

INTAREA  
Internet-Draft  
Updates: [4884](#) (if approved)  
Intended status: Standards Track  
Expires: June 18, 2018

R. Bonica  
R. Thomas  
Juniper Networks  
J. Linkova  
Google  
C. Lenart  
Verizon  
M. Boucadair  
Orange  
December 15, 2017

PROBE: A Utility For Probing Interfaces  
draft-ietf-intarea-probe-10

Abstract

This document describes a network diagnostic tool called PROBE. PROBE is similar to PING, in that it can be used to query the status of a probed interface. It differs from PING in that it does not require bidirectional connectivity between the probing and probed interfaces. Instead, PROBE requires bidirectional connectivity between the probing interface and a proxy interface. The proxy interface can reside on the same node as the probed interface or it can reside on a node to which the probed interface is directly connected. This document updates [RFC 4884](#).

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 <https://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 June 18, 2018.

Internet-Draft

PROBE

December 2017

## 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 (<https://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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">1.1.</a>	Terminology . . . . .	<a href="#">4</a>
<a href="#">1.2.</a>	Requirements Language . . . . .	<a href="#">4</a>
<a href="#">2.</a>	ICMP Extended Echo Request . . . . .	<a href="#">4</a>
<a href="#">2.1.</a>	Interface Identification Object . . . . .	<a href="#">6</a>
<a href="#">3.</a>	ICMP Extended Echo Reply . . . . .	<a href="#">7</a>
<a href="#">4.</a>	ICMP Message Processing . . . . .	<a href="#">9</a>
<a href="#">4.1.</a>	Code Field Processing . . . . .	<a href="#">10</a>
<a href="#">5.</a>	Use-Cases . . . . .	<a href="#">11</a>
<a href="#">6.</a>	Updates to <a href="#">RFC 4884</a> . . . . .	<a href="#">12</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">12</a>
<a href="#">8.</a>	Security Considerations . . . . .	<a href="#">13</a>
<a href="#">9.</a>	References . . . . .	<a href="#">14</a>
<a href="#">9.1.</a>	Normative References . . . . .	<a href="#">14</a>
<a href="#">9.2.</a>	Informative References . . . . .	<a href="#">15</a>
<a href="#">Appendix A.</a>	The PROBE Application . . . . .	<a href="#">15</a>
	Acknowledgments . . . . .	<a href="#">16</a>
	Authors' Addresses . . . . .	<a href="#">16</a>

[1.](#) Introduction

Network operators use PING [[RFC2151](#)] to test bidirectional connectivity between two interfaces. For the purposes of this document, we will call these interfaces the probing and probed interfaces. PING sends an ICMP [[RFC0792](#)] [[RFC4443](#)] Echo Request message from the probing interface to the probed interface. The

probing interface resides on a probing node while the probed interface resides on a probed node.

If the probed interface receives the ICMP Echo Request message, it returns an ICMP Echo Reply. When the probing interface receives the

ICMP Echo Reply, it has verified bidirectional connectivity between the probing and probed interfaces. Specifically, it has verified that:

- o The probing node can reach the probed interface
- o The probed interface is active
- o The probed node can reach the probing interface
- o The probing interface is active

This document describes a network diagnostic tool called PROBE. PROBE is similar to PING, in that it can be used to query the status of a probed interface. It differs from PING in that it does not require bidirectional connectivity between the probing and probed interfaces. Instead, PROBE requires bidirectional connectivity between the probing interface and a proxy interface. The proxy interface can reside on the same node as the probed interface or it can reside on a node to which the probed interface is directly connected. [Section 5](#) of this document describes scenarios in which this characteristic is useful.

Like PING, PROBE executes on a probing node. It sends an ICMP Extended Echo Request message from a local interface, called the probing interface, to a proxy interface. The proxy interface resides on a proxy node.

