BGP Usage for SDWAN Overlay Networks

draft-dunbar-bess-bgp-sdwan-usage-02

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

July 2019

SD-WAN Use Cases & Scenarios

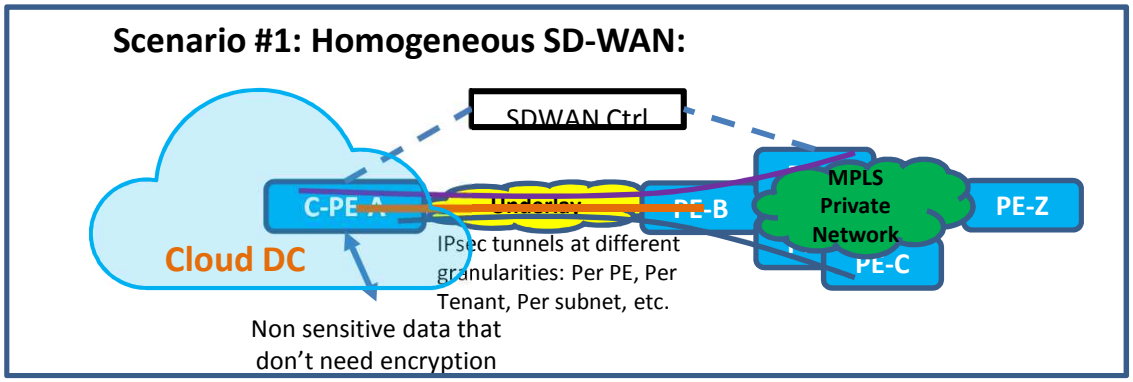

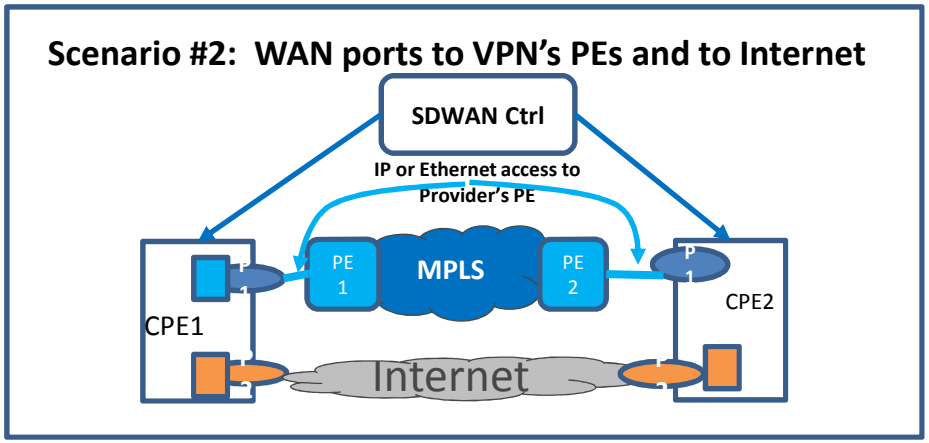

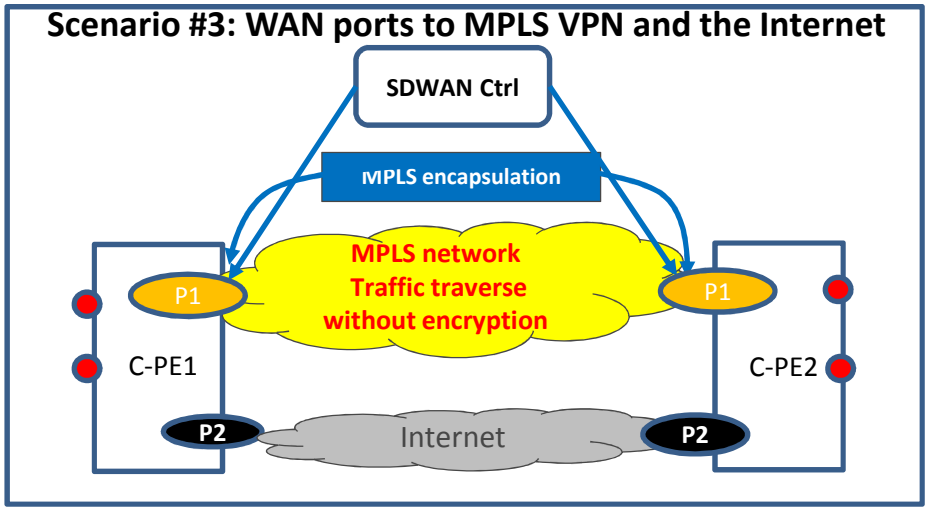

SDWAN over Hybrid Networks

Key Requirements

- Client Service Requirement
  - Client Interface: IP or Ethernet based
- SDWAN Node Provisioning
  1) ZTP,; 2) Auto-discovery of Network; 3) (Auto)-Provisioning for IPsec SAs (initial provisioning part); 4) Signaling of tenant's routes/info
- RR controlled Tenant scoped propagation

