

BGP Usage for SDWAN Overlay Networks

draft-dunbar-bess-bgp-sdwan-usage-08

Linda Dunbar & Jim Guichard (Futurewei)

Ali Sajassi (Cisco)

John Drake (Juniper)

Basil Najem (Bell Canada)

Ayan Barnerjee & Dave Carrel (Cisco)

July 2020

Key Items Discussed in the Draft

3. Use Case Scenario Description and Requirements.....	6
3.1. Requirements.....	6
3.1.1. Supporting Multiple SDWAN Segmentations.....	6
3.1.2. Client Service Requirement.....	6
3.1.3. Application Flow Based Segmentation.....	7
3.1.4. Zero Touch Provisioning.....	8
3.1.5. Constrained Propagation of SDWAN Edge Properties.....	9
3.2. Scenarios #1: Homogeneous WAN.....	10
3.3. Scenario #2: CPE based SDWAN over Hybrid WAN Underlay....	11
3.4. Scenario #3: Private VPN PE based SDWAN.....	14
4. BGP Walk Through.....	15
4.1. BGP Walk Through for Homogeneous SDWAN.....	15
4.2. BGP Walk Through for Application Flow Based Segmentation.	18
4.3. Client Service Provisioning Model.....	19
4.4. WAN Ports Provisioning Model.....	20
4.5. Why BGP as Control Plane for SDWAN?.....	20
5. SDWAN Traffic Forwarding Walk Through.....	21
5.1. SDWAN Network Startup Procedures.....	21
5.2. Packet Walk-Through for Scenario #1.....	22
5.3. Packet Walk-Through for Scenario #2.....	22
5.4. Packet Walk-Through for Scenario #3.....	24

key characteristics of “SDWAN” networks

- Augment of transport:
 - ❑ utilizing overlay paths over different underlay networks.
 - ❑ Among the multiple parallel overlay paths between any two SDWAN edges, some are private networks over which traffic can traverse with or without encryption, others require encryption, e.g. over untrusted public networks.
- Enable direct Internet access from remote sites, instead hauling all traffic to Corporate HQ for centralized policy control.
- Some traffic are routed based on application IDs instead of based on destination IP addresses.
- The Application Routing can also be based on specific performance criteria (e.g. packets delay, packet loos, jitter) to provide better application performance by choosing the right underlay that meets or exceeds the specified criteria.

To Support Multiple SDWAN Instances (segmentations)

- **Control Plane Messages:**

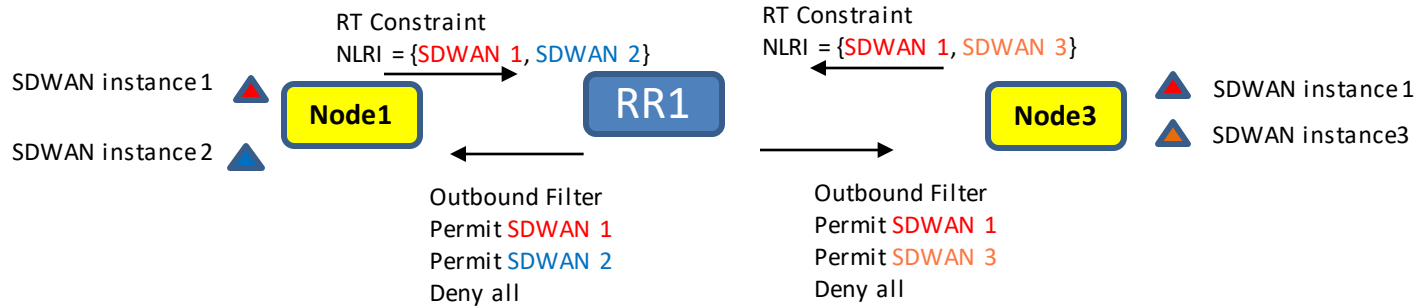
- Create a SDWAN Target ID in the BGP Extended Community to represent different SDWAN Segmentations
 - Same as Route Target, just use a different name to differentiate from VPN If a CPE supports traditional VPN with multiple VRFs, and supports multiple SDWAN Segmentations (instances).
- Overlay Topology:
 - Overlay connectivity is established based on SDWAN Route Target

