

# IESG LC: BFD for VXLAN

draft-ietf-bfd-vxlan

Santosh Pallagatti  
Sudarsan Paragiri  
Vengada Prasad Govindan  
Mallik Mudigonda  
Greg Mirsky

IETF-106 November 2019, Singapore

# Scope of the discussion

- Echo BFD mode for BFD over VXLAN
  - BFD Echo is outside the scope of this specification of BFD over VXLAN
- Selection and use of the destination MAC address in the inner Ethernet header
- Use of an IP address from the Internal host loopback address range as the destination IP address in the inner IP header
- Number of BFD sessions between the pair of VTEPs
- Use of Management VNI for BFD over VXLAN

# Destination MAC Address

- The choice of Destination MAC and Destination IP addresses for the inner Ethernet frame **MUST** ensure that the BFD Control packet is not forwarded to a tenant but is processed locally at the remote VTEP
- Not using the dedicated (assigned by IANA) MAC address
- Destination MAC: This **MUST NOT** be of one of tenant's MAC addresses. The destination MAC address **MAY** be the address associated with the destination VTEP. The MAC address **MAY** be configured, or it **MAY** be learned via a control plane protocol. The details of how the MAC address is obtained are outside the scope of this document
- VTEP **MUST** validate the packet. If the Destination MAC of the inner Ethernet frame matches one of the MAC addresses associated with the VTEP, the packet **MUST** be processed further. If the Destination MAC of the inner Ethernet frame doesn't match any of VTEP's MAC addresses, then the processing of the received VXLAN packet **MUST** follow the procedures described in Section 4.1 [RFC7348]. If the BFD session is using the Management VNI (Section 6), a BFD Control packet with unknown MAC address **MUST NOT** be forwarded to VMs.

# Destination IP Address

- The inner destination IP address SHOULD be set to one of the loopback addresses (127/8 range for IPv4 and 0:0:0:0:0:FFFF:7F00:0/104 range for IPv6)
- It is RECOMMENDED to allow addresses from the loopback range through a firewall only if it is used as the destination IP address in the inner IP header, and the destination UDP port is set to 3784

# Number of BFD Sessions and Management VNI

- Multiple BFD sessions may be running between two VTEPs. For demultiplexing packets with Your Discriminator equal to 0, a BFD session MUST be identified using the logical link over which the BFD Control packet is received. In the case of VXLAN, the VNI number identifies that logical link.
- The administrator MUST be able to select VNI number for a BFD session between two VTEPs
- An implementation MAY support the use of the Management VNI as control and management channel between VTEPs. The selection of the VNI number of the Management VNI MUST be controlled through a management plane. An implementation MAY use VNI number 1 as the default value for the Management VNI. All VXLAN packets received on the Management VNI MUST be processed locally and MUST NOT be forwarded to a tenant.

# Open questions

- Should the document define the default value for the destination MAC address in the inner Ethernet frame when the Management VNI is used for BFD over VXLAN?

The reason is to avoid the introduction of special per VNI behavior in VXLAN. Instead, the default MAC assigned for the Management VNI. Again, this MAC SHOULD NOT be one of the tenant's MACs.

- Stronger normative language for the use of VNI 1 as the default Management VNI. For example

Any VNI that is not provisioned for tenant systems MAY be used for the Management VNI, as the default VNI 1 is RECOMMENDED.

- The BFD Demand mode is not mentioned in the document explicitly. Is that the indication that it is outside the scope of the draft? If yes, then should that be stated explicitly, as is the case for BFD Echo?

# Next steps

- Discuss
- Decide
- WG LC II?

Thank you