EQUINIX
INTERNET EXCHANGE (EIE)

EQUINIX

Aug 2022

Agenda

- Equinix Internet Exchange (EIE)
- EIE Peering
- EIE VXLAN Implementation
- EIE Client Port Configuration & Onboarding
- EIE RPKI Integration
- EIE RTBH
- EIE with EVPN Control-plane

Equinix Internet Exchange(EIE)

Equinix Internet Exchange (EIE)

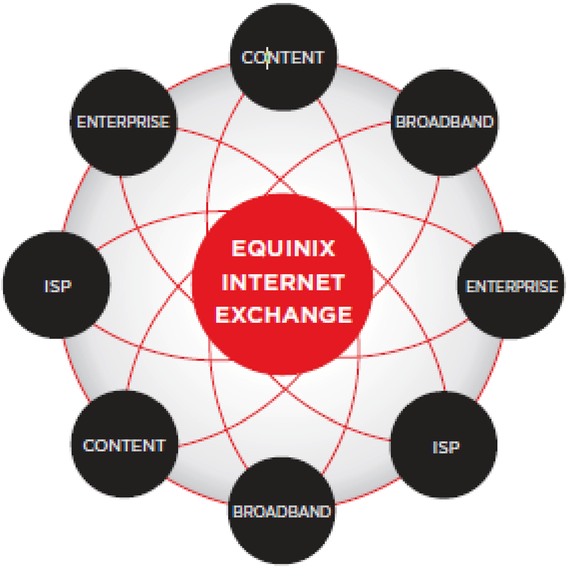

About Equinix Internet Exchange(EIE)

Equinix Internet Exchange enables networks, content providers and large enterprises to exchange internet traffic through the largest global peering solution across 43+ markets.
The Equinix Internet Exchange is a Layer 2 platform that enables interconnection (peering) between multiple networks in an operationally-efficient and cost-effective manner.

Benefits of Peering

Performance
Cost Reduction
Flexibility

EIE Features:

