

# BFD for Geneve

draft-xiao-nvo3-bfd-geneve-03

Xiao Min	ZTE
Greg Mirsky	ZTE
Santosh Pallagatti	VMware
Jeff Tantsura	Apstra

# A New Co-author Joins

- Welcome **Jeff Tantsura** to be co-author of this draft

# Summary of main updates

- Remove **BFD-MPLS-Geneve** encapsulation
  - This kind of encapsulation would very likely not to be used within DC
- Resolve one open issue on the mapping between VAP and VNI
  - RFC 8014 allows for N:1 mapping between VAP and VNI at one NVE

# Summary of this draft

- Key Point 1: Geneve BFD session originates and terminates at VAP of NVE
- Key Point 2: VAP and VNI have the relationship of N:1 within one NVE
- Key Point 3: Originating VAP decides Geneve BFD encapsulation and VNI
- Key Point 4: Peer VAP address can be obtained by management/control plane
- Key Point 5: Management VNI and BFD echo function are out of scope

# Key Point 1

- Geneve BFD session originates and terminates at VAP of an NVE
  - VAP MAC/IP address would be used as the inner MAC/IP address
  - If the terminating VAP has no IP address assigned, then set the IP DA as a special IP address (chosen from the 127/8 range for IPv4, and from the ::ffff:127.0.0.0/104 range for IPv6)
  - If the originating VAP has no IP address assigned, currently the IP address of the originating NVE is used, 0.0.0.0 for IPv4 and 0:0:0:0:0:0:0:0 for IPv6 seem more reasonable, any other alternatives?

# Key Point 2

- VAP and VNI have the relationship of N:1 within one NVE
  - Multiple BFD sessions for one VNI are allowed between a pair of NVEs
  - If the BFD packet is received with Your Discriminator equals to 0, VNI itself is not enough to demultiplex the received BFD packets, MAC/IP address and source UDP port are also needed
  - If the BFD packet is received with non-zero Your Discriminator, then the BFD session would be demultiplexed only by Your Discriminator

# Key Point 3

- Originating VAP decides the used Geneve BFD encapsulation and VNI
  - If the VAP that originates the BFD packets is used to encapsulate Ethernet data frames, then BFD packets are encapsulated using **BFD-Ethernet-Geneve**
  - If the VAP that originates the BFD packets is used to encapsulate IP data packets, then BFD packets are encapsulated using **BFD-IP-Geneve**
  - a BFD session can only be established between two VAPs that are mapped to **the same VNI** and use **the same way** to encapsulate data packets
  - It makes the Geneve OAM packets **fate-sharing** (a key characteristic) with Geneve Data packets

# Key Point 4

- Peer VAP address can be obtained by management or control plane
  - the encapsulation type and address of peer VAP can be obtained by **Netconf**
  - the encapsulation type and address of peer VAP can be obtained by **EVPN**
  - the encapsulation type and address of peer VAP can be obtained by **OVSDB**
  - the encapsulation type and address of peer VAP can be obtained by **OpenFlow**
  - .....

# Key Point 5

- Management VNI and BFD echo function are out of scope
  - Management VNI method uses a special independent VNI to perform OAM related functions
  - Management VNI can be used to achieve Geneve BFD
  - Management VNI can only check whether the Geneve tunnel works for the special VNI, so it's complementary to the method described in this draft
  - Currently BFD echo function is considered out of scope, is there a requirement to support it? And what about demand mode?

# Next steps

- Ask for WG adoption