The ICMP Extended Echo Request contains an ICMP Extension Structure and the ICMP Extension Structure contains an Interface Identification Object. The Interface Identification Object identifies the probed interface. The probed interface can reside on the proxy node or it can be directly connected to the proxy node.

When the proxy interface receives the ICMP Extended Echo Request, the proxy node executes access control procedures. If access is granted,

the proxy node determines the status of the probed interface and returns an ICMP Extended Echo Reply Message. The ICMP Extended Echo Reply indicates the status of the probed interface.

If the probed interface resides on the proxy node, PROBE determines the status of the probed interface as it would determine its oper-status [[RFC7223](#)]. If oper-status is equal to up (1), PROBE reports that the probed interface is active. Otherwise, PROBE reports that the probed interface is inactive.

If the probed interface resides on a node that is directly connected to the probed node, and the probed interface appears in the IPv4

Address Resolution Protocol (ARP) table [[RFC0826](#)] or IPv6 Neighbor Cache [[RFC4861](#)], PROBE reports interface reachability. Otherwise, PROBE reports that the table entry does not exist.

### [1.1.](#) Terminology

This document uses the following terms:

- o Probing interface - The interface that sends the ICMP Extended Echo Request
- o Probing node - The node upon which the probing interface resides
- o Proxy interface - The interface to which the ICMP Extended Echo Request message is sent
- o Proxy node - The node upon which the proxy interface resides
- o Probed interface - The interface whose status is being queried
- o Probed node - The node upon which the probed interface resides. If the proxy interface and the probed interface reside upon the same node, the proxy node is also the probed node. Otherwise, the proxy node is directly connected to the probed node.

### [1.2.](#) Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and

"OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

## 2. ICMP Extended Echo Request

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

Figure 1 depicts the ICMP Extended Echo Request message.




Figure 1: ICMP Extended Echo Request Message

IP Header fields:

- o Source Address: The Source Address identifies the probing interface. It MUST be valid IPv4 or IPv6 unicast address.
- o Destination Address: The Destination Address identifies the proxy interface. It can be a unicast, multicast or anycast address.

## ICMP fields:

- o Type: Extended Echo Request. The value for ICMPv4 is TTT0. <RFC Editor: Please replace TTT0 with the ICMPv4 type number for Extended Echo Request>. The value for ICMPv6 is TTT1. <RFC Editor: Please replace TTT1 with the ICMPv6 type number for Extended Echo Request> .
- o Code: MUST be set to 0 and MUST be ignored upon receipt.
- o Checksum: For ICMPv4, see [RFC 792](#). For ICMPv6, see [RFC 4443](#).
- o Identifier: An identifier to aid in matching Extended Echo Replies to Extended Echo Requests. May be zero.
- o Sequence Number: A sequence number to aid in matching Extended Echo Replies to Extended Echo Requests. May be zero.
- o Reserved: This field MUST be set to zero and ignored upon receipt.
- o L (local) - The L-bit is set if the probed interface resides on the proxy node. The L-bit is clear if the probed interface is directly connected to the proxy node.
- o ICMP Extension Structure: The ICMP Extension Structure identifies the probed interface.

[Section 7 of \[RFC4884\]](#) defines the ICMP Extension Structure. As per [RFC 4884](#), the Extension Structure contains exactly one Extension Header followed by one or more objects. When applied to the ICMP Extended Echo Request message, the ICMP Extension Structure MUST contain exactly one instance of the Interface Identification Object ([Section 2.1](#)).

If the L-bit is set, the Interface Identification Object can identify the probed interface by name, index or address. If the L-bit is clear, the Interface Identification Object MUST identify the probed interface by address.

If the Interface Identification Object identifies the probed interface by address, that address can be a member of any address family. For example, an ICMPv4 Extended Echo Request message can



| Address ....