- **Data Plane:**

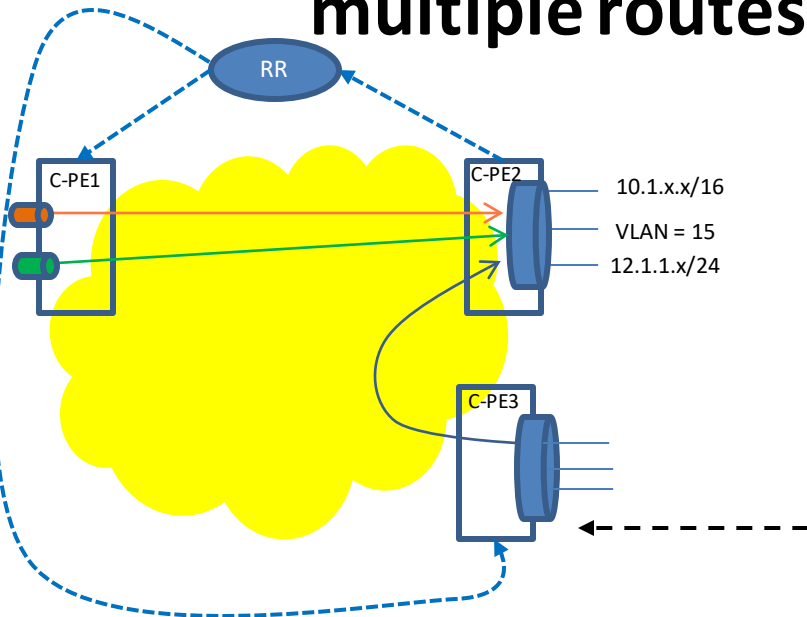
- Public network: SDWAN ID carried by the VNID (virtual network ID) or GRE keys carried by the inner encapsulation within IPsec ESP Tunnel.

Constrained Propagation of Clients routes/info

- Using RFC 4684 to constrain the distribution of BGP UPDATE to only a subset of SDWAN edges
- Using manually provisioned policies on RR to constrain the propagation of BGP UPDATE



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
-


SDWAN Instance 1
Import RT 1:1
Export RT 1:1

SDWAN Instance 2
Import RT 1:2
Export RT 1:2



BGP Walk Through for Homogeneous SD-WAN Client Routes with different Topologies & Policies

RR sends different UPDATE messages to different edges to reflect different topologies




One BGP UPDATE Message to C-PE1 to indicate the BLUE tunnel:

- 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 C-PE1 to C-PE2

One BGP UPDATE Message to C-PE3 to indicate the RED tunnel:

- multiple routes encoded in the MP-NLRI Path Attribute
 - VLAN #25
 - 22.1.1.x/24
- IPsec attributes are encoded in the Tunnel-Encap Path Attribute
 - IPsec attributes for C-PE3 to C-PE2

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 any nodes to establish IPsec tunnel C-PE-2 for the routes encoded in the MP-NLRI Path Attribute

