

BGP Usage for SDWAN Overlay Networks

draft-ietf-bess-bgp-sdwan-usage-01

Linda Dunbar & Jim Guichard (Futurewei)

Ali Sajassi (Cisco)

John Drake (Juniper)

Basil Najem (Bell Canada)

Ayan Barnerjee & Dave Carrel (Cisco)

Nov 2020

Key Items Discussed in the Draft

[3. Use Case Scenario Description and Requirements](#)

[3.1. Requirements](#)

[3.1.1. Supporting Multiple SDWAN Segmentations](#)

[3.1.2. Client Service Requirement](#)

[3.1.3. Application Flow Based Segmentation](#)

[3.1.4. Zero Touch Provisioning](#)

[3.1.5. Constrained Propagation of SDWAN Edge Properties](#)

[3.2. Scenario #1: Homogeneous WAN](#)

[3.3. Scenario #2: Hybrid WAN Underlay](#)

[3.4. Scenario #3: Private VPN PE based SDWAN](#)

[4. BGP Walk Through](#)

[4.1. BGP Walk Through for Homogeneous SDWAN](#)

[4.2. BGP Walk Through for Hybrid WAN Underlay](#)

[4.3. BGP Walk Through for Application Flow Based Segmentation](#)

[4.4. Client Service Provisioning Model](#)

[4.5. Underlay Network Properties Advertisement](#)

[4.6. Why BGP as Control Plane for SDWAN?](#)

[5. SDWAN Traffic Forwarding Walk Through](#)

[5.1. SDWAN Network Startup Procedures](#)

[5.2. Packet Walk-Through for Scenario #1](#)

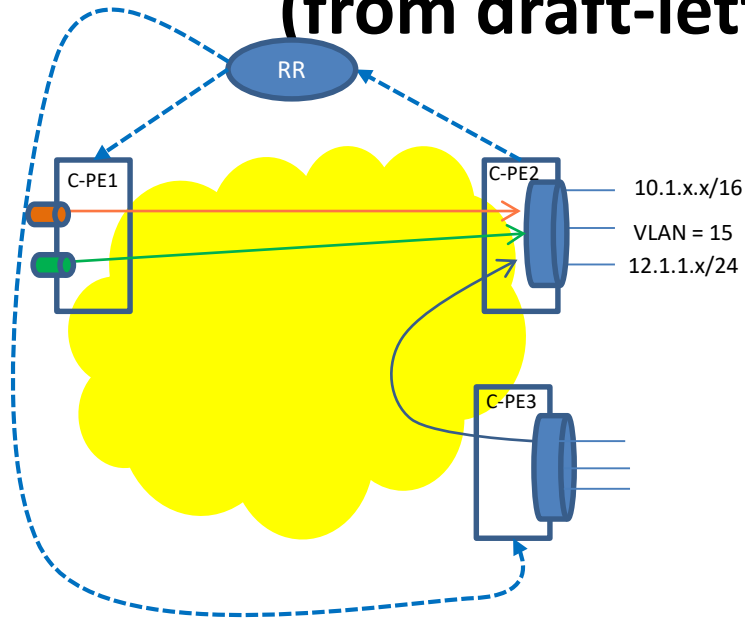
[5.3. Packet Walk-Through for Scenario #2](#)

[5.4. Packet Walk-Through for Scenario #3](#)

Major Changes since

- Reflecting draft-ietf-idr-tunnel-encaps-20
- To address suggestions from the BESS mailing list to use Recursive Lookup

BGP Walk Through for Homogeneous SD-WAN (from draft-ietf-bess-bgp-sdwan-usage-00)




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 Sub-TLVs [SECURE-EVPN]
-

Using Recursive Look-up with two UDPATEs

Homogeneous SD-WAN



UPDATE U1 from C-PE2 :

- multiple routes encoded in the MP-NLRI Path Attribute
 - 10.1.x.x/16
 - 12.1.1.x/24
- Next-hop: 2.2.2.2
- Encapsulation Extended Community: Tunnel-Type = IPsec
- Color Extended Community: Blue


UPDATE U2 from C-PE2 :

- MP-NLRI Path Attribute
 - 2.2.2.2
- Color Extended Community: Blue
- IPsec attributes are encoded in the Tunnel-Encap Path Attribute
 - IPsec SA attributes Sub-TLVs [SECURE-EVPN]

Homogeneous SD-WAN (from -00 version)

