

INTAREA
Internet-Draft
Updates: RFC [4884](#) (if approved)
Intended status: Standards Track
Expires: January 17, 2018

R. Bonica
R. Thomas
Juniper Networks
J. Linkova
Google
C. Lenart
Verizon
M. Boucadair
Orange
July 16, 2017

PROBE: A Utility For Probing Interfaces
draft-ietf-intarea-probe-01

Abstract

This document describes a network diagnostic tool called PROBE. PROBE is similar to PING, in that it can be used to test the status of a probed interface. It differs from PING in that it does not require bidirectional connectivity between the probing and probed interfaces. Alternatively, 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.

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 17, 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
1.1.	Terminology	4
1.2.	Requirements Language	4
2.	ICMP Extended Echo Request	4
2.1.	Interface Identification Object	6
3.	ICMP Extended Echo Reply	7
4.	ICMP Message Processing	8
4.1.	Code Field Processing	10
5.	Use-Cases	11
6.	Updates to RFC 4884	11
7.	IANA Considerations	11
8.	Security Considerations	12
9.	References	13
9.1.	Normative References	13
9.2.	Informative References	14
Appendix A.	The PROBE Application	14
	Acknowledgments	16
	Authors' Addresses	16

[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 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 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 test the status of a probed interface. It differs from PING in that it does not require bidirectional connectivity between the probing and probed interfaces. Alternatively, 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 message from a local interface, called the probing interface, to a proxy interface. The proxy interface resides on a probed 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 probed node or it can be directly connected to the probed node.

When the proxy interface receives the ICMP Extended Echo Request, it executes access control procedures. If access is granted, the probed 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 probed node, PROBE determines the status of the probed interface as it would determine its MIB-II [\[RFC2863\]](#) ifOperStatus. If ifOperStatus 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, PROBE reports that the interface is up if it appears in the IPv4 Address Resolution Protocol (ARP) table or the

IPv6 Neighbor Cache. Otherwise, it reports that the interface does not exist.

1.1. Terminology

This document uses the following terms:

- o Probing node - The node upon which PROBE executes
- o Probing interface - The interface from which an ICMP Extended Echo originates
- o Proxy interface - The interface to which the ICMP Extended Echo message is sent
- o Probed node - The node upon which the proxy interface resides
- o Probed interface - The interface whose status is being queried

1.2. 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](#)].

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.

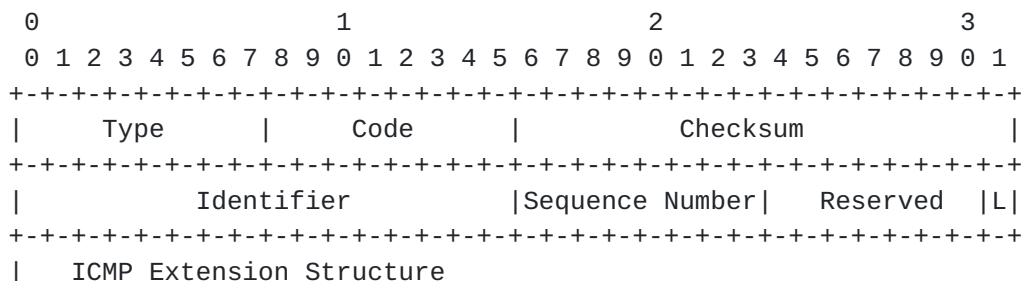


Figure 1: ICMP Extended Echo Request Message

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 TBD by IANA. The value for ICMPv6 is also TBD by IANA.
- o Code: 0
- 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 probed node. The L-bit is clear if the probed interface is directly connected to the probed 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 one or two instances of the Interface Identification Object ([Section 2.1](#)).

In most cases, a single instance of the Interface Identification Object identifies the probed interface. However, in some cases, a second instance is required for disambiguation.