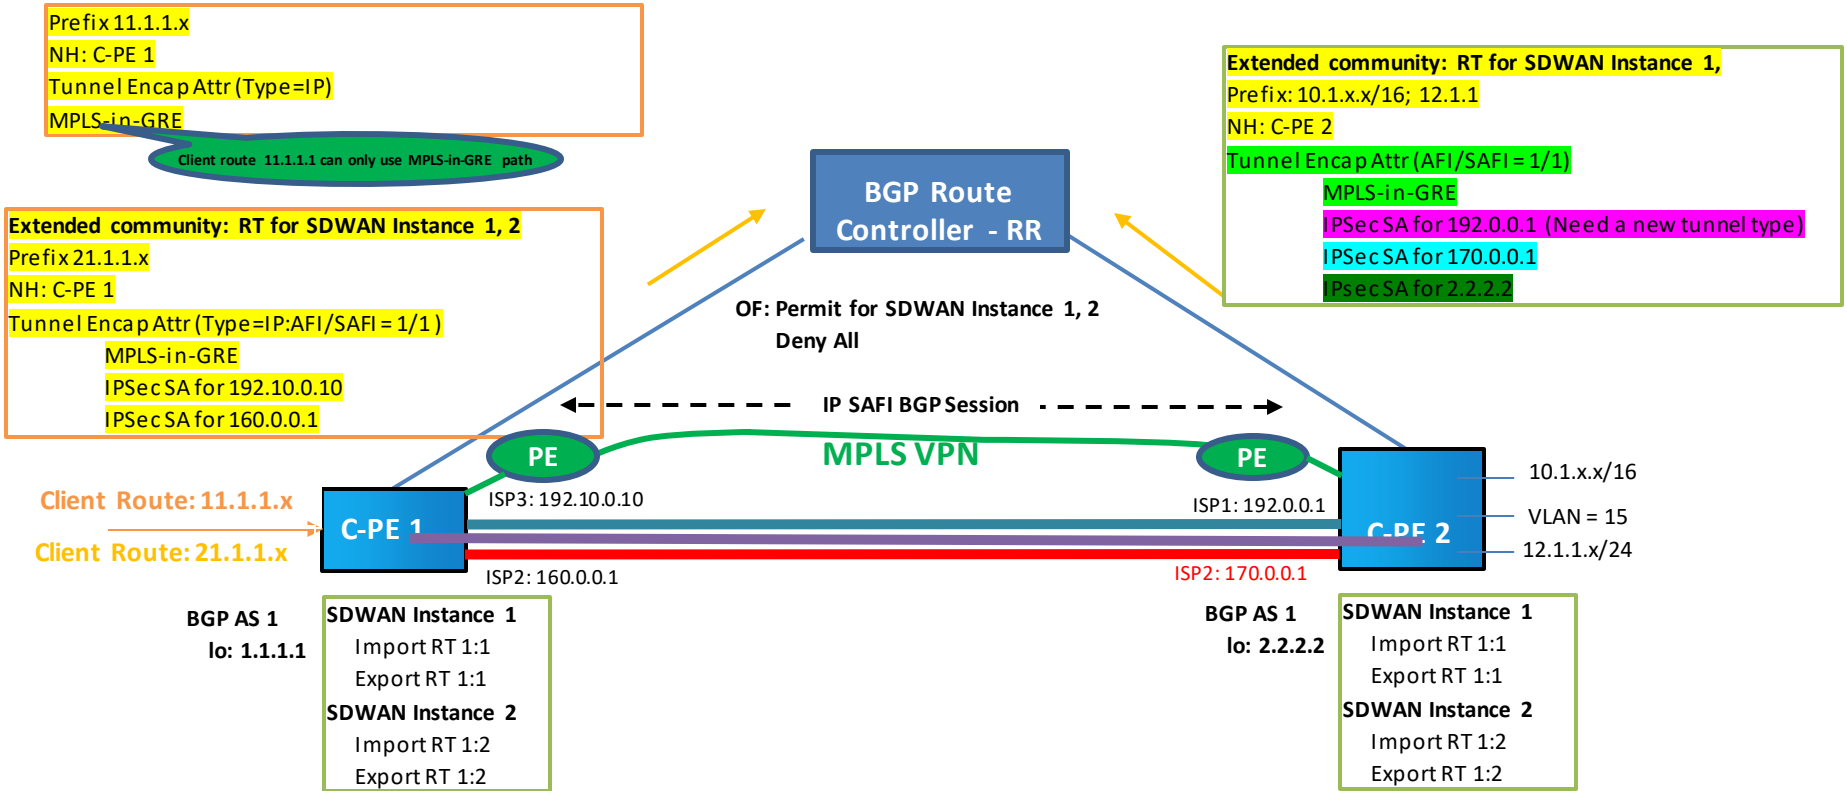
UPDATE 2:

- MP-NLRI Path Attribute
 - VLAN #15
- Tunnel-Encap
 - IPsec SA attributes for any nodes to establish IPsec tunnel C-PE-2 for the routes encoded in the MP-NLRI Path Attribute