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

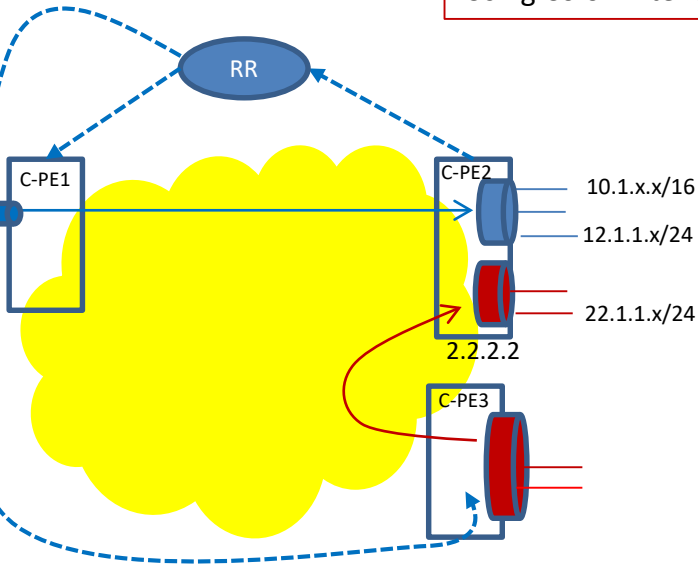
One BGP UPDATE Message to C-PE3 to indicate the RED tunnel:

- 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 Sub-TLVs for C-PE3 to C-PE2 [SECURE-EVPN]

Recursive Look-up: Homogeneous SD-WAN

Client Routes with different Topologies & Policies

Using Color Extended Community to reflect different topologies



UPDATE U1 (to indicate the BLUE tunnel going to C-PE1):

- MP-NLRI Path Attribute
 - 10.1.x.x/16; 12.1.1.x/24
- Next-hop: 2.2.2.2
- Encap Extended Community: IPsec
- Color Extended Community: Blue

UPDATE U2 (detailed IPsec Tunnel Attributes)

- MP-NLRI Path Attribute
 - 2.2.2.2
- Color: Blue
- IPsec attributes are encoded in the Tunnel-Encap Path Attribute
 - IPsec attributes for C-PE1 to C-PE2 [SECURE-EVPN]

UPDATE U1 (to indicate the RED tunnel to C-PE3)

- MP-NLRI Path Attribute
 - 22.1.1.x/24
- Encap Extended Community: IPsec or SDWAN
- Color Extended Community: RED

UPDATE U2 (detailed IPsec Tunnel Attributes)

- MP-NLRI Path Attribute
 - 2.2.2.2
- Color: RED
- IPsec attributes are encoded in the Tunnel-Encap Path Attribute
 - IPsec attributes for C-PE3 to C-PE2 [SECURE-EVPN]

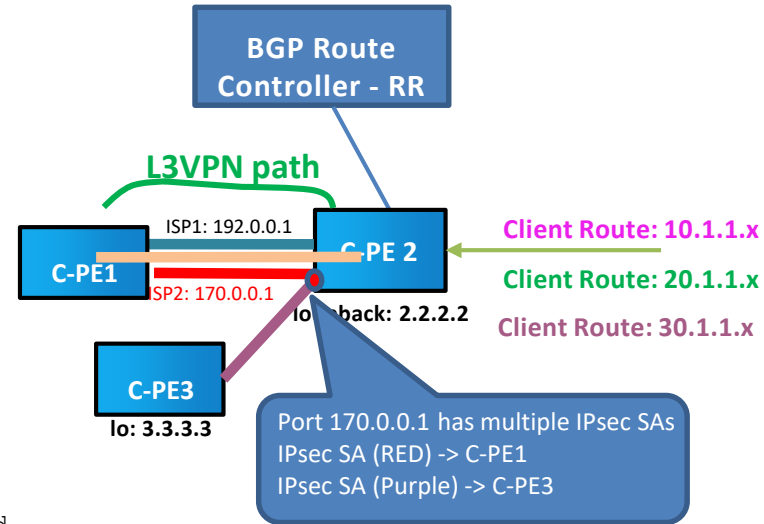
Hybrid SDWAN BGP Work-Through

UPDATE U1:

- MP-NLRI Path Attribute:
 - 10.1.x.x/16
 - 20.1.1.x/24
- Nexthop: 2.2.2.2 (C-PE2)
- Encapsulation Extended Community: Type = SDWAN-Hybrid
- Color Extended Community