If the L-bit is set, the Interface Identification Object identifies the probed interface by name, index or address. If the L-bit is clear, the Interface Identification Object identifies 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 carry an Interface Identification Object that identifies the probed interface by IPv4, IPv6 or IEEE 802 address. Likewise, an ICMPv6 Extended Echo Request message can carry an Interface Identification Object that identifies the probed interface by IPv4, IPv6 or IEEE 802 address.

2.1. Interface Identification Object

The Interface Identification Object identifies the probed interface by name, index, or address. Like any other ICMP Extension Object, it contains an Object Header and Object Payload. The Object Header contains the following fields:

- o Class-Num: Interface Identification Object. Value is TBD by IANA
- o C-type: Values are: (1) Identifies Interface By Name, (2) Identifies Interface By Index, and (3) Identifies Interface By Address
- o Length: Length of the object, measured in octets, including the object header and object payload.

If the Interface Identification Object identifies the probed interface by name, the object payload contains the human-readable interface name. The interface name SHOULD be the full MIB-II ifName, if less than 255 octets, or the first 255 octets of the ifName, if the ifName is longer. The interface name MAY be some other human-meaningful name of the interface. The interface name MUST be represented in the UTF-8 charset [[RFC3629](#)] using the Default Language [[RFC2277](#)].

If the Interface Identification Object identifies the probed interface by index, the length is equal to 8 and the payload contains the MIB-II ifIndex [[RFC 2863](#)].

If the Interface Identification Object identifies the probed interface by address, the payload is as depicted in Figure 2.

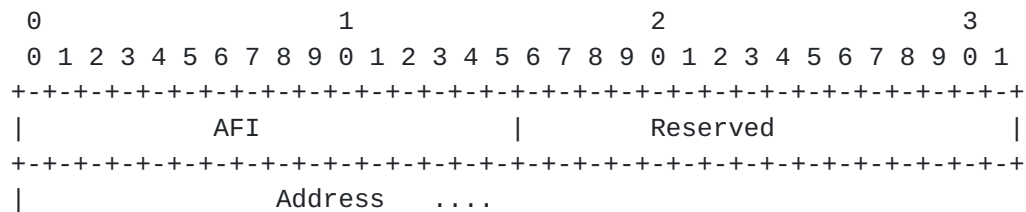


Figure 2: Interface Identification Object - C-type 3 Payload

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 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.

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.

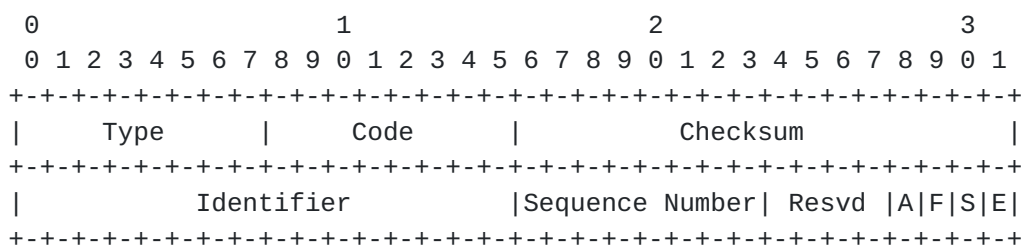


Figure 3: ICMP Extended Echo Reply Message

IP Header fields:

- o Source address: Copied from the Destination Address field of the invoking Extended Echo Request message
- o Destination address: Copied from the Source Address field of the invoking Extended Echo Request message

ICMP fields:

- o Type: Extended Echo Reply. The value for ICMPv4 is TBD by IANA. The value for ICMPv6 is also TBD by IANA
- o Code: (0) No Error, (1) Malformed Query, (2) No Such Interface, (3) Multiple Interfaces Satisfy Query
- o Checksum: For ICMPv4, see [RFC 792](#). For ICMPv6, see [RFC 4443](#)
- o Identifier: Copied from the Identifier field of the invoking Extended Echo Request packet
- o Sequence Number: Copied from the Sequence Number field of the invoking Extended Echo Request packet
- o Resvd - This field MUST be set to zero and ignored upon receipt
- o A (Active) - The A-bit is set if Code is equal to zero and the probed interface is active. Otherwise, the A-bit is clear.
- o F (IPv4) - The F-bit is set if the A-bit is also set and IPv4 is running on the probed interface. Otherwise, the F-bit is clear.
- o S (IPv6) - The S-bit is set if the A-bit is also set and IPv6 is running on the probed interface. Otherwise, the S-bit is clear.
- o E (Ethernet) - The E-bit is set if the A-bit is also set and IPv4 is running on the probed interface. Otherwise, the E-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