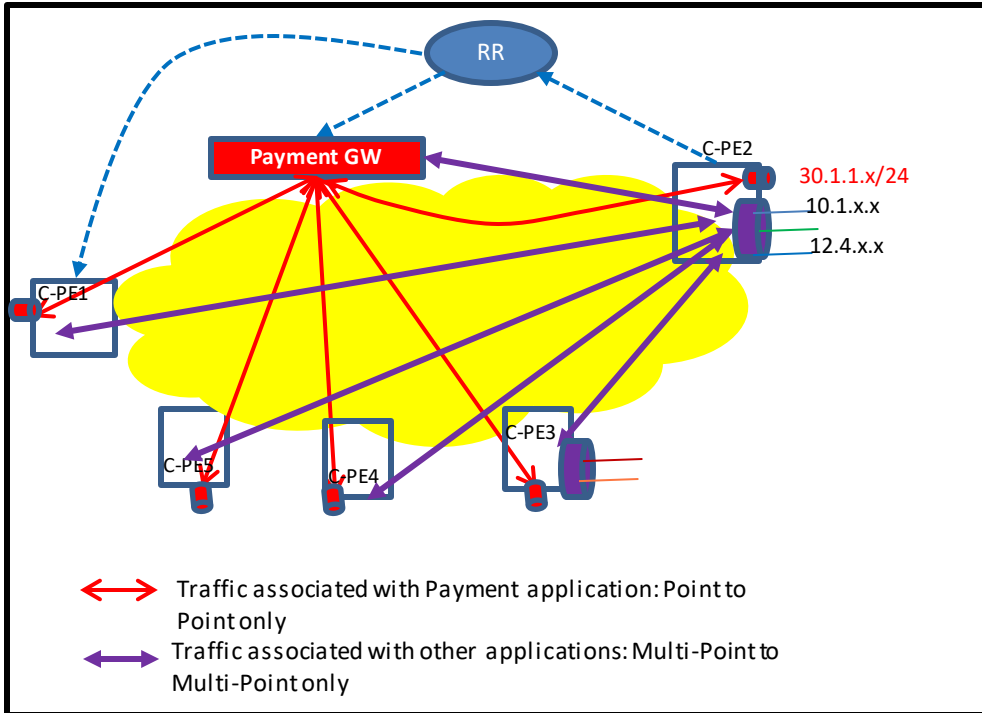
UPDATE 3:

- MP-NLRI Path Attribute
 - 12.1.1.x/24
- Tunnel-Encap
 - IPsec SA attributes for any nodes to establish IPsec tunnel C-PE-2 for the routes encoded in the MP-NLRI Path Attribute

BGP Walk Through for Scenario #2: C-PE with multiple Underlays



BGP Walk Through for Applications Based Segmentation in SDWAN



Assume Payment Application has different IP address than other segments, e.g. 30.1.1.x/24 for Payment application