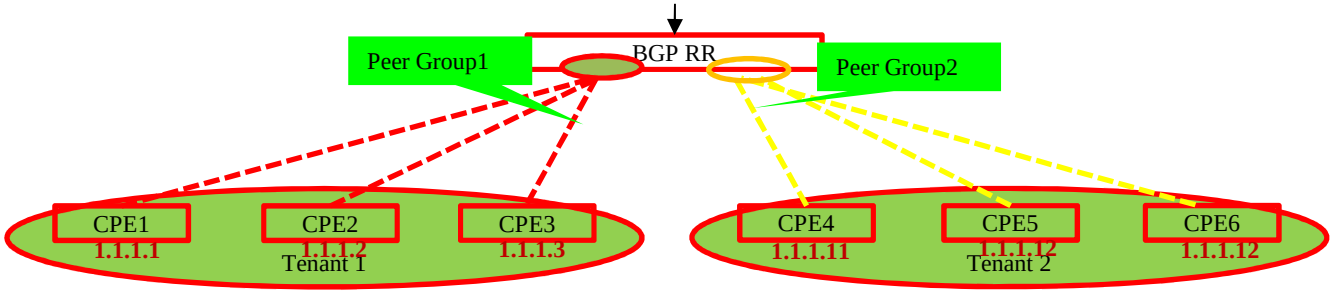

Key Characteristics of Scenario #1

- Key Use cases:
  - A small branch office connecting to its HQ offices via the Internet.
  - A store in a shopping mall may need to securely connect to its applications in one or more Cloud DCs via the Internet
- Key Characteristics:
  - All sensitive traffic to/from this small branch office has to be encrypted, which is usually achieved using IPsec SAs.
  - SDWAN Local Network Controller (RR)<-> C-PEs via untrusted public network,
    - requiring secure connection (TLS, DTLS, etc.)
- [SECURE-EVPN]: the granularity of the IPsec SAs for Homogeneous SDWAN can be per site, per subnet, per tenant, or per address.
- Need Controller managed re-key schemes for IPsec for all peers

