6man
Internet-Draft
Intended status: Standards Track
Expires: 28 August 2026

B. Varga
J. Halpern
Ericsson

24 February 2026

ICMP Error Handling for VPNs in SRv6 Networks
draft-varhal-6man-icmp-srv6-vpn-00

Abstract

This document specifies ICMP error handling in SRv6-based Virtual Private Networks.

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 28 August 2026.

Table of Contents

1. Introduction 2
2. Terminology 2
2.1. Terms Used in This Document 3
2.2. Abbreviations 3
2.3. Requirements Language 3
3. ICMP Error Handling for SRv6-VPNs 3
3.1. Overview 3
3.2. Details of ICMP Error Handling for VPNs in SRv6 Networks 4
3.3. Characteristics of the solution 6
4. Security Considerations 6
5. IANA Considerations 6
6. Acknowledgements 6
7. Normative References 6
Authors' Addresses 7

1. Introduction

Troubleshooting is an essential part of all IP networks. IP VPN service of transport networks always represented a special challenge to provide ICMP error handling, as the hosts in a VPN are not part of the transport network.

For VPNs the main challenge to use ICMP for connectivity check and fault localization is that network internal nodes (a.k.a. P routers) are not service (i.e., VPN) aware. Therefore, P routers cannot route VPN specific ICMP error messages back to the original source of the packet, that triggered the generation of the ICMP error message. Furthermore, in SRv6 networks the P routers may be IPv6-only (i.e., may not support the protocol (e.g. IPv4) or address space used by the VPN clients).

The Uniform Model defined in [RFC3443] allows visibility of traversed nodes outside the provider network.

This document proposes a solution that is capable to allow (1) ping or trace for VPN endpoints with transport network visibility; and (2) find broken link or node within the transport network (e.g., SRv6 domain).

2. Terminology

2.1. Terms Used in This Document

This document uses the Segment Routing terminology established in [RFC8402] and in [RFC9252]. The reader is assumed to be familiar with those documents and their terminology.

2.2. Abbreviations

The following abbreviations are used in this document:

SID Segment Identifier.

SRH SRv6 header.

SRv6 Segment Routing over IPv6.

2.3. 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.

3. ICMP Error Handling for SRv6-VPNs

3.1. Overview

While a solution for diagnostics in MPLS VPNs has been created, the solution created for MPLS based VPN ping or traceroute has many inherited drawbacks. MPLS technology has its special encapsulation, i.e., the MPLS header is a label stack. In case of MPLS, P routers have no options to identify the ingress of the MPLS tunnel, as labels in the header point towards the network egress point. This characteristic restricts the possible solutions to provide VPN specific ICMP handling in MPLS networks.

IPv6 encapsulation used by SRv6 has a srcIP field referring to the originator of the IP packet, i.e., the ingress endpoint of the SRv6 tunnel. Therefore the MPLS restriction does not have to apply for SRv6 networks. The solution described here takes advantages from this presence of the ingress endpoint information to provide an optimal method and to be fully inline with [RFC4443].

Node functions in the described method are as follows:

1. Ingress PE (ingress node of the SRv6 tunnel): VPN packet encapsulation follows [RFC3443] Uniform model. The node adds VPN specific information to the encapsulated packet (i.e., srcIP=VPN-specific-SID of the Ingress PE) and forwards it over the SRv6 network.

2. P node = Originator of the ICMP error (within the SRv6 domain): it does standard [RFC4443] operation, so an ICMP error message is sent to the originator (i.e., the ingress PE) of the SRv6 encapsulated packet, that caused the ICMP message generation (e.g., as the Hop Limit of the packet expired).

3. Ingress PE: it processes the ICMP error message and forwards it to the original source of the (passenger) packet (located within the VPN context). This function is called as VPN-associated-ICMP-leak-function and is described in detail in Section 3.2.

3.2. Details of ICMP Error Handling for VPNs in SRv6 Networks

Figure 1 shows the reference topology used to describe the ICMP error handling. Packet processing works as follows:

1. HostA sends packet to HostB.

2. PE1 encapsulates the packet in an SRv6 tunnel (Uniform model used). The srcIP of the encapsulation is a VPN specific SID of PE1.