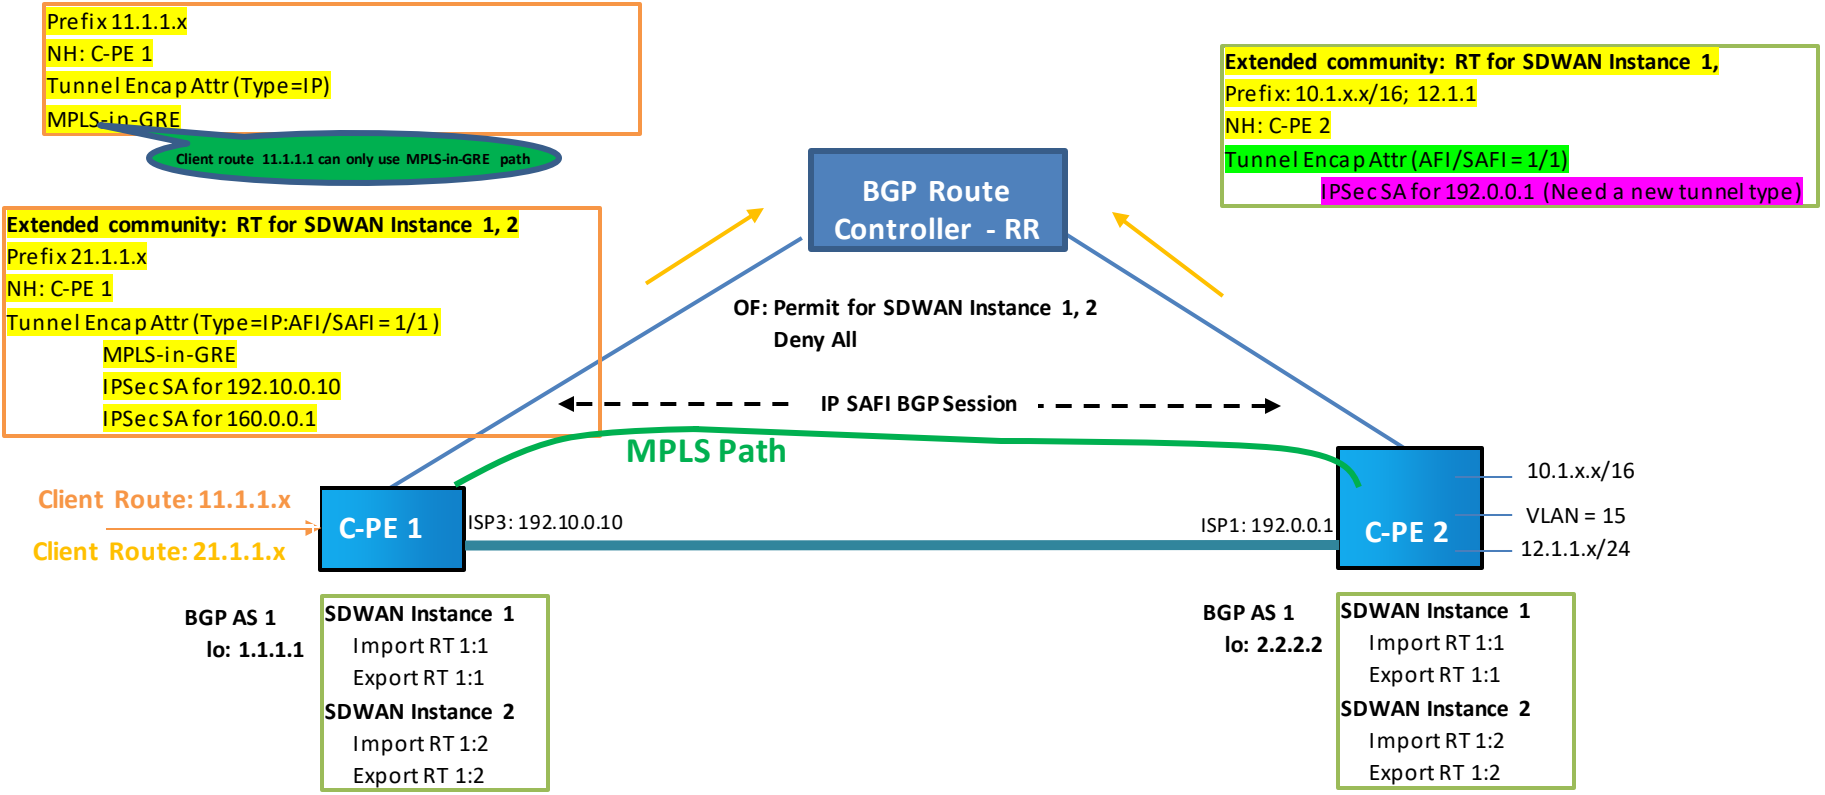
BGP UPDATE #1 from C-PE2 to RR for the RED P2P topology (only propagated to Payment GW node):

- MP-NLRI Path Attribute:
 - 30.1.1.x/24
- 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:
 - Any node to C-PE2

BGP Walk Through for Scenario #3: PE based SDWAN



Next Step

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