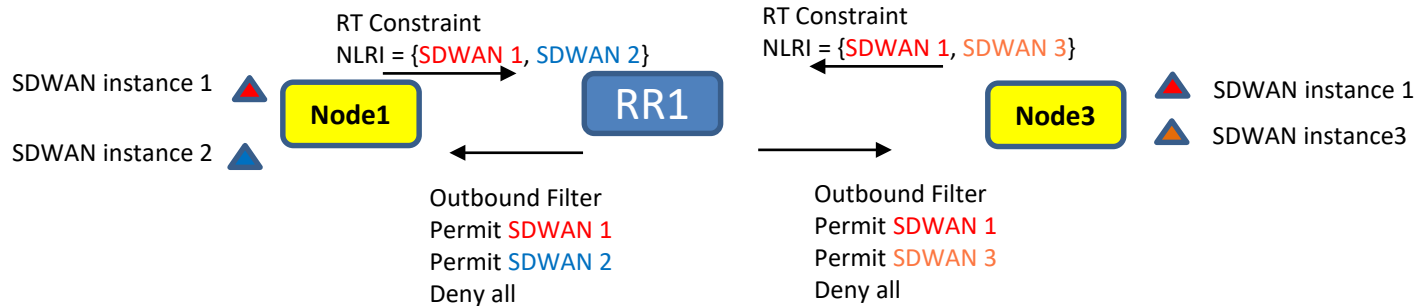
UPDATE U2:

- MP-NLRI Path Attribute:
 - 2.2.2.2 (C-PE2)
- Tunnel Encap Path Attributes: tunnel-type = SDWAN-Hybrid
 - Pre-configured IPsec SA IDs
 - Hybrid underlay tunnels
- Color Extended Community



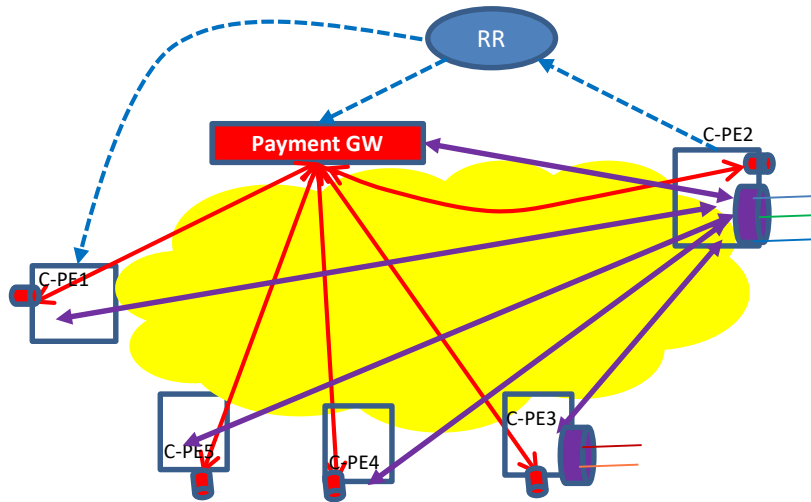
To prevents the propagation of SDWAN Update to edges that are not interested or not authorized

- **Static** (For small number of SDWAN edge nodes)
 - Policies can be configured on RR for RR to build outbound route filter, making RR as Route Controller (RC)
- **Dynamic:** using BGP RT constrained distribution (RFC4684)
 - A SDWAN edge sends RT Constraint (RTC) NLRI to the RR.
 - The RR installs an outbound route filter



Encoding

To differentiate SDWAN Segmentations



- ↔ Traffic associated with Payment application: Point to Point only
- ↔ Traffic associated with other applications: Multi-Point to Multi-Point only

Encoding for SDWAN Segmentation for Control Plane

- **To differentiate BGP Update by different SDWAN Segmentations:**
 - 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).
 - When the SDWAN Target ID is used,
 - Use the similar approach as VPN Label carried by NLRI Path Attribute [RFC8277] to identify routes belonging to different SDWAN Segmentations.
 - The MPLS VPN SAFI 128 & Route Distinguisher can be used for routes belonging to different SDWAN instances.
 - Question:
 - Use RFC7153, especially for 4 octet AS, use the Low 6 bits value or the Low Type.

RFC4360: Extended Community for SDWAN Route Target

```
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| Type high   | Type low(*) |                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               |                               | Value
|                               |                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

question:

- assign a new 6-bit VALUE for “SDWAN Instance Target” after the I & T bits?
- Or Request a new value of the low- order octet of the Type field for this community (different from the VPN Route Target 0x02)?

SDWAN Segmentation Identifier in Data Plane

- For a SDWAN edge node which can be reached by both MPLS and IPsec path, the client packets reached by MPLS network will be encoded with the MPLS Labels based on the scheme specified by RFC8277.
- For GRE Encapsulation within IPsec tunnel, the GRE key field can be used to carry the SDWAN Segmentation ID. For NVO (VxLAN, GENEVE, etc.) encapsulation within the IPsec tunnel, Virtual Network Identifier (VNI) field is used to carry the SDWAN Segmentation ID.
- [Note: the SDWAN Segmentation ID is same as EVI in EVPN, or VNI if VxLAN is used].

Next Step

- **Need Feedback**
- **WG LC?**
- **Why**
 - Give a clear picture on how BGP is used to scale SDWAN to the industry