BFD for VXLAN

(draft-spallagatti-bfd-vxlan-02.txt)

Santosh Pallagatti
Saji Basil
Sudarsan Paragiri
Vengada Prasad Govindan
Mallik Mudigonda
Greg Mirsky

Use case

- VMs OAM aggregation
- Layer 2 VMs.
- Fault localization.
- Service node reachability.
- p2p BFD for now and p2mp for future study.

Deployment

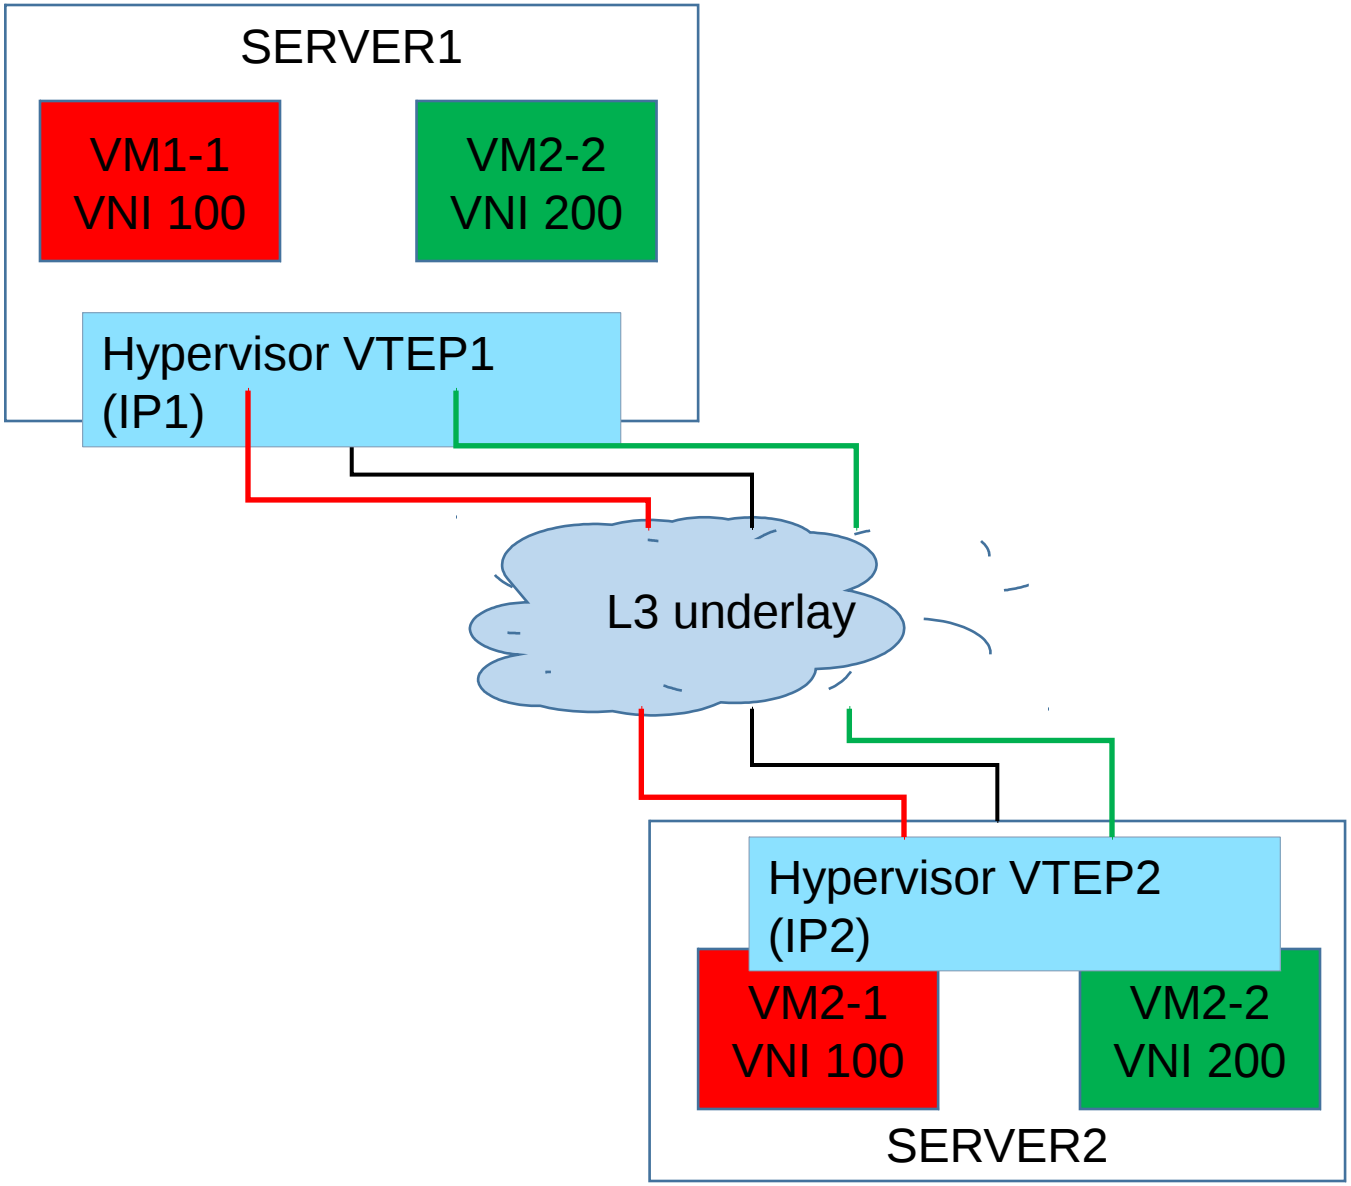
SERVER1
VM1-1
VNI 100
VM2-2
VNI 200
Hypervisor VTEP1
(IP1)
L3 underlay
Hypervisor VTEP2
(IP2)
VM2-1
VNI 100
VM2-2
VNI 200
SERVER2

BFD packet encap

Header	Field	Description
Inner MAC header	Destination MAC	This MUST be a well-known MAC [TBD] OR the MAC address of the destination VTEP.
	Source MAC	MAC address of the originating VTEP.
Inner IP header	Source IP address	IP address of the originating VTEP.
	Destination IP address	IP address of the terminating VTEP.
	TTL	This MUST be set to 1.
GPE header	O bit	MUST be set.

Note: We are still discussing on inner IP destination IP address.

VXLAN path and reserved VNI

- It is expected that implementation should ensure BFD packet would traverse the same path as any other VXLAN packet within the system when
  - BFD packet is sent out from VTEP
  - BFD packet is received on VTEP
- One way to aggregate BFD sessions between VTEP’s is to have BFD session established for VNI 0.
- VNI 0 can also be used to establish BFD session between VTEP and service node.

Reception of BFD packet

- Inner MAC set to well-known or receiving VTEP then packet should not be forwarded to VM.
- Use inner VNI as the key to demultiplex received BFD packet when Your Discriminator in BFD packet is set to 0.
- Reverse path over IP or MAY be directed over another VXLAN tunnel (for future work).

Next steps

- Welcome comments from the WG
- Asking WG to consider adoption of the draft
- Thank you