

Virtual Topologies for Service Chaining in BGP IP MPLS VPNs

(draft-rfernando-l3vpn-service-chaining-03)

Dhananjaya Rao
Rex Fernando
Luyuan Fang
Maria Napierala
Ning So
Adrian Farrel

IETF 88, November 2013
Vancouver

About this draft

- Propose techniques built upon BGP/IP MPLS VPN control plane mechanisms to construct virtual topologies for service chaining.
- The virtual service topologies interconnect network zones and constrain the flow of traffic between these zones via a sequence of service nodes
- Two approach described: routing control plane, and by network orchestration to realize these virtual service topologies.
- 04 hanges: Adrian joined as co-author and updated the draft.

Terminologies

- A network zone: a logical grouping of physical assets that supports certain applications. Hosts can communicate freely within a zone.
- Service-PE: An IP VPN PE to which a service is attached. Direct incoming traffic from other PEs or from attached hosts to the service node via an MPLS VPN label or IP lookup.
- Service node: A physical or virtual service appliance/application. E.g. FWs, LBs, and DPIs. The service node acts as a CE.
- Service chain: A sequence of service nodes that interconnect the zones containing the source and destination hosts. Unidirectional and creates a one way traffic flow between source zone and destination zone.
- Virtual service topology: a sequence of service-PEs and their attached service nodes created in a specific order. A service topology is constructed via one or more routes that direct the traffic flow among the PEs that form the service chain.
- Service-topology-RT: A BGP route attribute that identifies the specific service topology.

Virtual Machine Intra-zone routing

- In a data center, servers host VMs where end applications reside
 - each application VM is a CE from an IP BGP VPN perspective
- A collection of CE/VMs that can communicate freely form a *zone*
 - A PE creates a VRF for its attached CE/VMs in a zone
- Intra-zone connectivity achieved by designating a RT per zone (zone-RT)
 - Applied on all PE VRFs that terminate the CE/VMs that belong to the zone

Inter-zone Routing and Traffic Forwarding

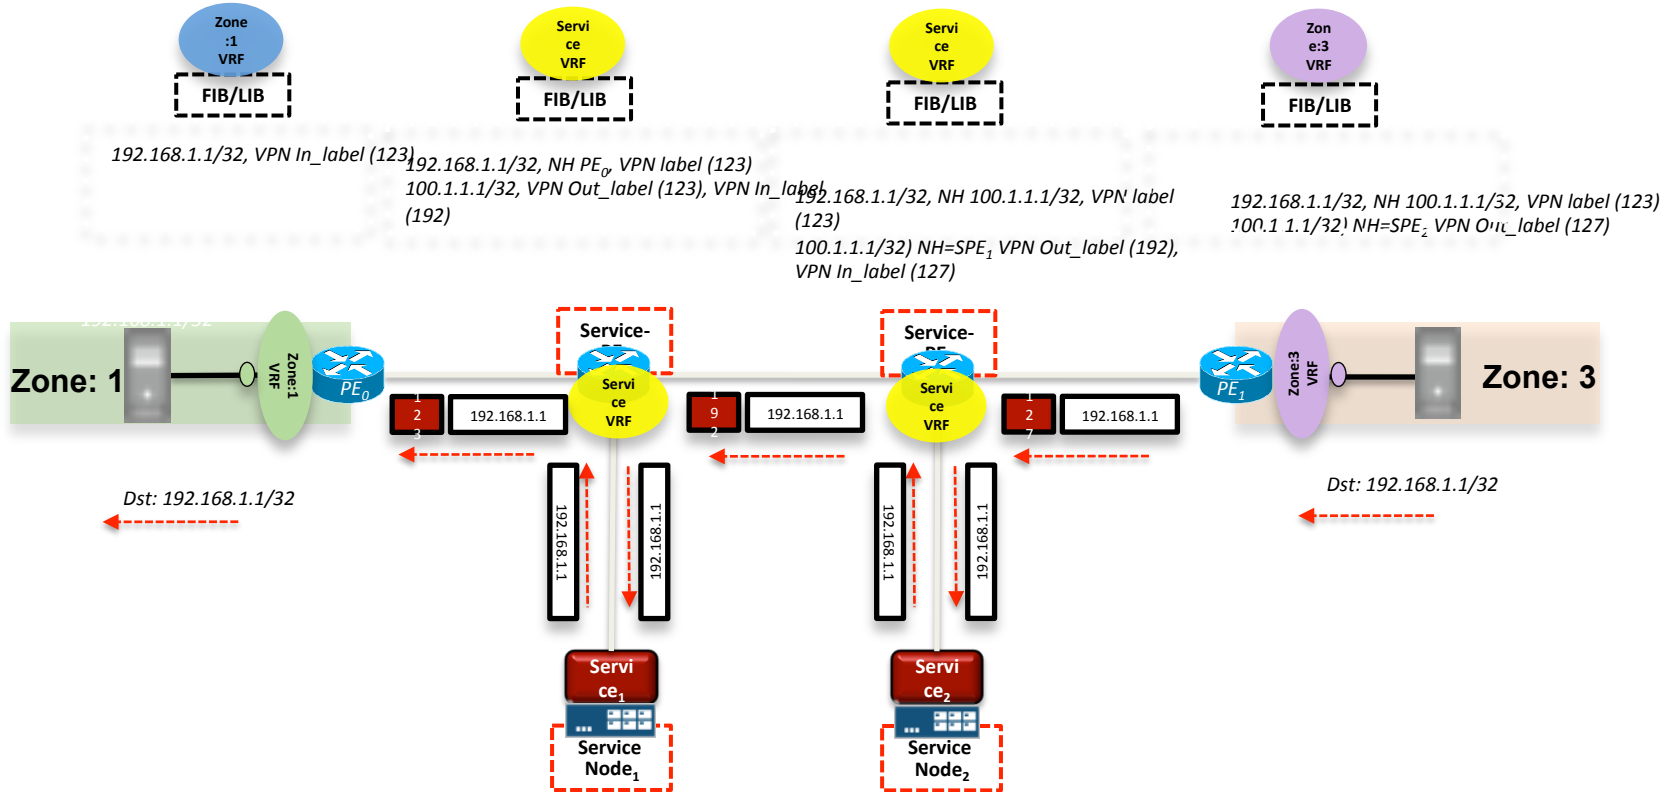
- Apply network policies and services in a specific order
- Service nodes may be VMs spread across the data center
- Inter-zone traffic must follow a predetermined service path and forwarding through one or more service nodes
- A sequence of service-PEs and their attached service nodes creates a unidirectional service chain or topology
- Two step process:
 - Virtual Service Topology construction
 - Inter-zone Routing and Service Chaining

Zone: 1

Zone: 3

Inter-zone Routing

Inter-zone Traffic Forwarding



Orchestration Driven Approach

- Remove the need for the zone or service PEs to determine the appropriate next-hops based on the specified service node sequence.
- The central orchestrator performs the necessary policy computations, and construct the forwarding tables for the various VRFs at the PEs.
- The orchestrator communicates with the various PEs (typically virtual PEs on the end-servers) to populate the forwarding tables.

Next Steps

- Add Nabil back to the draft, as orchestration approach has been added
- More feedbacks welcome
- Ready to ask for WG adoption and quickly progress to last call