Private Vlan
MPLE
RTBH

Global Scale of Deployment

5 continents, 23 countries, 43 metros, 193 IBXs

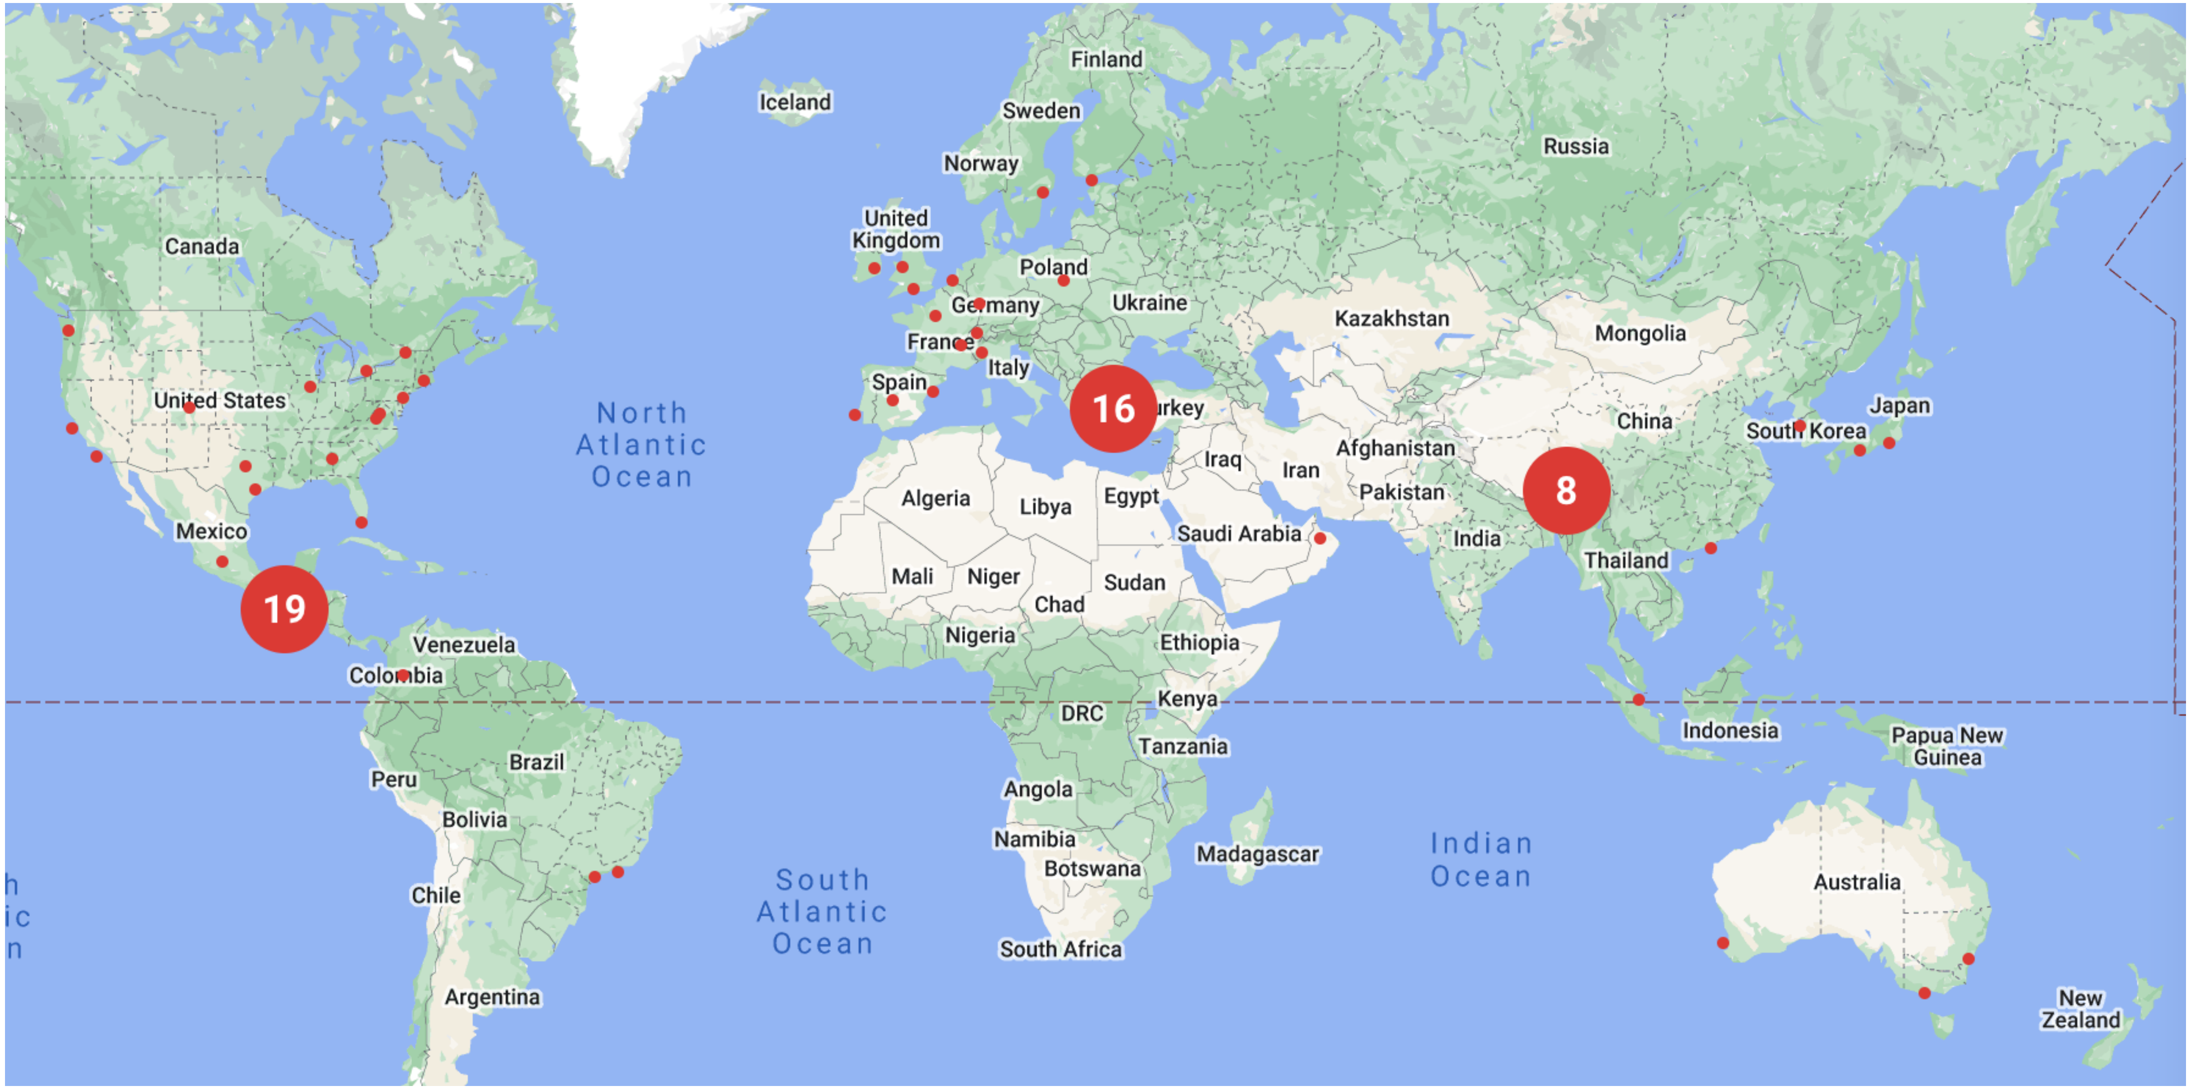

APAC Scale of Deployment

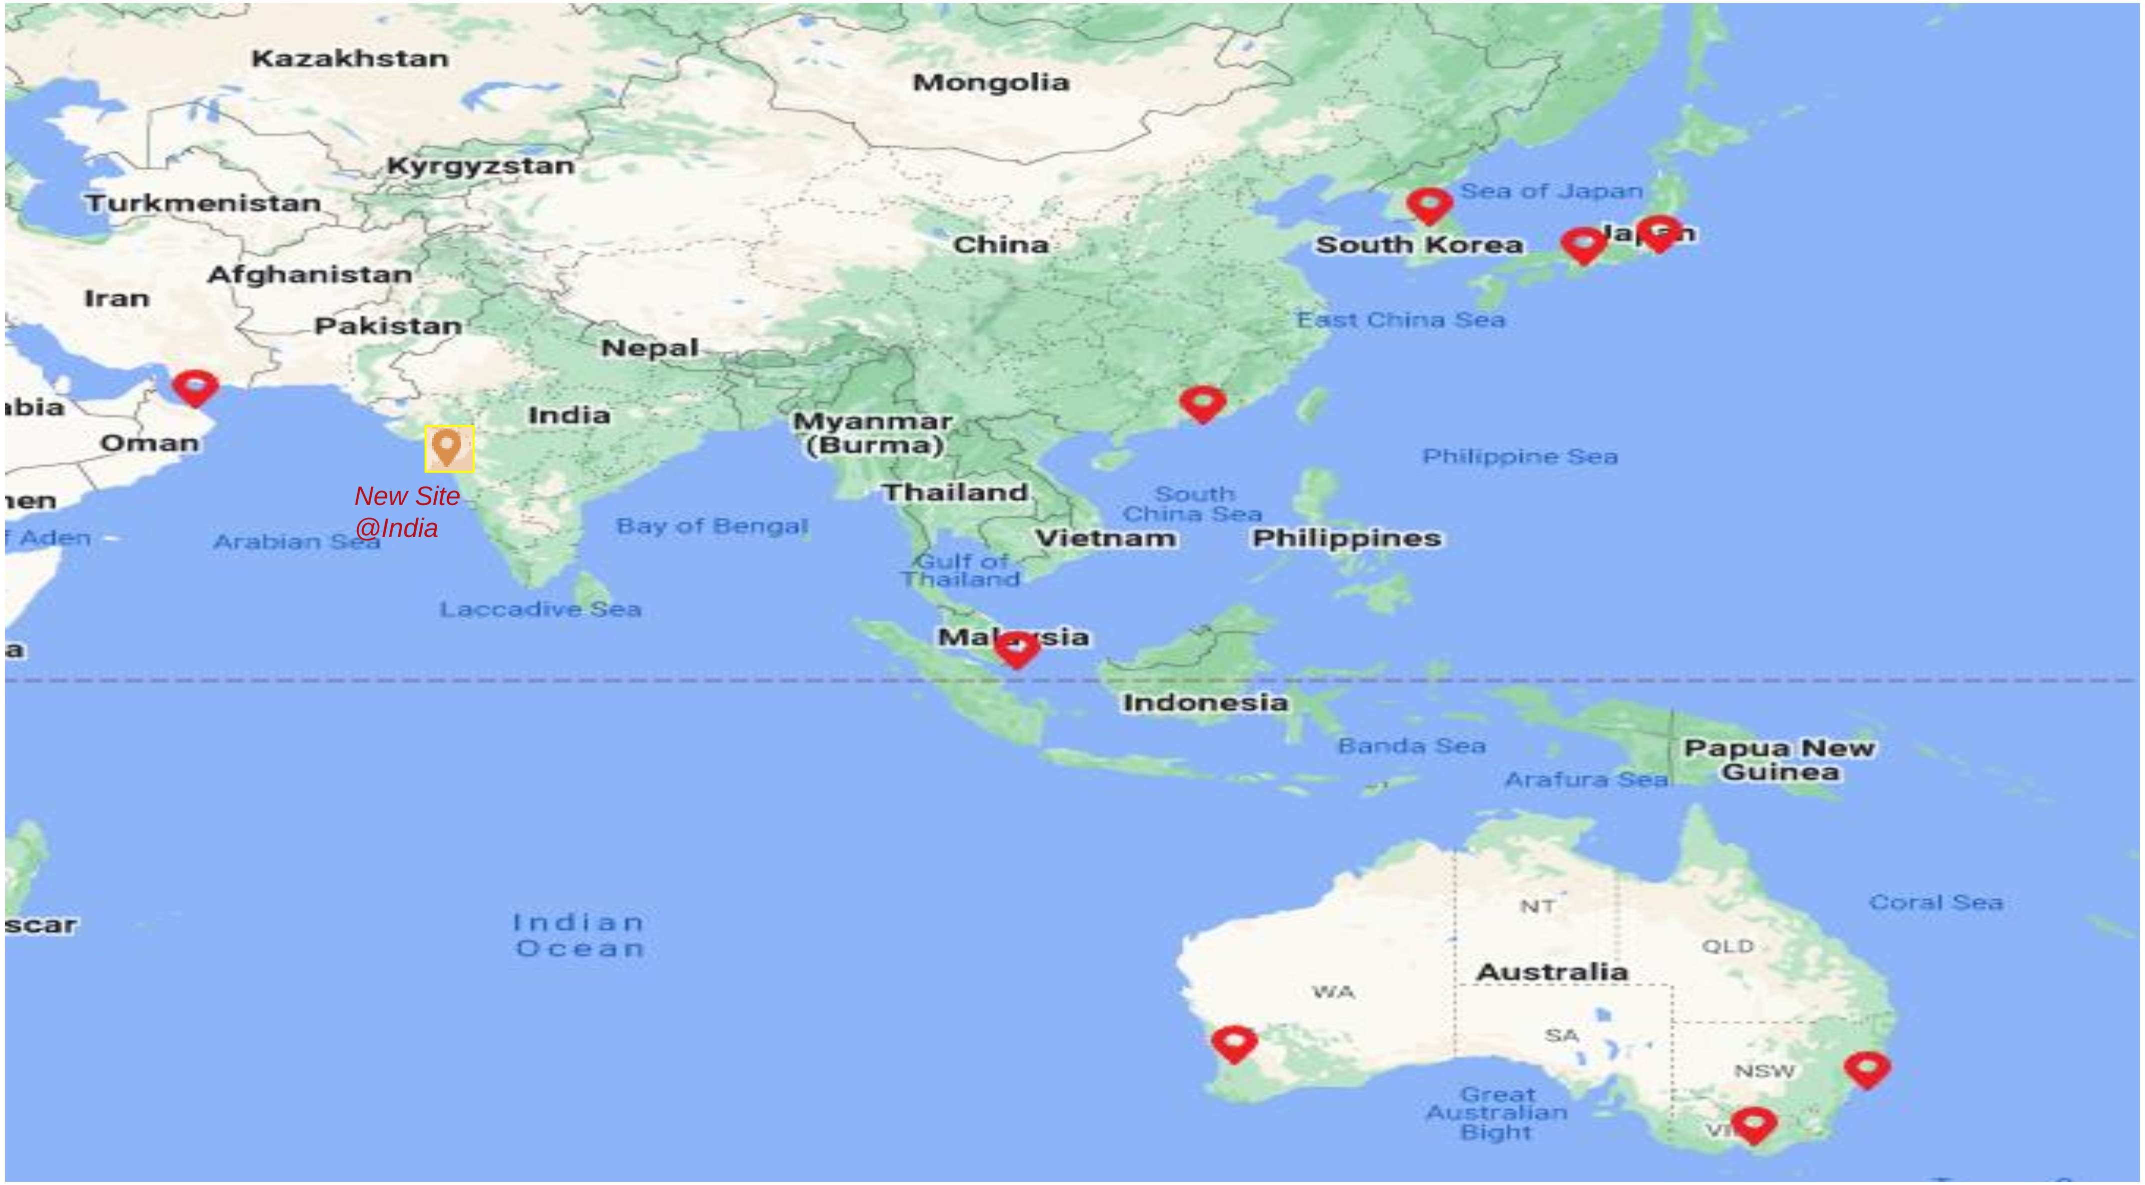

New Site @India

EIE VXLAN Implementation

VXLAN Implementation

Virtual Extensible LAN (VXLAN) is a network virtualization technology to address the limitation of VLAN range and unused (blocking) port in spanning tree topology.

