

6MAN
Internet-Draft
Updates: 4884 (if approved)
Intended status: Standards Track
Expires: 20 April 2026

Y. Liu
ZTE
Y. Liu
China Mobile
17 October 2025

Extending ICMPv6 for SRv6-related Information Validation
draft-liu-6man-icmp-verification-08

Abstract

This document introduces the mechanism to verify the data plane against the control plane and detect data plane failures in IPv6/SRv6 networks by extending ICMPv6 messages.

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 20 April 2026.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements Language	4
2. ICMPv6 Validation Request	4
2.1. Validation Information Object	5
3. ICMPv6 Validation Reply	6
4. ICMPv6 Validation Message Processing	6
4.1. Sending a Validation Request	7
4.2. Receiving a Validation Request	7
4.3. Sending a Validation Reply	8
4.3.1. Return Code	9
4.4. Receiving a Validation Reply	9
5. Updates to RFC 4884	9
6. IANA Considerations	9
7. Security Considerations	10
8. Acknowledgement	10
9. References	11
9.1. Normative References	11
9.2. Informative References	11
Authors' Addresses	13

1. Introduction

An SR-MPLS SID/MPLS label can be related with various FEC information, e.g, VPN IP prefix [RFC4365], EVPN service information [RFC7432], flex algorithms[RFC9350] and etc. Most of these information would be advertised via control plane protocols(e.g, IGP, BGP, etc).

Procedures for simple and efficient mechanisms to verify the data plane against the control plane using LSP Ping in MPLS network are well defined in [RFC8029]. Normally, when a new feature is introduced and the MPLS label is associated with new information, the LSP Ping mechanism is still applicable by defining new FEC sub-TLV with the new information encoded in it.

[RFC9489] defines procedures to detect data plane failures using LSP Ping in MPLS networks deploying EVPN. Figure 1 is an unicast data plane connectivity check scenario provided in [RFC9489]. CE1 is dual-homed to PE1 and PE2, PE1 and PE2 both announced a MAC route for CE1 with the same C-MAC but different RDs. On PE3, when an operator performs a connectivity check for the C-MAC address on PE1, the operator initiates an LSP Ping request with the Target FEC Stack TLV containing the C-MAC address, the corresponding RD and other necessary fields in the packet. The egress PE will process the packet and perform checks for the EVPN-related information carried in Target FEC Stack TLV.




Figure 1: EVPN Network

For VPN/EVPN services over SRv6 as in [RFC9252], the requirements to detect data plane failures are similar. Besides VPN/EVPN information, IPv6 addresses (mainly SRv6 SID), can be related with other information/functions such as flex algorithm [RFC9350] [RFC9502], SRv6 endpoint behaviors [RFC8986], service functions [I-D.ietf-spring-sr-service-programming] and so on. Operators may want to validate these information as well. But there's no such validation mechanism in IPv6/SRv6 yet.

[RFC9259] describes how the existing ICMPv6 mechanisms for ping and traceroute can be used in an SRv6 network for data plane connectivity check purpose. This document introduces the mechanism to verify the data plane against the control plane and to detect data plane failures in IPv6/SRv6 networks by extending ICMPv6 messages.

Editor's Note: Instead of extending ICMPv6 Node Information Query (or NI Query) and the Node Information Reply (or NI Reply) based on [RFC4620], this document introducing ICMPv6 Validation Request and ICMPv6 Validation Reply messages by defining two new types of ICMPv6 messages taking example from [RFC8335]. The reason is that NI Query and NI Reply are originally defined for discovering information about nodes, such as names and addresses, while this document aims to provide an IP-related information validation mechanism, which makes RFC4620 not quite suitable.

1.1. 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. ICMPv6 Validation Request

The Validation Request message is defined for ICMPv6[RFC4443]. Like any ICMPv6 message, the ICMPv6 Validation Request message is encapsulated in an IPv6 header.