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 If the Code Field is equal to No Error (0) and the L-bit is clear, set the A-Bit.
- o If the Code Field is equal to No Error (0) and the L-bit is set and the probed interface is active, set the A-bit.
- o If the A-bit is set, set the F-bit, S-bit and E-bit as appropriate. Otherwise, clear the F, S and E bits.
- 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 Malformed Query (1) 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 an Interface Identification Object
- o The ICMP Extension Structure contains more than two Interface Identification Objects
- 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 No Such Interface (2) if any of the following conditions apply:

- o The L-bit is set and the ICMP Extension Structure does not identify any local interfaces
- o The L-bit is clear and the address or addresses found in the Interface Identification object appear in neither the IPv4 Address Resolution Protocol (ARP) Table nor the IPv6 Neighbor Cache

The Code field MUST be set to Multiple Interfaces Satisfy Query (3) if any of the following conditions apply:

- o The L-bit is set and the ICMP Extension Structure identifies more than one local interfaces
- o The L-bit is clear and the address or addresses found in the Interface Identification object map to multiple IPv4 ARP or IPv6 Neighbor Cache entries

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

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 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).
- o Add an entry to the "Internet Control Message Protocol version 6 (ICMPv6) Parameters" registry, representing the Extended Echo Request. 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: (0) No Error, (1) Malformed Query, (2) No Such Interface, (3) Multiple Interfaces Satisfy Query. Protocol Flag Bit mappings are as follows: Bit 0 (IPv4), Bit 1 (IPv6), Bit 2 (Ethernet), Bits 3-15 (Reserved).
- o Add an entry to the "Internet Control Message Protocol version 6 (ICMPv6) Parameters" registry, representing the Extended Echo Reply. This entry has the following codes: (0) No Error, (1) Malformed Query, (2) No Such Interface, (3) Multiple Interfaces Satisfy Query. Protocol Flag Bit mappings are as follows: Bit 0 (IPv4), Bit 1 (IPv6), Bit 2 (Ethernet), Bits 3-15 (Reserved).
- 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)

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

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 ifName, by ifIndex, 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, <<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>>.
- [RFC2277] Alvestrand, H., "IETF Policy on Character Sets and Languages", [BCP 18](#), [RFC 2277](#), DOI 10.17487/RFC2277, January 1998, <<http://www.rfc-editor.org/info/rfc2277>>.
- [RFC2863] McCloghrie, K. and F. Kastenholz, "The Interfaces Group MIB", [RFC 2863](#), DOI 10.17487/RFC2863, June 2000, <<http://www.rfc-editor.org/info/rfc2863>>.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, [RFC 3629](#), DOI 10.17487/RFC3629, November 2003, <<http://www.rfc-editor.org/info/rfc3629>>.
- [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, <<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

- [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>>.
- [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. 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, waits for the timer to expire. If an expected ICMP Extended Echo Reply arrives while PROBE is waiting for the timer to expire, 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 is local to 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 probed node. It is **FALSE** if the proxy interface resides on the probed node and the probed interface is directly connected to the probed node.

The probed interface is the interface whose status is being queried. It is identified by 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 ifIndex

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 destination interface address and an IPv6 probed interface identifier

Acknowledgments

Thanks to Sowmini Varadhan, Jeff Haas, Carlos Pignataro, Jonathan Looney and Joe Touch for their thoughtful review of this document.

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, Karnataka 560103
India

Email: rejithomas@juniper.net

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

Email: furry@google.com

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

Email: chris.lenart@verizon.com

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