Key Characteristics of Scenario #2: Multi-players

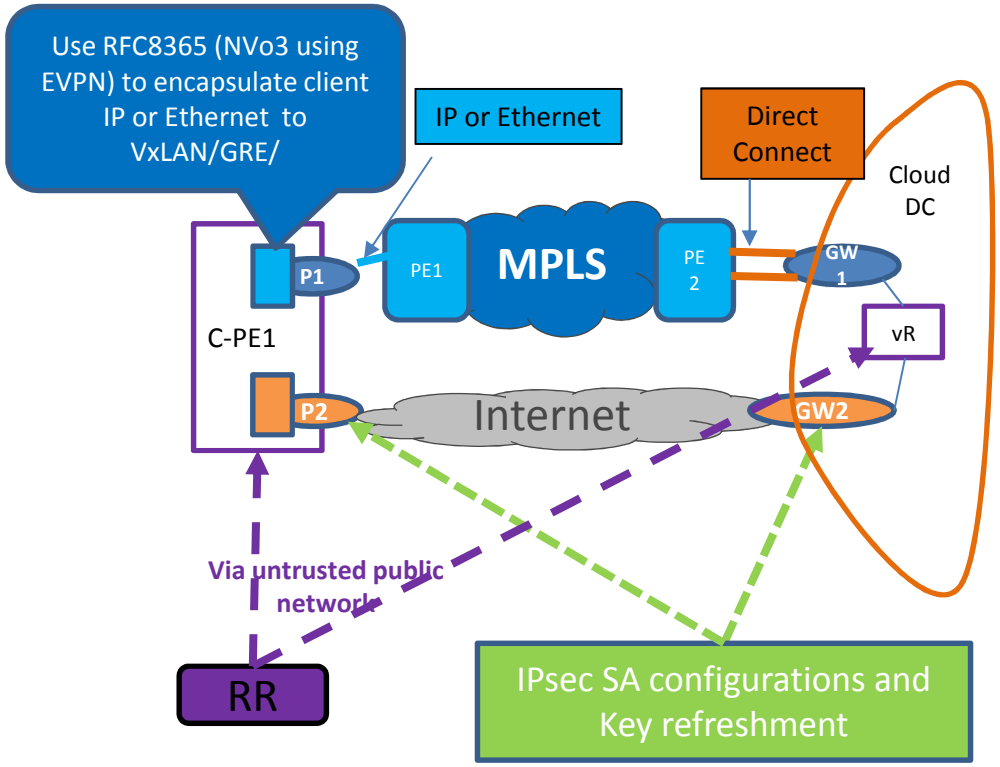

- Policies for flows:
  - Which flows over a private network without encryption (for better performance)
  - Which flows over any networks as long as the packets of the flows are encrypted when traversing untrusted networks
  - not needing encryption at all
  - Flows traverse specific geographic location (with more than one SDWAN nodes) → grouping
- For a flow traversing multiple segments, such as A<->B<->C<->D, can traverse different underlays in different segments
  - Services may not be congruent
  - Underlay port can be local decision
- the connections between BGP RR & CPEs are over public network, therefore requires TLS, DTLS, or IPsec.

Application components can be available from multiple Cloud DCs.

