

BGP-LS Filters

A Framework for Network Slicing and Enhanced VPNs

draft-drake-bess-enhanced vpn-02

Adrian Farrel (adrian@olddog.co.uk)

John Drake (jdrake@juniper.net)

Luay Jalil (luay.jalil@verizon.com)

Avinash Lingala (ar977m@att.com)

Network Slicing/Enhanced VPN

- Described in draft-ietf-teas-enhanced-vpn
- A set of VPN overlay networks over a common underlay network
 - Different performance and scaling properties
 - Each requires a subset, either dedicated or shared, of the underlay network resources
 - First proposed for 5G mobile networks (network slicing)
- Requires a mechanism to inform each VPN of the underlay network resources (links and nodes) it can use
- Each VPN uses only those resources
 - Segment routing or RSVP-TE
- Approach is a generalization of SR TE policy
 - Supports MP2MP, P2MP, and P2P bi-directional topologies
 - Integrates BGP-VPN support (L3VPN, EVPN)
 - DSCP and Color based forwarding

Approach

- Network Controller performs global optimization across all enhanced VPNs
 - [Source, Destination] traffic matrix for each VPN
 - Performance and scaling properties for each VPN
 - Underlay network topology (BGP-LS northbound)
- Network controller assigns P router link bandwidth to DSCP or adjacency SID
- Each enhanced VPN gets a subset of the underlay network resources
 - BGP-LS Filter
 - PEs have a complete view of underlay network resources, BGP-LS Filter acts as a filter
 - Encoded w/ BGP-LS (southbound)
- Granularity
 - All PEs in a VPN
 - Set of PEs in a VPN
 - Individual PE in a VPN
- Distribute BGP-LS Filter to PEs using standard VPN tools
 - Route targets (RTs), route reflectors, RT constraints

Approach – Continued


- BGP-LS
 - Integrates easily /w VPNs
 - Supports segment routing and RSVP-TE
 - Supports inter-AS
 - Supports both IS-IS and OSPF
 - No new invention required
- Follows VPN scaling model
 - Per-VPN information restricted to its PEs
 - P-routers have no per-VPN information
- Non-enhanced VPN PEs don't receive BGP-LS Filters
 - Network controller ensures underlay network resources for enhanced VPNs have higher IGP and TE metric
 - Enhanced VPNs may use non-enhanced VPN underlay network resources

BGP-LS Filter

- Provides connectivity between the PEs in a VPN or set of VPNs
- Should provide multiple paths between a given ingress/egress PE pair
- Tagged with the RTs of the VPNs whose PEs are to import it
- When sending a packet an ingress PE
 - Selects a path to the egress PE using the BGP-LS Filters for the packet's VPN
 - Adds a segment routing header to the packet *OR*
 - Places the packet in an RSVP-TE LSP using that path
 - Can use any path computation it wishes as long as path is constrained
- An underlay network link or node may appear in multiple BGP-LS Filters
 - Requires AFI (BGP-LS) and SAFI (BGP-LS-VPN) w/ a different RD for each BGP-LS Filter
- Requires a new attribute for a BGP-LS UPDATE

BGP-LS Filter example

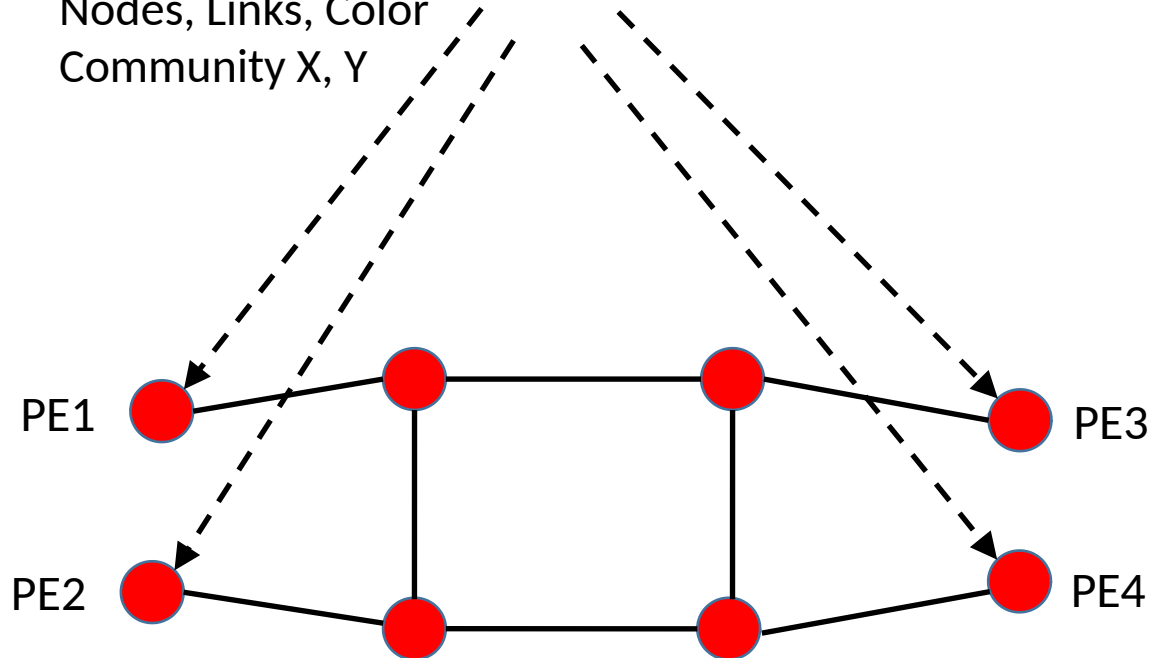
Base Network Topology



BGP-LS from Controller:

Filter filter1:

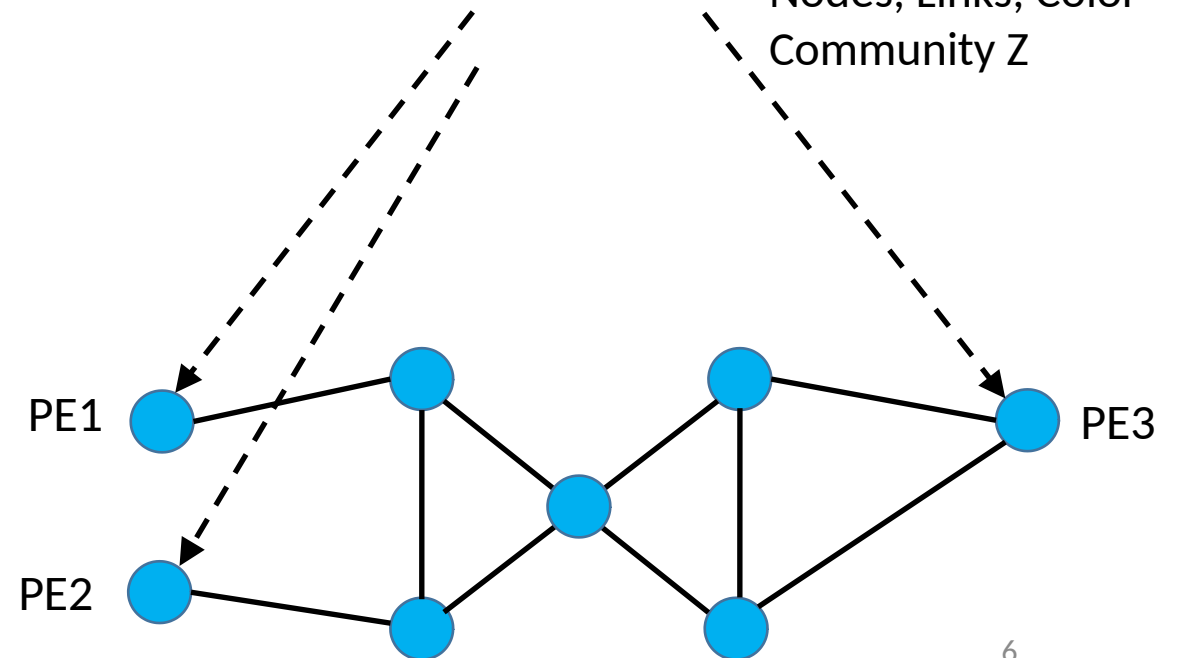
Nodes, Links, Color
Community X, Y



BGP-LS from Controller:

Filter filter2:

Nodes, Links, Color
Community Z



BGP-LS Filter Attribute

- Assumptions:
 - A PE may import more than one BGP-LS Filter
 - A BGP-LS Filter may change over time
 - A given BGP-LS Filter may contain multiple BGP-LS UPDATES
 - The BGP-LS UPDATES for a given BGP-LS Filter may arrive in any order
 - The BGP-LS UPDATES for multiple BGP-LS Filters may be interleaved
- This requires:
 - A Filter identifier
 - A version number
 - A Filter type (MP2MP, P2MP, P2P (uni- or bi-directional))
 - Total # of fragments
 - Fragment # of the current fragment
- Also may contain a DSCP and/or a Color List
 - If present, indicates which DSCPs and/or Colors may use it

Details

- Cannot use a BGP-LS Filter until it is completely assembled
- Underlay network resource can only be used if active
- Underlay network topology is the union of all imported BGL-LS Filters
- Filter precedence: P2P, P2MP, MP2MP
- New version of BGP-LS Filters when
 - The PE membership of a given VPN changes
 - Network controller must track PE membership
 - The underlay network topology changes
 - Not required if connectivity still exists
- Enhanced VPNs may use non-enhanced VPN underlay network resources when:
 - A new PE is activated
 - No connectivity to an egress PE
- Should be corrected when new BGP-LS Filters are received

Plans

- Check back with draft-ietf-teas-enhanced-vpn
- Work with authors of related drafts...
 - draft-dong-spring-sr-for-enhanced-vpn
 - draft-dong-teas-enhanced-vpn-control-plane
 - draft-dong-idr-bgppls-sr-enhanced-vpn

... We believe that the concepts are similar and that our draft is a functional superset
- Iterate the draft
- Discuss with implementers