BGP Usage for SDWAN Overlay Networks

draft-dunbar-bess-bgp-sdwan-usage-03

Linda Dunbar & Jim Guichard (Futurewei)
Ali Sajassi (Cisco)
John Drake (Juniper)
Ayan Barnerjee & Dave Carrel (Cisco)

Nov 2019

BGP Walk Through for Homogeneous SD-WAN multiple routes aggregated in one IPsec

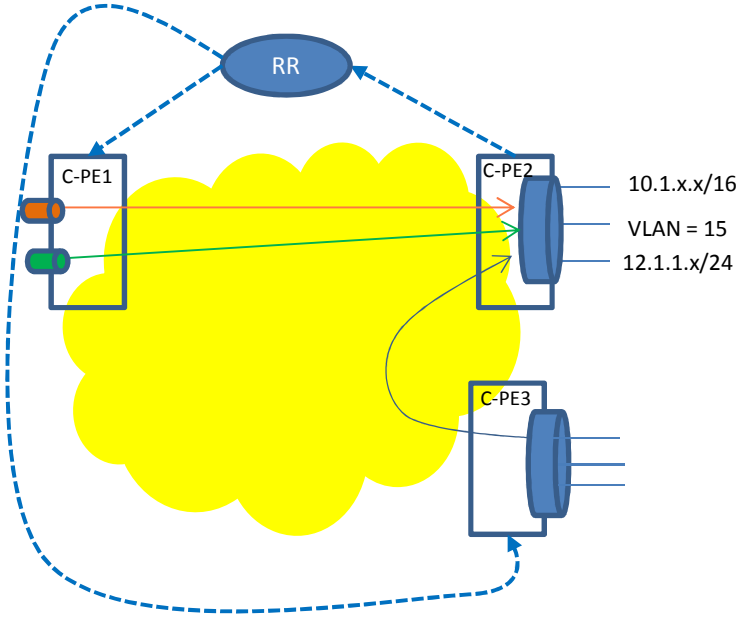

One BGP UPDATE Message from C-PE2 to RR:

- multiple routes encoded in the MP-NLRI Path Attribute
  - 10.1.x.x/16
  - VLAN #15
  - 12.1.1.x/24
- IPsec attributes are encoded in the Tunnel-Encap Path Attribute
  - IPsec attributes for all possible remote nodes, or
  - IPsec attributes for specific remote nodes, or
  - IPsec attributes for specific remote subnets
  - ….

BGP Walk Through for Homogeneous SD-WAN per Route Encryption (Fine-Grained)

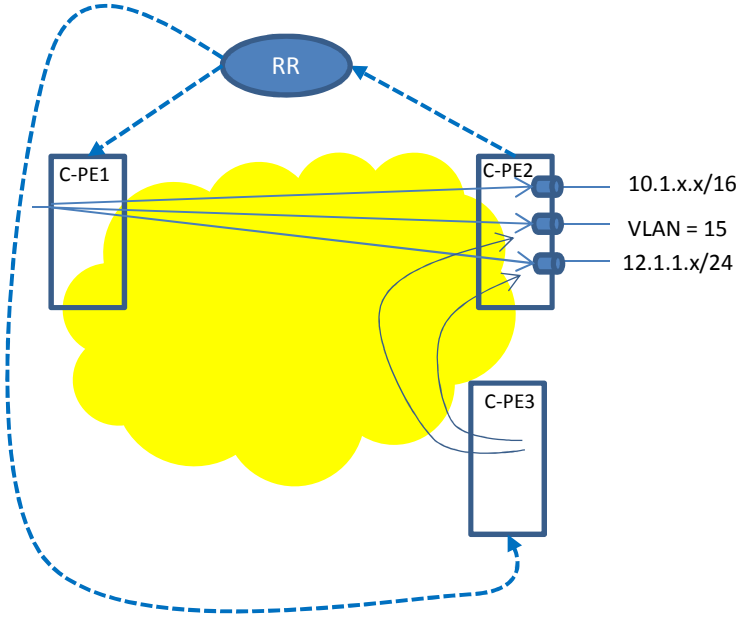

Three Separate BGP UPDATE messages from C-PE2 to RR:

UPDATE 1:

- MP-NLRI Path Attribute
  - 10.1.x.x/16 encoded
- Tunnel-Encap
  - IPsec SA attributes for C-PE1->C-PE2 & C-PE3-> C-PE-2
  - Optional: IPsec attributes for tunnel from specific Subnet A on C-PEx

UPDATE 2:

- MP-NLRI Path Attribute
  - VLAN #15
- Tunnel-Encap:
  - IPsec SA attributes for C-PE1->C-PE2 & C-PE3-> C-PE-2
  - Optional: IPsec attributes for tunnel from specific Subnet A on C-PEx

UPDATE 3:

- MP-NLRI Path Attribute
  - 12.1.1.x/24
- Tunnel-Encap:
  - IPsec SA attributes for C-PE1->C-PE2 & C-PE3-> C-PE-2