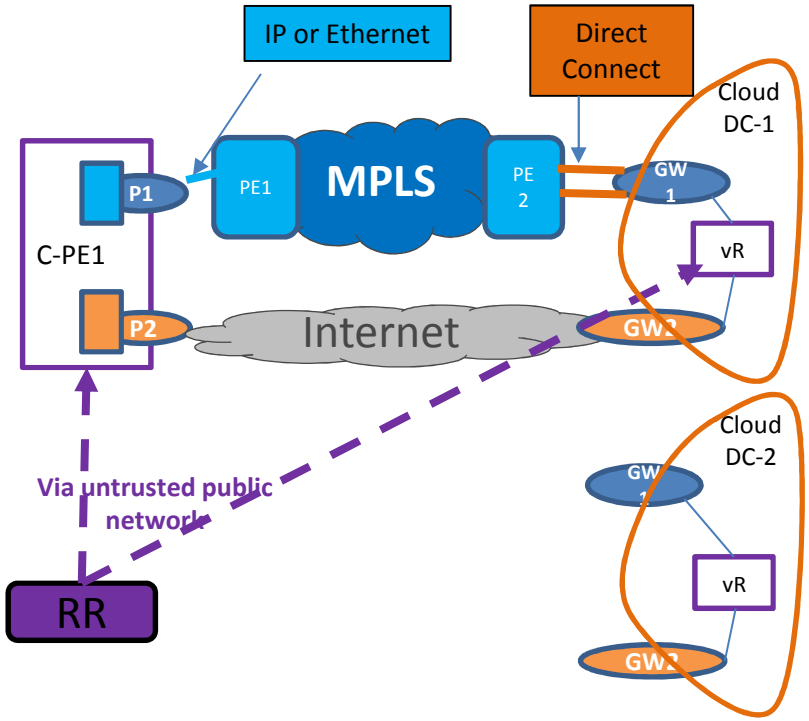

“Site” (or location) has to be present

Key Characteristics of Scenario #3

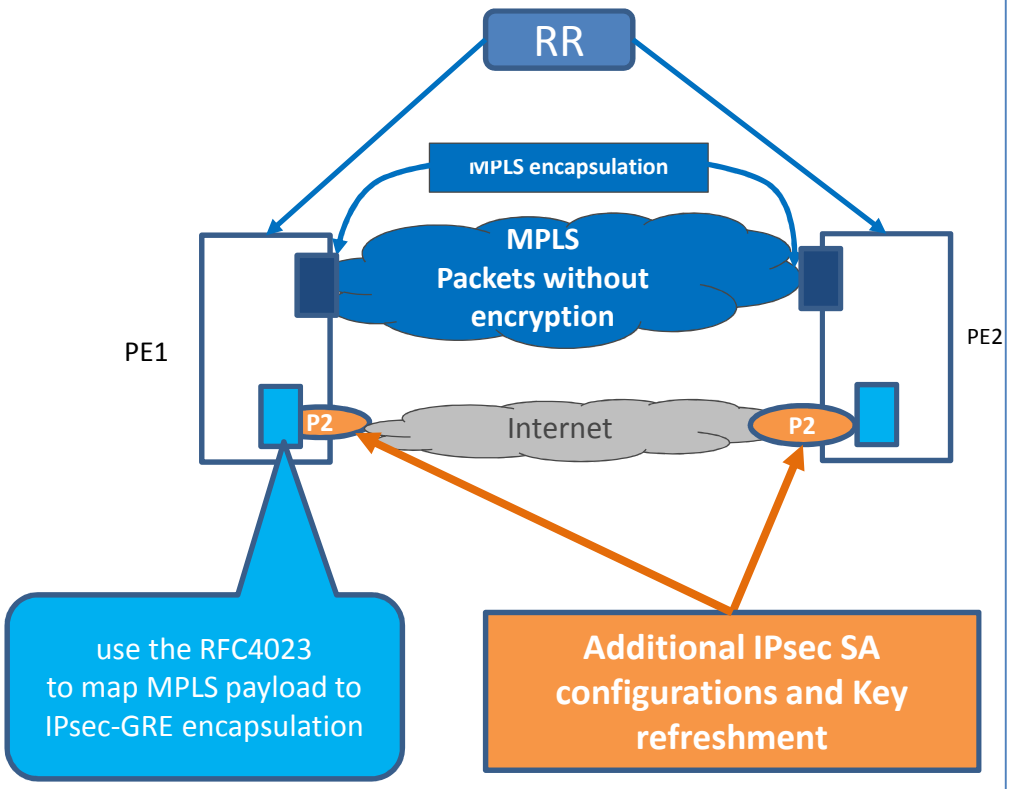

- PEs continue having MPLS encapsulation handoff to existing paths.
- The BGP RR is connected to PEs in the same way as VPN, i.e. via the trusted network.
- For the added Internet ports, PEs have IP handoff.
  - PEs can have the option to encapsulate the MPLS payload in IP, as specified by RFC4023.
- The ports facing public internet might get IP addresses assigned by ISPs, which may not be in the same address domain as PEs'.
- Ports facing public internet are not secure
  - could face spoofing, or DDOS attacks
  - Extra consideration must be given when injecting the new routes from public network into VRFs.
- the performance SLA is not guaranteed over public internet.
  - clients may have policies only allowing some flows to be offloaded to internet path.