The structure of ICMPv6 Validation Request is shown in Figure 2, where:

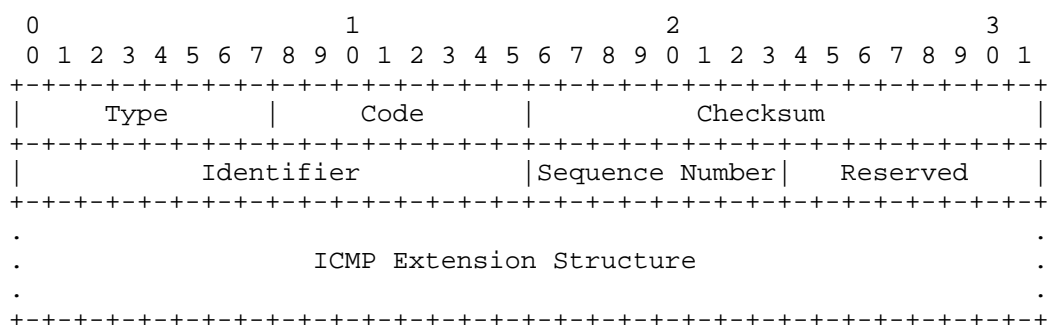


Figure 2: Validation Request

- * Type: The value is TBD1.
- * Code: MUST be set to 0 and MUST be ignored upon receipt.
- * Checksum: For ICMPv6, see [RFC4443].
- * Identifier: An Identifier to aid in matching Validation Replies to Validation Requests. May be zero.
- * Sequence Number: A Sequence Number to aid in matching Validation Replies to Validation Requests. May be zero.
- * Reserved: This field MUST be set to 0 and ignored upon receipt.

- * ICMP Extension Structure: The ICMP Extension Structure carries the information that needs to be verified. Section 7 of [RFC4884] defines the ICMP Extension Structure. As per [RFC4884], the Extension Structure contains one Extension Header followed by one or more objects. When applied to the ICMP Validation Request message, the ICMP Extension Structure MUST only contain one or more instance of the Validation Information Objects as defined in section 2.1.

2.1. Validation Information Object

The Validation Information Object is shown in Figure 3, where:

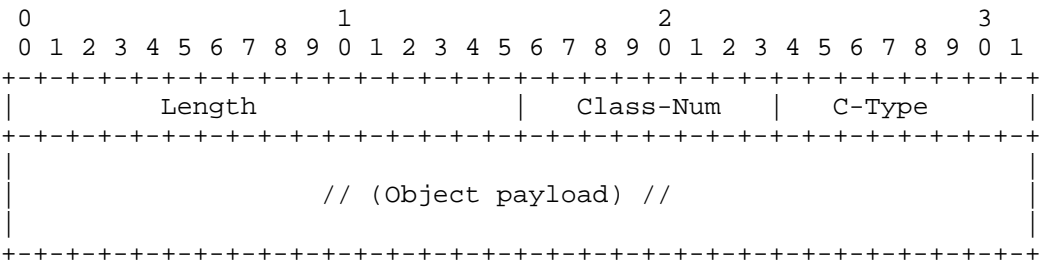


Figure 3: Validation Information Object

- * Length: Length of the object, measured in octets, including the Object Header and Object Payload.
- * Class-Num: Validation Information Object. The value is TBD2.
- * Object payload: Variable-length field. C-Type-specific data.
- * C-Type: For this object, the C-Type is used to indicate the type of the information that needs to be verified. This document only defines the C-Type value 0 as shown below:

C-Type	Object Payload
-----	-----
0	Revsered

This document only defines fundamental packet formats and processing rules, the detailed C-Type values and the corresponding information carried in object payload is out of the scope of the document and would be defined in separate documents. For example, [I-D.liu-bess-srv6-evpn-validation] provides solutions to detect data plane failures for EVPN over SRv6 leveraging the mechanism defined in this document.

3. ICMPv6 Validation Reply

The Validation Reply message is defined for ICMPv6. Like any other ICMPv6 messages, the Validation Reply message is encapsulated in an IPv6 header. Figure 4 describes the ICMPv6 Validation Reply message.




Figure 4: Validation Reply

ICMP fields:

- * Type: Validation Reply. The value is TBD3.
- * Code: Values are
 - (0) Validation passed
 - (1) Malformed request received
 - (2) One or more of the objects were not understood
 - (3) Information mismatch
- * Checksum: For ICMPv6, see [RFC4443].
- * Identifier: Copied from the Identifier field of the invoking Validation Request packet.
- * Sequence Number: Copied from the Sequence Number field of the invoking Validation Request packet.

4. ICMPv6 Validation Message Processing

4.1. Sending a Validation Request

A node that originates an ICMPv6 validation request message SHOULD first determine which IPv6 address/SRV6 SID needs to be verified with what information. How the sender node get the information is out of scope of the document.