Figure 2: Interface Identification Object - C-type 3 Payload

Payload fields are defined as follows:

- o Address Family Identifier (AFI): This 16-bit field identifies the type of address represented by the Address field. All values found in the IANA registry of Address Family Numbers (available from <<https://www.iana.org/assignments/address-family-numbers/address-family-numbers.xhtml>>) are valid in this field.
- o Address Length - Number of significant bytes contained by the Address field. (The address field contains significant bytes and padding bytes)
- o Reserved: This field MUST be set to zero and ignored upon receipt.
- o Address: This variable-length field represents an address associated with the probed interface. If the address field would not otherwise terminate on a 32-bit boundary, it MUST be padded with zeros.

### 3. ICMP Extended Echo Reply

The ICMP Extended Echo Reply message is defined for 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 3 depicts the ICMP Extended Echo Reply message.





- o Res - This field MUST be set to 0 and ignored upon receipt.
- o A (Active) - The A-bit is set if Code is equal to zero, the probed interface resides on the proxy node, and the probed interface is active. Otherwise, the A-bit is clear.
- o 4 (IPv4) - The 4-bit is set if the A-bit is also set and IPv4 is running on the probed interface. Otherwise, the 4-bit is clear.
- o 6 (IPv6) - The 6-bit is set if the A-bit is also set and IPv6 is running on the probed interface. Otherwise, the 6-bit is clear.

#### 4. ICMP Message Processing

When a node receives an ICMP Extended Echo Request message and any of the following conditions apply, the node MUST silently discard the incoming message:

- o The node does not recognize ICMP Extended Echo Request messages
- o The node has not explicitly enabled ICMP Extended Echo functionality
- o The incoming ICMP Extend Echo Request carries a source address that is not explicitly authorized for the incoming ICMP Extended Echo Request L-bit setting
- o The incoming ICMP Extend Echo Request carries a source address that is not explicitly authorized for the incoming ICMP Extended Echo Request type (i.e., by ifName, by IfIndex, by Address)
- o The Source Address of the incoming messages is not a unicast address

Otherwise, when a node receives an ICMPv4 Extended Echo Request, 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

---

Internet-Draft

PROBE

December 2017

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

- o Hop Limit is 255
- o Next Header is ICMPv6

In either case, the responding node MUST:

- o Copy the source address from the Extended Echo Request message to the destination address of the Extended Echo Reply
- o Copy the destination address from the Extended Echo Request 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 Request message to the Extended Echo Reply
- o Copy the Sequence Number from the Extended Echo Request message to the Extended Echo Reply
- o Set the Code field as described [Section 4.1](#)
- o Set the State field to 0.
- o Clear the A-bit, the 4-bit and the 6-bit.
- o If the Code Field is equal to (0) No Error and the L-bit is set and the probed interface is active, set the A-bit. Also set the 4-bit and the 6-bit as appropriate.
- o If the Code Field is equal to (0) No Error and the L-bit is clear, set the State field to reflect the state of the ARP Table or Neighbor Cache entry that represents the probed interface.

- o Set the checksum appropriately
- o Forward the ICMP Extended Echo Reply to its destination

#### 4.1. Code Field Processing

The Code field MUST be set to (1) Malformed Query if any of the following conditions apply:

- o The ICMP Extended Echo Request does not include an ICMP Extension Structure
- o The ICMP Extension Structure does not include exactly one Interface Identification Object
- o The L-bit is clear and the Interface Identification Object identifies the probed interface by ifName or ifIndex
- o The query is otherwise malformed

The Code field MUST be set to (2) No Such Interface if the L-bit is set and the ICMP Extension Structure does not identify an interface that resides on the proxy node.

The Code field MUST be set to (3) No Such Table Entry if the L-bit is clear and the address found in the Interface Identification Object does not appear in the IPv4 Address Resolution Protocol (ARP) Table or the IPv6 Neighbor Cache.

