMPv6 Type Number" registry, representing the Extended Echo Reply. This entry has the following codes: Inactive, IPv4_active, IPv6_acive, IPv4_and_IPv6_active, Interface_does_not_exist, and Query_not_supported.

Note to RFC Editor: this section may be removed on publication as an RFC.

7. Security Considerations

7.1. Probing by ifName and ifIndex

Many implementations encode the following information in an ifName:

- o Interface type (e.g.., Gigabit Ethernet, SONET, T1)
- o Location on chassis (i.e., slot identifier)
- o Location on line card (i.e., port identifier)
- o Location on port (i.e., logical port identifier)

While an operator may have a requirement to probe ports using eping, that operator may not want to expose the above mentioned information.

Therefore, by default, implementations SHOULD NOT support probing by ifName. However, probing by ifName can be enabled through configuration.

Likewise, the ability to probe by if ifIndex may enable certain information to be disclosed to attackers. Therefore, by default, implementations SHOULD NOT support probing by ifIndex. However, probing by ifIndex can be enabled through configuration.

8. Acknowledgements

Thanks to Jeff Haas for his thoughtful review of this document.

9. References

9.1. Normative References

[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>.

[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>.

[RFC5837] Atlas, A., Ed., Bonica, R., Ed., Pignataro, C., Ed., Shen, N., and JR. Rivers, "Extending ICMP for Interface and Next-Hop Identification", RFC 5837, DOI 10.17487/RFC5837, April 2010, <http://www.rfc-editor.org/info/rfc5837>.

9.2. Informative References

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <http://www.rfc-editor.org/info/rfc1918>.

[RFC2151] Kessler, G. and S. Shepard, "A Primer On Internet and TCP/IP Tools and Utilities", FYI 30, RFC 2151, DOI 10.17487/RFC2151, June 1997, <http://www.rfc-editor.org/info/rfc2151>.

[RFC3927] Cheshire, S., Aboba, B., and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses", RFC 3927, DOI 10.17487/RFC3927, May 2005, <http://www.rfc-editor.org/info/rfc3927>.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <http://www.rfc-editor.org/info/rfc4193>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <http://www.rfc-editor.org/info/rfc4291>.

[RFC4594] Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes", RFC 4594, DOI 10.17487/RFC4594, August 2006, <http://www.rfc-editor.org/info/rfc4594>.

Appendix A. An Appendix

Authors' Addresses

Ron Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, Virginia 20171
USA

Email: rbonica@juniper.net

Reji Thomas
Juniper Networks
Elnath-Exora Business Park Survey
Bangalore, Kanata 560103
India

Email: rejithomas@juniper.net