An ICMPv6 validation request contains one or more Validation Information objects, depending on how the user wants to do the validation. For example, an SRv6 service SID is related with an endpoint behavior and an IPv4 VPN prefix, if one wants to verify both information of the SID via one request message, an ICMPv6 validation request is sent with two validation information objects in it. Or one may choose to send two individual ICMPv6 validation requests, each carries one validation information object to verify these two information separately.

The target IP is the IP address/SRV6 SID to be verified and MUST be a unicast address. The ICMPv6 validation request is sent with the target IP address/SRV6 SID set as the destination address of the IP header field without SRH for the SRv6 best-effort connectivity case, or set as the last segment with SRH. The Source Address of the ICMPv6 packet MUST be a unicast address belonging to the node.

The Hop Limit SHOULD be set to 255 to prevent transit nodes from processing the validation request.

4.2. Receiving a Validation Request

All transit nodes process the validation request message like any other IPv6 data packets and hence do not require any change.

As specified in [RFC4443], if a router receives a packet with a Hop Limit of zero, or if a router decrements a packet's Hop Limit to zero, it MUST discard the packet and originate an ICMPv6 Time Exceeded message with Code 0 to the source of the packet. The source address SHOULD be set as a local address of the router.

The target node is a node receiving an validation request where the target IP of that message is locally configured as a segment or local interface.

When the validation request packet arrives at the target node, and any of the following conditions apply, the node MUST silently discard the incoming message:

- * The node does not recognize ICMPv6 Validation Request messages.

- * The node has not explicitly enabled ICMPv6 Validation functionality.
- * The incoming ICMPv6 Validation Request carries a Source Address that is not explicitly authorized for the incoming ICMP Validation Request type.
- * The Source Address of the incoming message is not a unicast address.
- * The Destination Address of the incoming message is not a unicast address.

Otherwise, if the packet is well formed, the target node verifies the information encoded in the Validation Information Object against the corresponding local information and state.

4.3. Sending a Validation Reply

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

- * Copy the Source Address from the Validation Request message to the Destination Address of the Validation Reply.
- * Copy the Destination Address from the Validation Request message to the Source Address of the Validation Reply.
- * Set the Hop Limit to 255
- * Set the Next Header to ICMPv6.
- * Set the DiffServ codepoint to CS0 [RFC4594].
- * Set the ICMP Type to Validation Reply.
- * Copy the Identifier from the Validation Request message to the Validation Reply.
- * Copy the Sequence Number from the Validation Request message to the Validation Reply.
- * Set the Code field as described in Section 4.3.1
- * Set the Checksum appropriately.
- * Forward the ICMP Validation Reply to its destination.

4.3.1. Return Code

The Code field MUST be set to 0 if all the the information encoded in the Validation Information Object is consistent with the the corresponding local information on the target node.

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

- * The ICMP Request does not include an ICMPv6 Extension Structure.
- * The ICMP Extension Structure does not only include the Validation Information Object(s).
- * The query is otherwise malformed.

The Code field MUST be set to 2 if one or more of the objects are not understood by the node.

The Code field MUST be set to 3 if the information in the Validation Information Object(s) is not consistent with the local information and validation is not passed.

4.4. Receiving a Validation Reply

A node should only receive a validation reply in response to a validation request that it sent. Thus, on receipt of a validation reply, the node should parse the packet to ensure that it is well-formed, then attempt to match up the validation reply with a validation request that it had previously sent, using the Identifier and Sequence Number. If no match is found, the node ignores the echo reply.

5. Updates to RFC 4884

Section 4.6 of [RFC4884] provides a list of extensible ICMP messages (i.e., messages that can carry the ICMP Extension Structure). This document adds the ICMPv6 Validation Request message and the ICMPv6 Validation Reply message to that list.

6. IANA Considerations

This document requests the following actions from IANA:

- * Add an entry to the "ICMPv6 "type" Numbers" registry, representing the Validation Request. This entry has one code 0.

- * Add an entry to the "ICMPv6 "type" Numbers" registry, representing the Validation Reply. This entry has the following codes:

(0) Validation passed

(1) Malformed request received

(2) One or more of the objects were not understood

(3) Information mismatch

- * Add an entry to the "ICMP Extension Object Classes and Class Subtypes" registry, representing the Validation Information Object with C-types:

(0) Reserved

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.