3. Encapsulated packet reaches P2 where Hop Limit expires.

4. P2 generates an ICMP Error Message and sends it to PE1, using the VPN specific SID as a dstIP.

5. PE1 process the ICMP Error Message according to its VPN-associated-ICMP-leak-function and identifies the related VPN instance.

6. PE1 send the processed ICMP Error Message to HostA.

7. HostA is informed about the Hop Limit expire event and its network location (i.e., P2).

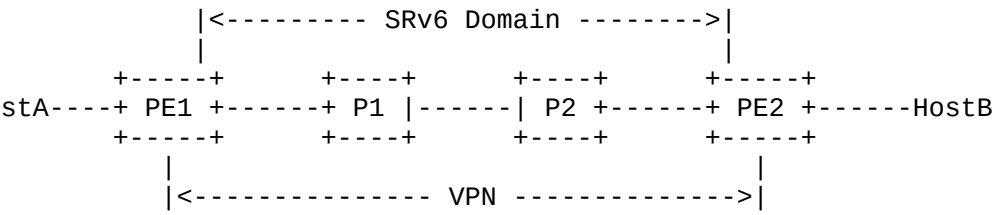

Figure 1: ICMP Error Handling for VPNs in SRv6 Networks: Reference Topology

The VPN-associated-ICMP-leak-function operation contains the followig steps:

1. It processes the received ICMP error message (originated e.g., from a P node within the SRv6 domain).

2. It identifies the related VPN, based on the VPN specific srcIP value in the SRv6 encapsulation of the received ICMP error message.

3. It modifies the ICMP error message:

   * It removes the SRv6 domain specific encapsulation/header(s) of the received ICMP error message.

   * It identifies the VPN specific source of the original packet that caused the ICMP error message, based on the invoking packet header part of the ICMP error message payload.

   * It removes the SRv6 domain specific header(s) from the invoking packet header part of the ICMP error message payload.

   * It creates a new header for the ICMP error message, where srcIP=SRv6Originator-of-the-ICMP-error-message and dstIP=SourceIP-of-the-invoking-packet.

4. Forwards the modified ICMP error message according to the local VPN routing table (vrf).

In case of an IPv4-VPN service, a translation of involved IP addresses is needed (between the related IPv6 and IPv4 addresses).

The VPN-associated-ICMP-leak-function may translate the IP address of the SRv6Originator-of-the-ICMP-error-message (e.g., a P node) to limit the VPN specific visibility characteristics. For example, if the SRv6 domain operator does not want to export the real SID values of the domain nodes.

3.3. Characteristics of the solution

ICMP Error Handling for VPNs in SRv6 Networks has the following characteristics:

* It is compliant to existing standards, like [RFC4443].

* It eliminates the shortcomings of the MPLS based solutions, as (1) it works in case of failures between ingress-PE and egress-PE and (2) it supports direct localization of failures.

* It does not result in additional complexity on P nodes.

* It makes P nodes service agnostic and allows building IPv6-only core networks.

* It uses a SID as a source address.

* It can hide the SIDs used inside the SRv6 domain and can provide different visibility for served VPNs if needed.

4. Security Considerations

This document does not impose any additional security challenges to be considered beyond the security threats described in [RFC9252].

5. IANA Considerations

This document makes no IANA requests.

6. Acknowledgements

Authors extend their appreciation to Janos Farkas and Ferenc Fejes for their insightful comments and productive discussion that helped to improve the document.

7. 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>.

[RFC3443] Agarwal, P. and B. Akyol, "Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks", RFC 3443, DOI 10.17487/RFC3443, January 2003, <https://www.rfc-editor.org/info/rfc3443>.

[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>.

[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>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <https://www.rfc-editor.org/info/rfc8402>.

[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>.

Authors' Addresses

Balazs Varga
Ericsson
Email: balazs.a.varga@ericsson.com

Joel Halpern
Ericsson
Email: joel.halpern@ericsson.com