SDWAN Site

Multiple SDWAN nodes sharing a common property: e.g. geographic location

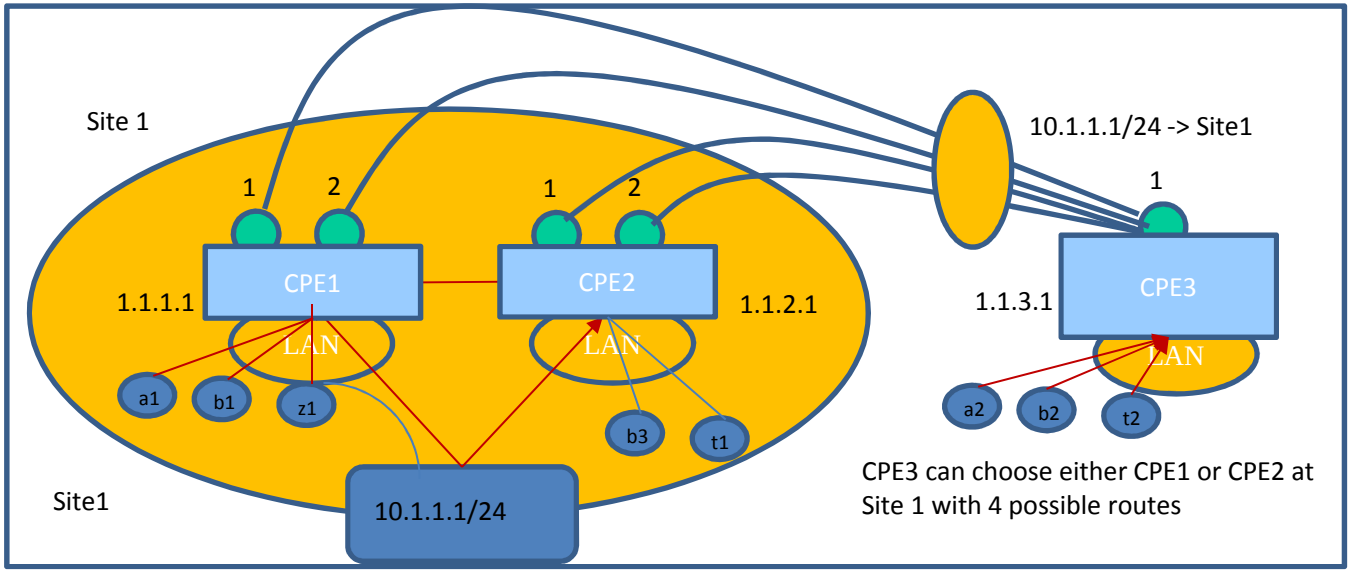

SDWAN Traffic Forwarding Walk Through

- **SDWAN Network Startup Procedures**
  - For Scenario #1: a SDWAN edge node in a shopping mall or Cloud DC can be added or removed on demand. The Zero Touch Provisioning is required for the node startup.
  - For Scenario #2: this can be Data Centers or enterprises upgrading their CPEs to add extra bandwidth via public internet in addition to VPN services that they already purchased. Before the node powers up or upgraded, there should be links connected to the PEs of a provider VPNs.
  - For Scenario #3, the Internet facing WAN ports are added to (or removed from) existing VPN PEs.

- **Client Service Provisioning Model**
  - The provisioning tasks described in Section 4 of RFC8388 are the same for the SDWAN client traffic
  - When client traffic are multi-homed to two (or more) C-PEs, the Non-Service-Specific parameters need to be provisioned per the Section 4.1.1 of RFC8388
- **WAN Ports Provisioning**

Three different Components of SD-WAN Control Plane

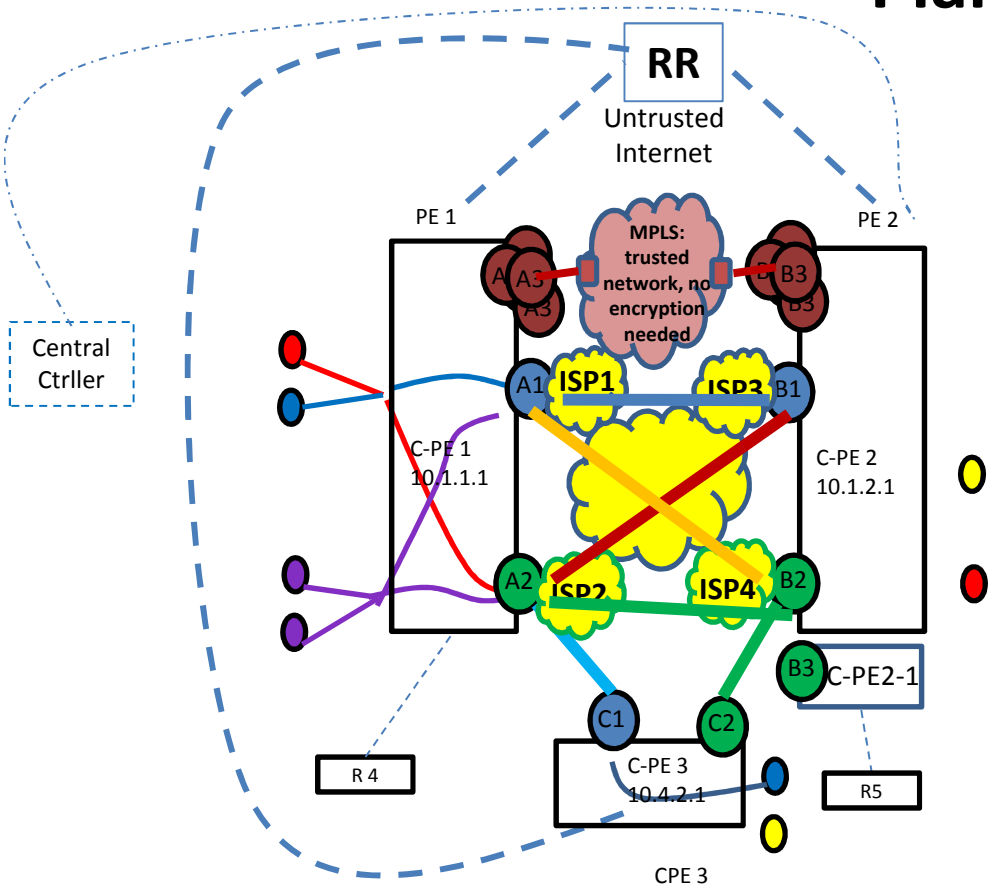

1. **End Node Registration:**
   SD-WAN node's private address and WAN Ports/Addresses registration to the SD-WAN Controller.
   It is for informing the SD-WAN controller and potential peers of the underlay networks to which the CPE is connected.
2. **Controller facilitated IPsec SA association establishment among WAN Ports**
3. **Attached routes distribution using BGP RR:**
   - **EVPN**
   - **IPVPN**
   - **Or something else**

Questions?