

ICMP Error Handling for VPNs in SRv6 Networks (for VPN ping/trace ...)



Authors: Balázs Varga, Joel Halpern
Draft: draft-varhal-6man-icmp-srv6-vpn

IETF-125: IPv6 Maintenance

2026-03-18

Overview

draft-varhal-6man-icmp-srv6-vpn



- We propose:
 - This draft specifies ICMP error handling in SRv6-based Virtual Private Networks.
 - The solution only changes ingress-PE behavior.
- Goal:
 - Provide transport network visibility for VPN endpoints (i.e., via ping/trace).
 - Support finding broken link/node within the transport network (i.e., SRv6 domain).

The VPN ping/trace challenge

P nodes are not VPN aware



- Challenge:
 - P nodes are not VPN aware (can not route VPN specific ICMP error messages)
 - P nodes may be IPv6-only (no IPv4 support)
- Goal:
 - Provide transport network visibility for VPN endpoints (i.e., via ping/trace).
 - Support finding broken link/node within the transport network (i.e., SRv6 domain).
- Essence of the solution :
 - Defines an ICMP processing function on the edge node (ingress-PE node)
 - Egress-PE is NOT involved! We try to find forwarding issue(s) to egress-PE node ...

Note on history ...



- Problem space already discussed in detail for MPLS networks (for example):
 - MPLS technology has its special encapsulation, i.e., the MPLS header is a label stack.
 - In case of MPLS,
 - P nodes have no options to identify the ingress of the MPLS tunnel, as labels in the header point towards the egress point.
 - This characteristic of MPLS encapsulation restricted the possible solutions to provide VPN specific ICMP handling.
- SRv6 differs:
 - IP header has a srcIP field referring to the originator of the IP packet, i.e., the ingress SRv6 tunnel endpoint, therefore the **MPLS restriction does not have to apply for SRv6 networks**. This impacts the possible solution space ...




The concept:

Sending ICMP Error to the ingress-PE with "ICMP-process-function"

- Node roles of the solution
 1. Ingress-PE: VPN packet encapsulation follows RFC3443 Uniform model + adds VPN specific information to the encapsulated packet (srcIP=VPN-specific-SID of the Ingress PE)
 2. P node = Originator of the ICMP error (within the SRv6 domain): standard RFC4443 operation, ICMP error message is sent to the originator of the encapsulated packet (i.e., ingress-PE)
 3. Ingress-PE: processes of the ICMP error message and forwards it to the original source of the (inner) packet (located within the VPN)
- We propose:
 - An ICMP error message processing function on ingress-PE (i.e., ICMP-process-function).
 - A method to modify and forward the ICMP Error Message in a VPN specific way.