The Code field MUST be set to (4) Multiple Interfaces Satisfy Query if any of the following conditions apply:

- o The L-bit is set and the ICMP Extension Structure identifies more than one interface that resides in the proxy node
- o The L-bit is clear and the address found in the Interface Identification Object maps to multiple IPv4 ARP or IPv6 Neighbor Cache entries

Otherwise, the Code field MUST be set to (0) No Error

## 5. Use-Cases

In the scenarios listed below, network operators can use PROBE to determine the status of a probed interface, but cannot use PING for the same purpose. In all scenarios, assume bidirectional connectivity between the probing and proxy interfaces. However, bidirectional connectivity between the probing and probed interfaces is lacking.

- o The probed interface is unnumbered
- o The probing and probed interfaces are not directly connected to one another. The probed interface has an IPv6 link-local address, but does not have a more globally scoped address

- o The probing interface runs IPv4 only while the probed interface runs IPv6 only
- o The probing interface runs IPv6 only while the probed interface runs IPv4 only
- o For lack of a route, the probing node cannot reach the probed interface.

## 6. Updates to [RFC 4884](#)

[Section 4.6 of RFC 4884](#) provides a list of extensible ICMP messages (i.e., messages that can carry the ICMP Extension Structure). This document adds the ICMP Extended Echo Request message and the ICMP Extended Echo Reply message to that list.

## 7. IANA Considerations

This document requests the following actions from IANA:

- o Add an entry to the "ICMP Type Number" registry, representing the Extended Echo Request. This entry has one code (0) No Error.
- o Add an entry to the "ICMPv6 "type" Numbers" registry, representing the Extended Echo Request. This entry has one code (0) No Error.

As ICMPv6 distinguishes between informational and error messages, and this is an informational message, the value must be assigned from the range 128-255.

- o Add an entry to the "ICMP Type Number" registry, representing the Extended Echo Reply. This entry has the following codes: (0) No Error, (1) Malformed Query, (2) No Such Interface, (3) No Such Table Entry, (4) Multiple Interfaces Satisfy Query.
- o Add an entry to the "ICMPv6 "type" Numbers" registry, representing the Extended Echo Reply. This entry has the following codes: (0) No Error, (1) Malformed Query, (2) No Such Interface, (3) No Such Table Entry, (4) Multiple Interfaces Satisfy Query.
- o Add an entry to the "ICMP Extension Object Classes and Class Subtypes" registry, representing the Interface Identification Object. It has C-types Reserved (0), Identifies Interface By Name (1), Identifies Interface By Index (2), Identifies Interface By Address (3). C-Type values are assignable on a first-come-first-serve (FCFS) basis with a range of 0-255.

All codes mentioned above are assigned on a First Come First Serve (FCFS) basis with a range of 0 -255.

## [8.](#) Security Considerations

The following are legitimate uses of PROBE:

- o to determine the operational status of an interface
- o to determine which protocols (e.g., IPv4, IPv6) are active on an interface

However, malicious parties can use PROBE to obtain additional information. For example, a malicious party can use PROBE to discover interface names. Having discovered an interface name, the malicious party may be able to infer additional information. Additional information may include:

- o interface bandwidth
- o the type of device that supports the interface (e.g., vendor

identity)

- o the operating system version that the above-mentioned device executes

Understanding this risk, network operators establish policies that restrict access to ICMP Extended Echo functionality. In order to enforce these policies, nodes that support ICMP Extended Echo functionality MUST support the following configuration options:

- o Enable/disable ICMP Extended Echo functionality. By default, ICMP Extended Echo functionality is disabled.
- o Define enabled L-bit settings. By default, L-bit set is enabled and L-bit clear is disabled.
- o Define enabled query types (i.e., by name, by index, by address). By default, all query types are disabled.
- o For each enabled query type, define the prefixes from which ICMP Extended Echo Request messages are permitted
- o For each interface, determine whether ICMP Echo Request messages are accepted

When a node receives an ICMP Extended Echo Request message that it is not configured to support, it MUST silently discard the message. See [Section 4](#) for details.