Applications Based Segmentation in SDWAN

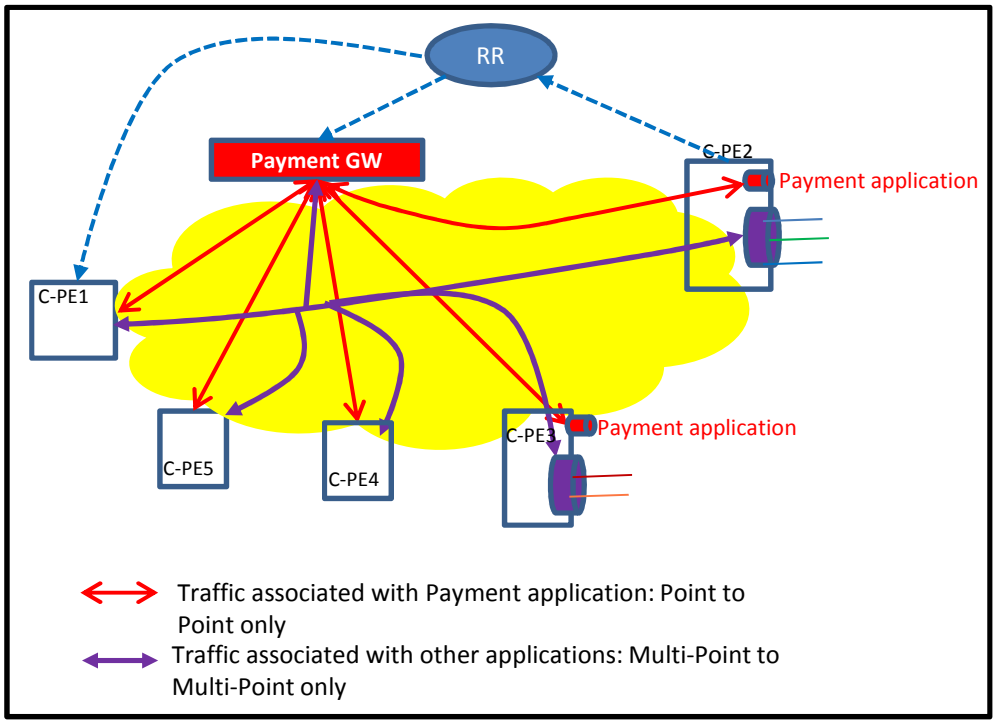

Characteristics:

- RED route (payment applications) can only be propagated to “Payment GW”
- Purple routes needs to be propagated to all other nodes

BGP Walk Through for Applications Based Segmentation in SDWAN

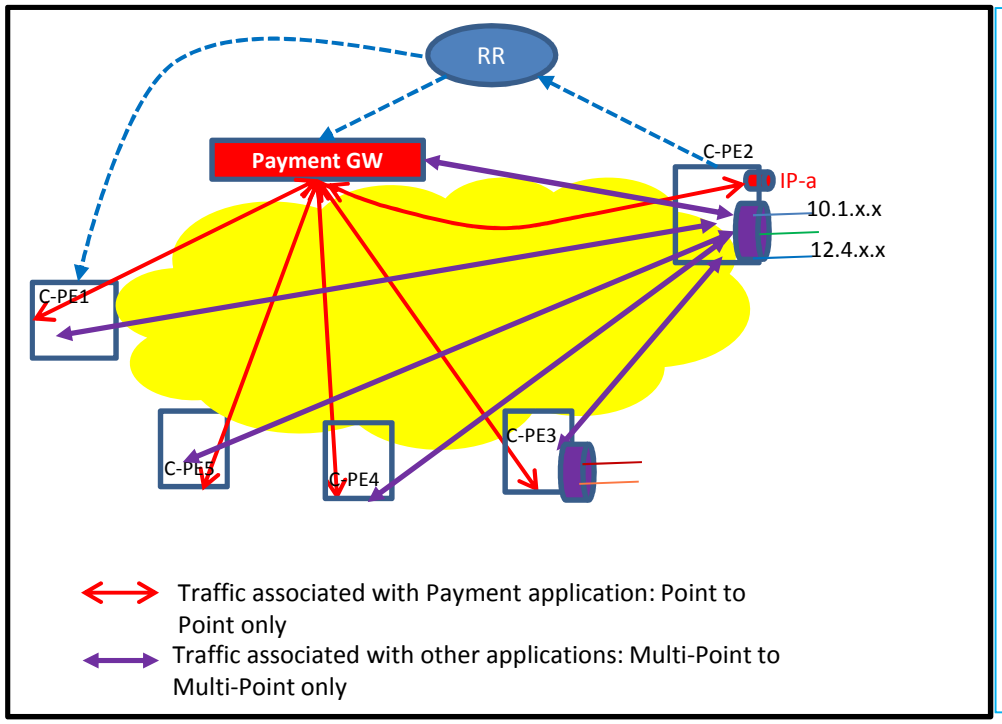

Assume Payment Application has different IP address than other segments: IP-a for Payment application

BGP UPDATE #1 from C-PE2 to RR for the RED P2P topology:

- MP-NLRI Path Attribute:
  - IP-a
- Tunnel Encap Path Attribute
  - IPsec Attributes for PaymentGW ->C-PE2