"ICMP-process-function" based solution

Overview



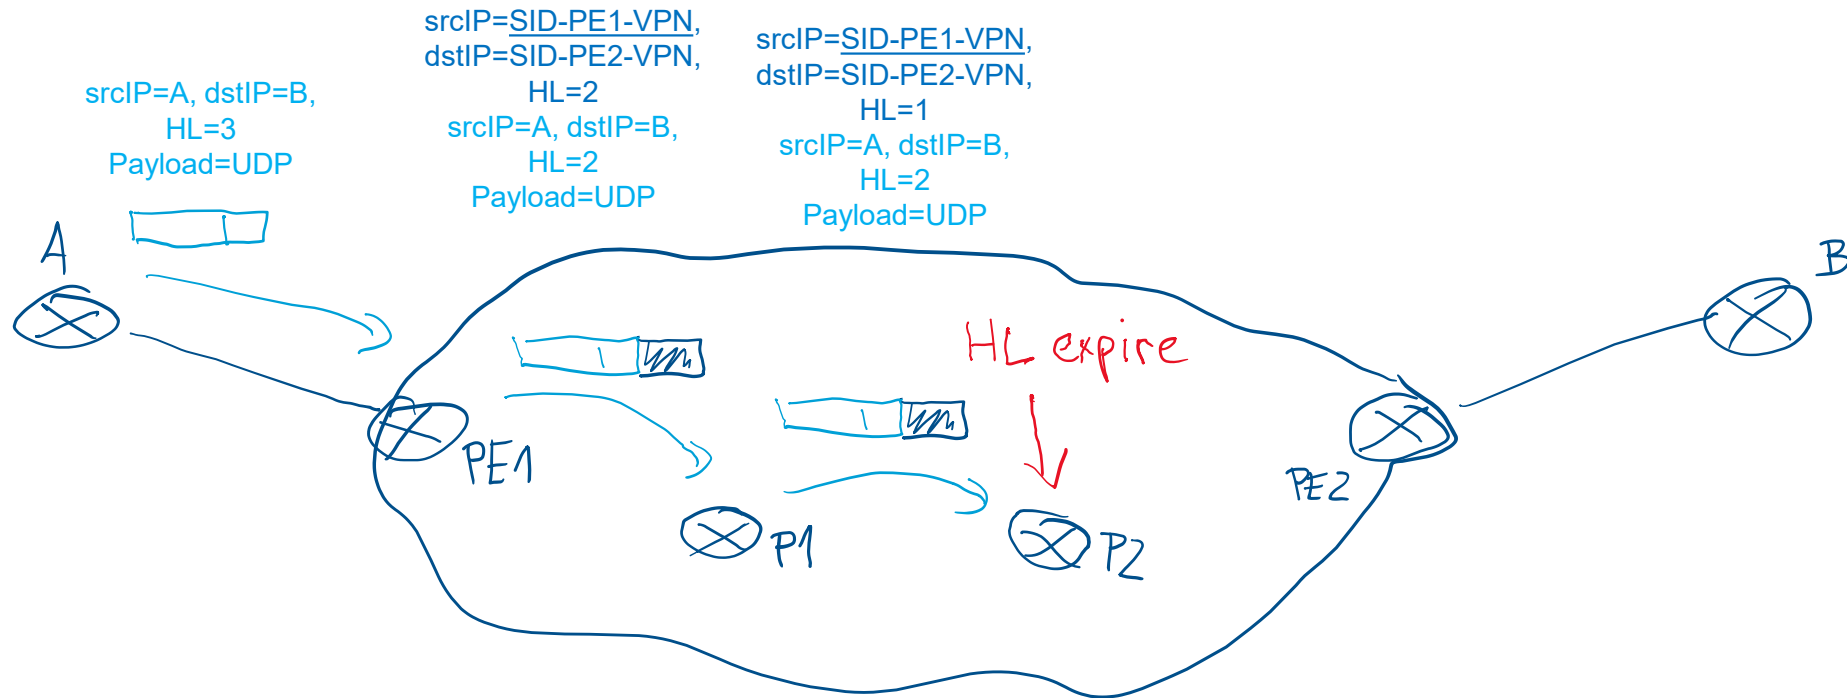
"ICMP-process-function" based solution

Overview



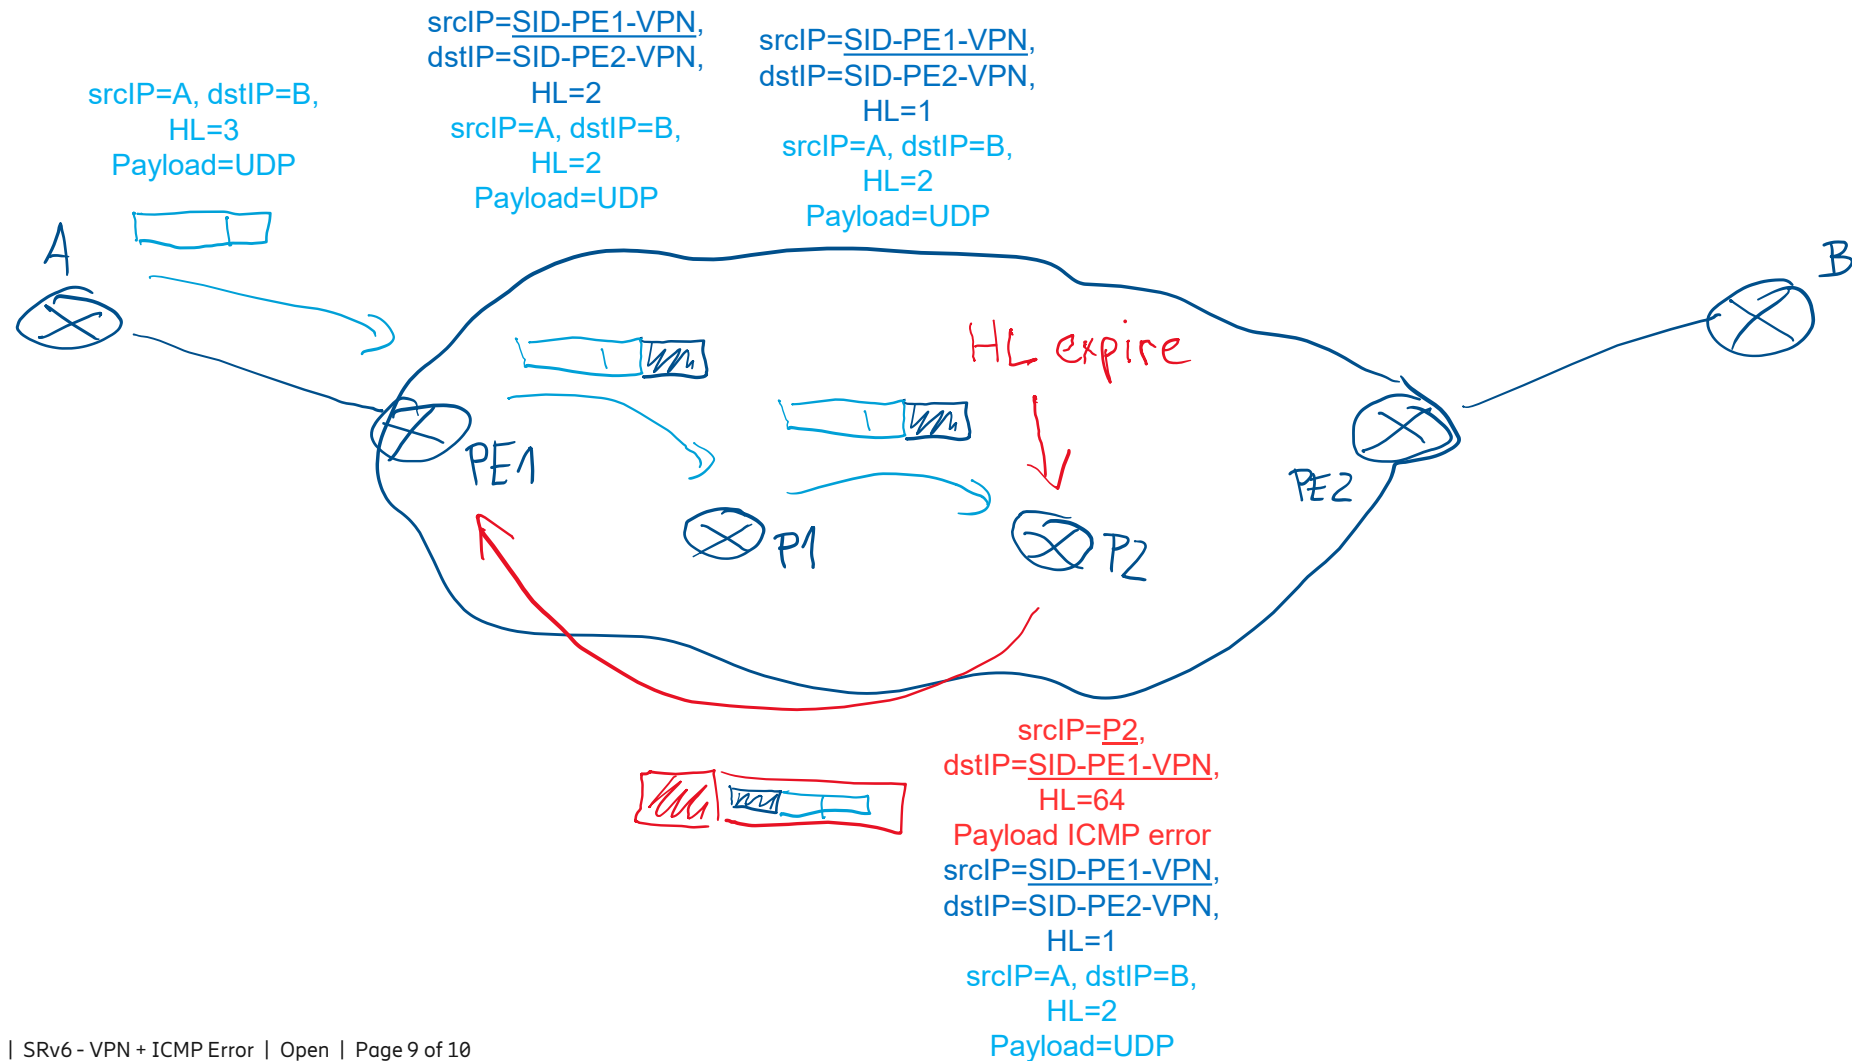
"ICMP-process-function" based solution

Overview



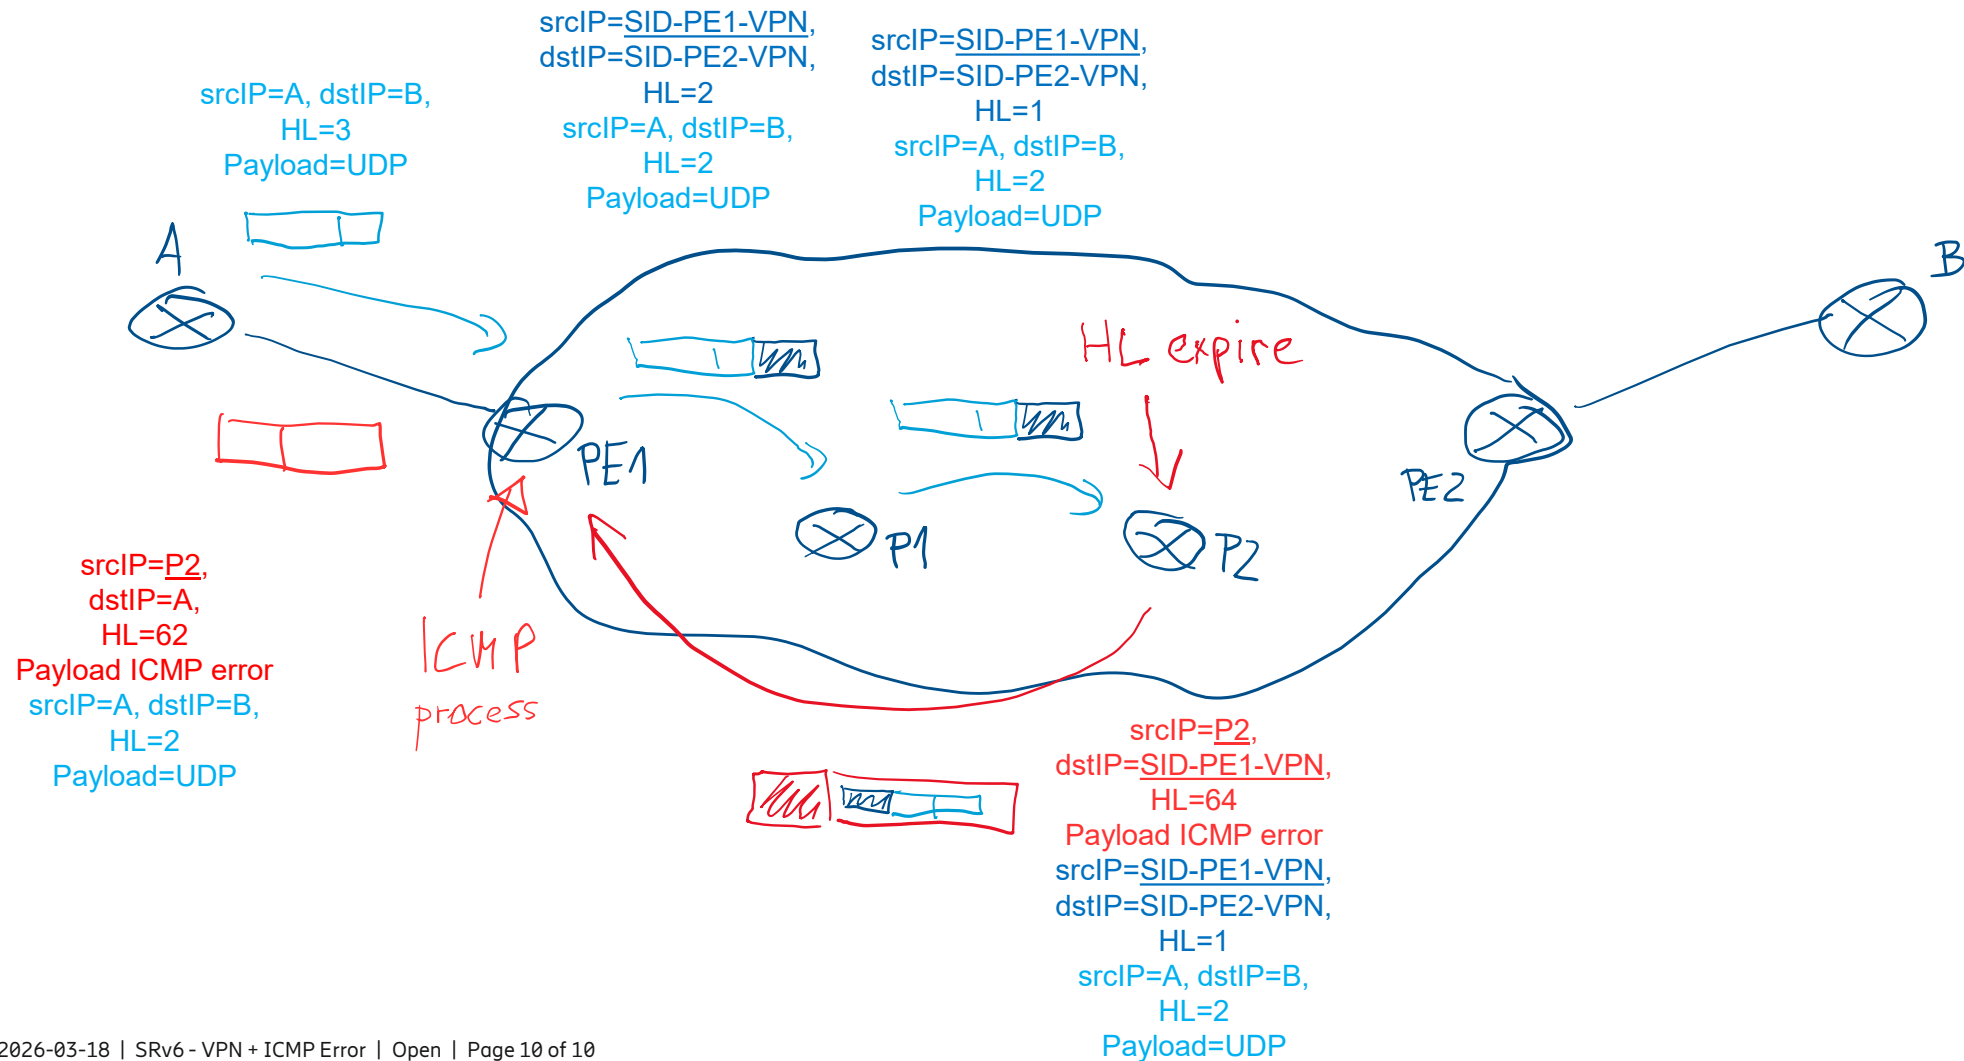
"ICMP-process-function" based solution

Overview



"ICMP-process-function" based solution

Overview



“ICMP-process-function” operation

Details ...



- ICMP-process-function on ingress-PE node:
 - A. Process the received ICMP error message (originated e.g., from a P node within the SRv6 domain).
 - B. Identifies the related VPN, based on the VPN specific srcIP of the SRv6 encapsulation (= srcIP of the received ICMP error message).
