Virtual Topologies for Service Chaining in BGP IP/MPLS VPNs

draft-rfernando-bess-service-chaining-00
(previously draft-rfernando-l3vpn-service-chaining-04)

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

About this draft

- Describes the concept of Virtual Service Topologies in BGP based VPN domains
- Explains functions at PEs to efficiently constrain routing and traffic flow over these topologies
- Uses standard BGP IP/MPLS VPN constructs and procedures for control and data plane at PEs
- Describes traditional routing and network controller based approaches
- Proposes incremental extensions to PE control plane behavior to facilitate traditional approach

Draft History

- First presented at IETF 85 in Atlanta (November 2012)
- Added additional co-authors
- Incorporated couple of rounds of comments and feedback
- Changed from Informational to Standards Track
  - Proposed incremental change in PE local VPN procedures
  - No protocol changes
- Version -04 presented at Toronto
  - Received good support for WG adoption
- Current version posted as draft-rfernando-bess-service-chaining-00

Virtual Service Topology

- A virtual network topology is divided into network zones (source, destination)
- Inter-zone traffic may need application of network policies and set of services
- Service nodes may be VMs spread across a physical network, attached to various PEs
- Inter-zone traffic must follow a specific service path and forwarding through a sequence of one or more service nodes
- The sequence of service nodes creates a unidirectional service chain in the topology

Topology Illustration

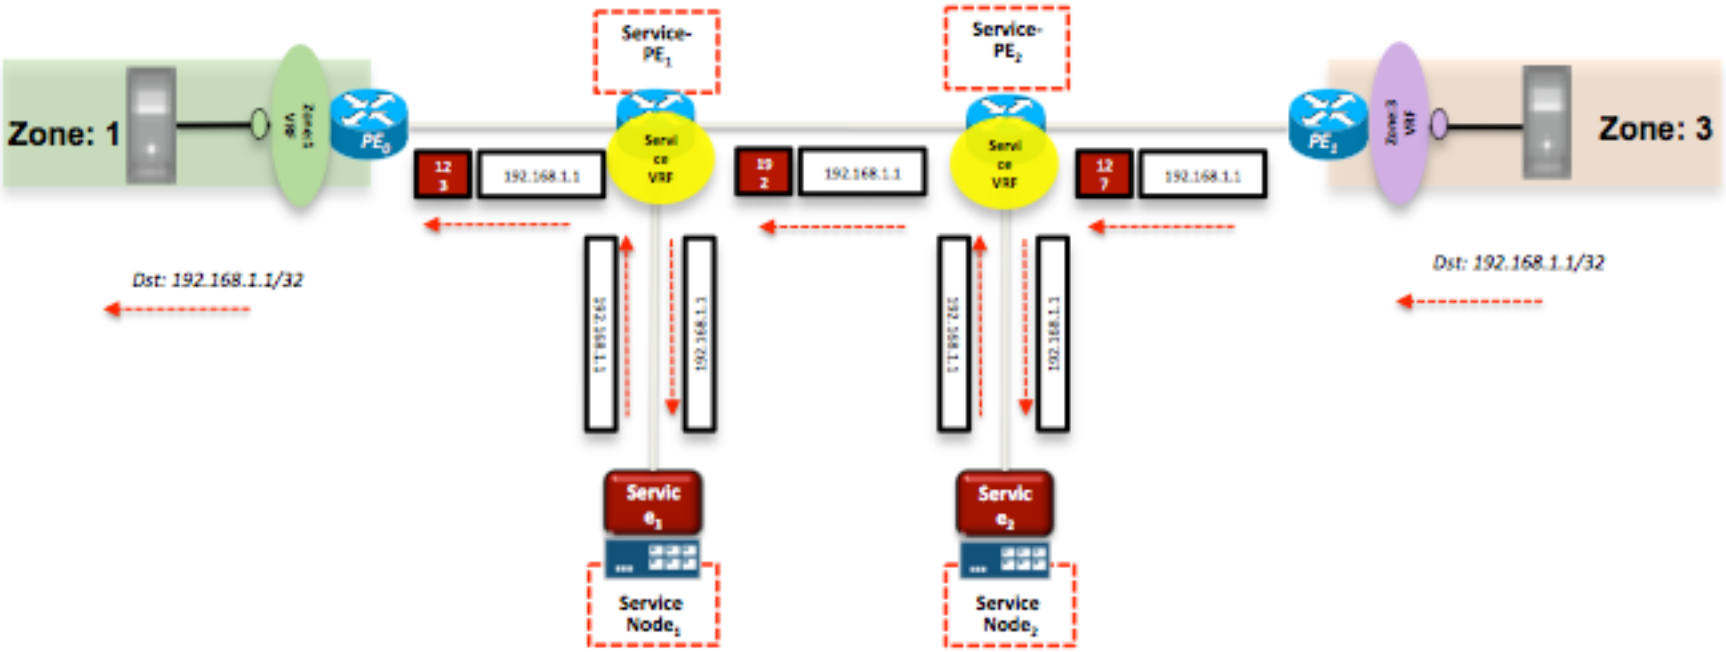

Virtual Topology Routing/Forwarding Steps

- Service Chain provisioning
  - <Service-topology-RT, Service-node-sequence>
    - Service-node-sequence: List of IP addresses
  - Provisioned/signaled to relevant set of PEs for a VPN
- Service node reachability
  - Route to a service node's IP address originated by attached PE as a VPN route
  - Each node serves as a Service-next-hop for the actual inter-zone (VPN) prefixes
- Service Topology Next-hop Resolution
  - Zone prefixes advertised with Service-topology-RT attached
  - PEs hosting service nodes & source zone resolve routes via appropriate Service-next- hop

Virtual Topology Forwarding

- Steer traffic into service chain
  - Destination based
  - Flowspec based

- Traffic flow through service chain
  - VPN (virtual network) forwarding
  - Multiple encapsulations supported

Deployment options

- BGP VPN signaling, PE mechanisms
  - Leverages PE VPN BGP route resolution
- Network controller driven approach
  - Topology resolution decisions taken at controller
  - Prescriptive route entries signaled to PEs
  - Signaling uses MP-BGP and Netconf/Restconf

PE behavior modifications

- Source zone, service PEs must ignore VPN label and next-hop carried with inter-zone routes exported by destination PEs
- Instead, resolve next-hop in VRF using appropriate service node in sequence
- Currently specified as a local PE procedure change
  - Could define protocol extension

Next Steps

- WG adoption call was planned at last IETF
- Resolve differences with draft-mackie ?