BGP UPDATE #2 from C-PE2 to RR for the routes to be reached by Purple:

- MP-NLRI Path Attribute:
  - 10.1.x.x
  - 12.4.x.x

- TunnelEncap Path Attribute:
  - IPsec SA-12 (C-PE1->C-PE2)
  - IPsec SA-32 (C-PE3->C-PE2)
  - IPsec SA-42 (C-PE4->C-PE2)
  - IPsec SA-52 (C-PE5->C-PE2)
  - IPsec SA-G2 (Payment GW -> C-PE2)

Purple Topology Supported by Multiple IPsec tunnels stitched together

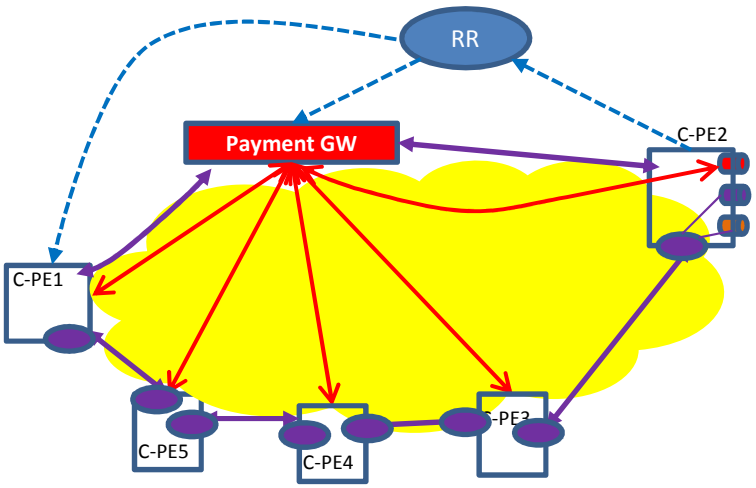

IPsec tunnel as transport that aggregate traffic to different destinations. Each edge route the traffic just like VRF.

BGP UPDATE from C-PE2 to RR for the IPsec C-PE3 -> C-PE2 segment:

- Port Addr encoded in MP-NLRI Path Attribute
- Port based IPsec tunnels encoded in the Tunnel Encap Path Attribute
  - IPsec SA-P-G32 (C-PE3>C-PE2)

BGP UPDATE from C-PE3 to RR for the IPsec C-PE2 -> C-PE3 segment:

- Port Addr encoded in MP-NLRI Path Attribute
- Port based IPsec tunnel encoded in the Tunnel Encap Path Attribute
  - IPsec SA-P-G23 (C-PE2>C-PE3)

BGP UPDATE from C-PE3 to RR for the IPsec C-PE4 -> C-PE3 segment:

- Port Addr encoded in MP-NLRI Path Attribute
- Port based IPsec tunnel encoded in the Tunnel Encap Path Attribute
  - IPsec SA-P-G43 (C-PE4>C-PE3)

BGP UPDATE from C-PE4 to RR for the IPsec C-PE3 -> C-PE4 segment
BGP UPDATE from C-PE4 to RR for the IPsec C-PE5 -> C-PE4 segment

Analysis of Multiple IPsec tunnels stitched together

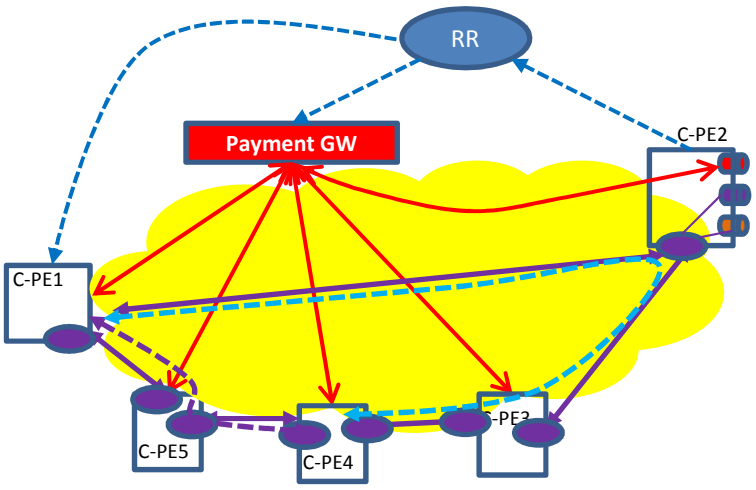

IPsec tunnel as transport that aggregate traffic to different destinations. Each edge route the traffic just like VRF.

Pros

- Each C-PEs has much smaller number of IPsec Tunnels to manage
- Less BGP UPDATE messages for keys refreshment and management
- Deterministic route to each destination.

Cons:

- Some C-PEs need to forward packets to another IPsec tunnel,
- Those C-PEs need to process those BGP UPDATE to build the forwarding table.
- There could be multiple ways for packets from A->B.
- TE needs to be enforced by the Controller

Next Step

- **WG Adoption.**
- **Why**
  - Give a clear picture on how BGP is used to scale SDWAN to the industry