

IESG LC: BFD for VXLAN

draft-ietf-bfd-vxlan

Santosh Pallagatti
Sudarsan Paragiri
Vengada Prasad Govindan
Mallik Mudigonda
Greg Mirsky

IETF-105 July 2019, Montreal

Status

- Comments from:
 - Erik Kline (Gen-ART)
 - Jürgen Schönwälder (Ops-Dir)
 - Shawn M. Emery (Sec-Dir)
 - Joel Halpern (RTG-Dir)
 - Olivier Bonaventure (TSV-ART)
 - Carlos Pignataro
 - Dinesh Dutt
 - T.Sridhar
 - Sam Aldrin

Many thanks to all for their reviews, comments, suggestions.
And, the most of, for their patience.

Resolved comments

- Added NVE, VNI to Terminology
- Merged Section 4.1 into Section 4 (we used to have Section 4.2 but not anymore)
- Multiple BFD sessions between a pair of VTEPs:

OLD TEXT:

The implementation SHOULD have a reasonable upper bound on the number of BFD sessions that can be created between the same pair of VTEPs.

NEW TEXT:

If the implementation supports establishing multiple BFD sessions between the same pair of VTEPs, there SHOULD be a mechanism to control the maximum number of such sessions that can be active at the same time.

- Removed request to allocate the dedicated MAC address – not using the dedicated MAC address as an option in the destination MAC address of the inner Ethernet frame
- Added description of Tenant-VTEP-VTEP-Tenant e2e monitoring:
 - single BFD session – VTEPs are transit nodes as any other tenant-addressed data
 - concatenated BFD sessions – follow Section 6.8.17 RFC 5880

Open questions

- Destination MAC address of the inner Ethernet frame:
 - Remove using the dedicated MAC address
 - Can a MAC of the remote VTEP be used?
 - Seems as “borrowing” address from a tenant space
 - Inconclusive responses from VXLAN experts – 1:1
 - Suggested alternatives:
 - Use Management VNI, e.g., VNI 1 (similar to VLAN 1)
 - Allocate OAM flag – going towards VXLAN-GPE?
- Demultiplexing BFD control packet with Your Discriminator == 0
OLD TEXT:

For such packets, the BFD session MUST be identified using the inner headers, i.e., the source IP, the destination IP, and the source UDP port number present in the IP header carried by the payload of the VXLAN encapsulated packet.

NEW TEXT:

For such packets, the BFD session MUST be identified using the source UDP port number present in the IP header carried by the payload of the VXLAN encapsulated packet.
- Track:
 - RFC 7348 is Informational – should BFD over VXLAN be Informational or Standard?
- Echo BFD in the document:
 - Interest to add to the document?
 - Leave the current text that states that Echo BFD is outside the scope or remove and add to Introduction section:

This specification describes procedures only for BFD asynchronous mode.

Next steps

- Discuss
- Decide
- WG LC II?

Thank you