 - C. Modifies the ICMP error message
 - 1) Removes the SRv6 domain specific encapsulation(s)/header(s) of the received ICMP error message.
 - 2) 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.
 - 3) Removes the SRv6 domain specific header(s) from the “invoking packet header” part.
 - 4) Creates a new header for the ICMP error message, where srcIP=Originator-of-the-ICMP-error-message, dstIP=sourceIP-of-the-invoking-packet.
 - D. Forwards the modified ICMP error message according to the local VPN routing table (vrf).
- Note: In case of IPv4-VPN the IPv4 address of the Originator node is used OR the rules of RFC7600 (Section 4.8) applies (+ draft-ietf-intarea-extended-icmp-nodeid).
- Optional: ICMP-process-function may translate the IP address of the “Originator-of-the-ICMP-error-message” to limit the VPN specific visibility characteristics (e.g., the SRv6 domain owner does not want to export the real SID values of the domain nodes).

Benefits of the Solution

Supporting direct localization of failures ...



- It is compliant to existing standards, like RFC4443.
- No additional complexity on P nodes, no involvement of egress-PE.
- It eliminates the shortcomings of an MPLS like solutions.
 - It works in case of failures between ingress-PE and egress-PE.
 - It supports localization of failures.
- It makes P nodes service agnostic. It allows building IPv6-only core networks.
- It can hide IP addresses used inside the SRv6 domain. It can provide different visibility for served VPNs.

Next-steps

More discussions



- Looking for feedbacks:
 - Discussion already exists on the 6man (+spring) mailing list. Thanks ...
- Encouraging further discussion on:
 - Scenarios we intend to cover (e.g., complex multi-level encapsulations like TI-LFA).
 - Evaluation criteria (e.g., keep it simple).
- Note:
 - There is an alternative proposal, that changes P node behavior and involves egress-PE node in forwarding ICMP error messages: [draft-ali-6man-srv6-vpn-icmp-error-handling](#)