PROBE must not leak information about one Virtual Private Network (VPN) into another. Therefore, when a node receives an ICMP Extended Echo Request and the proxy interface is in a different VPN than the probed interface, the node MUST return an ICMP Extended Echo Reply with error code equal to (2) No Such Interface.

In order to protect local resources, implementations SHOULD rate-limit incoming ICMP Extended Echo Request messages.

## [9.](#) References

## 9.1. Normative References

- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, [RFC 792](#), DOI 10.17487/RFC0792, September 1981, <<https://www.rfc-editor.org/info/rfc792>>.
- [RFC0826] Plummer, D., "Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", STD 37, [RFC 826](#), DOI 10.17487/RFC0826, November 1982, <<https://www.rfc-editor.org/info/rfc826>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://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", STD 89, [RFC 4443](#), DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [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, <<https://www.rfc-editor.org/info/rfc4884>>.
- [RFC7223] Bjorklund, M., "A YANG Data Model for Interface Management", [RFC 7223](#), DOI 10.17487/RFC7223, May 2014, <<https://www.rfc-editor.org/info/rfc7223>>.

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.



## 9.2. Informative References

- [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, <<https://www.rfc-editor.org/info/rfc2151>>.
- [RFC4594] Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes", [RFC 4594](#), DOI 10.17487/RFC4594, August 2006, <<https://www.rfc-editor.org/info/rfc4594>>.

## Appendix A. The PROBE Application

The PROBE 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, PROBE emits an ICMP Extended Echo Request, decrements the counter, sets a timer and waits. The ICMP Extended Echo Request includes an Identifier and a Sequence Number.

If an ICMP Extended Echo Reply carrying the same Identifier and Sequence Number arrives, PROBE relays information returned by that message to its user. However, on each iteration of the loop, PROBE waits for the timer to expire, regardless of whether an Extended Echo Reply message arrives.

PROBE accepts the following parameters:

- o Count
- o Wait
- o Probing Interface Address
- o Hop Count
- o Proxy Interface Address
- o Local
- o Probed Interface Identifier

Count is a positive integer whose default value is 3. Count determines the number of times that PROBE 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.

Probing Interface Address specifies the source address of ICMP Extended Echo Request. The Probing Interface Address MUST be a unicast address and MUST identify an interface that resides on the probing node.

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

Local is a boolean value. It is TRUE if the proxy and probed interfaces both reside on the same node. Otherwise, it is FALSE.

The Probed Interface Identifier identifies the probed interface. It is one of the following:

- o an interface name
- o an address from any address family (e.g., IPv4, IPv6, IEEE 802, 48-bit MAC, 64-bit MAC)
- o an if-index

If the Probed Interface Identifier is an address, it does not need to be of the same address family as the proxy interface address. For example, PROBE accepts an IPv4 Proxy Interface Address and an IPv6 Probed Interface Identifier

## Acknowledgments

Thanks to Sowmini Varadhan, Jeff Haas, Carlos Pignataro, Jonathan Looney, Dave Thaler, Mikio Hara, Joel Halpern, Yaron Sheffer, Stefan Winter, Jean-Michel Combes, Amanda Barber and Joe Touch for their thoughtful review of this document.

## Authors' Addresses

Internet-Draft

PROBE

December 2017

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

Email: [rbonica@juniper.net](mailto:rbonica@juniper.net)

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

Email: [rejithomas@juniper.net](mailto:rejithomas@juniper.net)

Jen Linkova  
Google  
1600 Amphitheatre Parkway  
Mountain View, California 94043  
USA

Email: [furry@google.com](mailto:furry@google.com)

Chris Lenart  
Verizon  
22001 Loudoun County Parkway  
Ashburn, Virginia 20147  
USA

Email: [chris.lenart@verizon.com](mailto:chris.lenart@verizon.com)

Mohamed Boucadair  
Orange  
Rennes 35000  
France

Email: mohamed.boucadair@orange.com

Bonica, et al.

Expires June 18, 2018

[Page 17]