7. Security Considerations

Security considerations discussed in [RFC4443] and [RFC4884] apply to this document.

To protect against unauthorized sources using validation request messages to obtain network information, it is RECOMMENDED that implementations provide a means of checking the source addresses of validation request messages against an access list before accepting the message.

The validation mechanism SHOULD be only used in the limited domain. The validation request contains the control plane information, policies should be implemented on the edge devices of the domain to prevent the information from being leaked into other domains.

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

8. Acknowledgement

The authors would like to thank Greg Mirsky, Bruno Decraene, Matthew Bocci, Zafar Ali, Ali Sajassi and Jorge Rabadan for their comments and suggestions.

9. References

9.1. Normative References

- [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>>.
- [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>>.
- [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>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

9.2. Informative References

- [I-D.ietf-spring-sr-service-programming] Clad, F., Xu, X., Filsfils, C., Bernier, D., Li, C., Decraene, B., Ma, S., Yadlapalli, C., Henderickx, W., and S. Salsano, "Service Programming with Segment Routing", Work in Progress, Internet-Draft, draft-ietf-spring-sr-service-programming-11, 23 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-sr-service-programming-11>>.
- [I-D.liu-bess-srv6-evpn-validation] Liu, Y., "Data Plane Failure Detection Mechanisms for EVPN over SRv6", Work in Progress, Internet-Draft, draft-liu-bess-srv6-evpn-validation-01, 2 August 2024, <<https://datatracker.ietf.org/doc/html/draft-liu-bess-srv6-evpn-validation-01>>.

- [RFC4365] Rosen, E., "Applicability Statement for BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4365, DOI 10.17487/RFC4365, February 2006, <<https://www.rfc-editor.org/info/rfc4365>>.
- [RFC4594] Babiarez, 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>>.
- [RFC4620] Crawford, M. and B. Haberman, Ed., "IPv6 Node Information Queries", RFC 4620, DOI 10.17487/RFC4620, August 2006, <<https://www.rfc-editor.org/info/rfc4620>>.
- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, DOI 10.17487/RFC5036, October 2007, <<https://www.rfc-editor.org/info/rfc5036>>.
- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.
- [RFC8029] Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", RFC 8029, DOI 10.17487/RFC8029, March 2017, <<https://www.rfc-editor.org/info/rfc8029>>.
- [RFC8335] Bonica, R., Thomas, R., Linkova, J., Lenart, C., and M. Boucadair, "PROBE: A Utility for Probing Interfaces", RFC 8335, DOI 10.17487/RFC8335, February 2018, <<https://www.rfc-editor.org/info/rfc8335>>.
- [RFC9252] Dawra, G., Ed., Talaulikar, K., Ed., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)", RFC 9252, DOI 10.17487/RFC9252, July 2022, <<https://www.rfc-editor.org/info/rfc9252>>.
- [RFC9259] Ali, Z., Filsfils, C., Matsushima, S., Voyer, D., and M. Chen, "Operations, Administration, and Maintenance (OAM) in Segment Routing over IPv6 (SRv6)", RFC 9259, DOI 10.17487/RFC9259, June 2022, <<https://www.rfc-editor.org/info/rfc9259>>.

- [RFC9350] Psenak, P., Ed., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", RFC 9350, DOI 10.17487/RFC9350, February 2023, <<https://www.rfc-editor.org/info/rfc9350>>.
- [RFC9489] Jain, P., Sajassi, A., Salam, S., Boutros, S., and G. Mirsky, "Label Switched Path (LSP) Ping Mechanisms for EVPN and Provider Backbone Bridging EVPN (PBB-EVPN)", RFC 9489, DOI 10.17487/RFC9489, November 2023, <<https://www.rfc-editor.org/info/rfc9489>>.
- [RFC9502] Britto, W., Hegde, S., Kaneriya, P., Shetty, R., Bonica, R., and P. Psenak, "IGP Flexible Algorithm in IP Networks", RFC 9502, DOI 10.17487/RFC9502, November 2023, <<https://www.rfc-editor.org/info/rfc9502>>.

Authors' Addresses

Yao Liu
ZTE
Nanjing
China
Email: liu.yao71@zte.com.cn

Yisong Liu
China Mobile
China
Email: liuyisong@chinamobile.com