VXLAN provides network overlay capability which allows layer-2 connectivity across layer 3 IP networks, by encapsulation technique to encapsulate layer 2 Ethernet frames within layer 4 UDP datagrams.

Head-end replication (HER) is used to flood Broadcast, unknown-unicast and multicast traffic.

- Overhead 8 byte for VXLAN header.
- Simple configuration, easy to implement, non-complicate in troubleshooting.
- VXLAN traffic follows underlay forwarding path, ECMP or LAG hashing works as it is.
- HER provides traditional mac-address learning (flood & learn).
- EVPN eliminate HER and provides mac-address learning at control-plane (instead of flood & learn at data plane).
- EVPN + VXLAN will be the best choice for scalable layer2 distributed network with additional features.
    - ✓ Faster convergence, mac-address withdrawal (minimize transient unknown unicast traffic due to port went down).
    - ✓ Scalability from small environments to cloud infrastructure (multi-pod, multi-sites).
    - ✓ Active-active multi-homing (provides redundant client access ports).

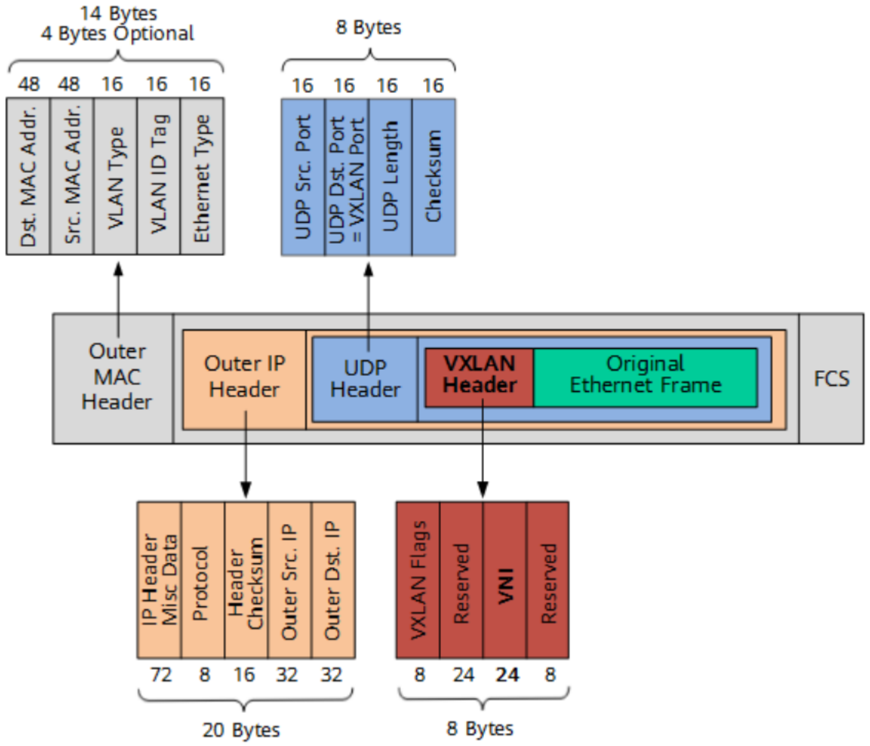

VXLAN in Internet Exchange Switch Fabric

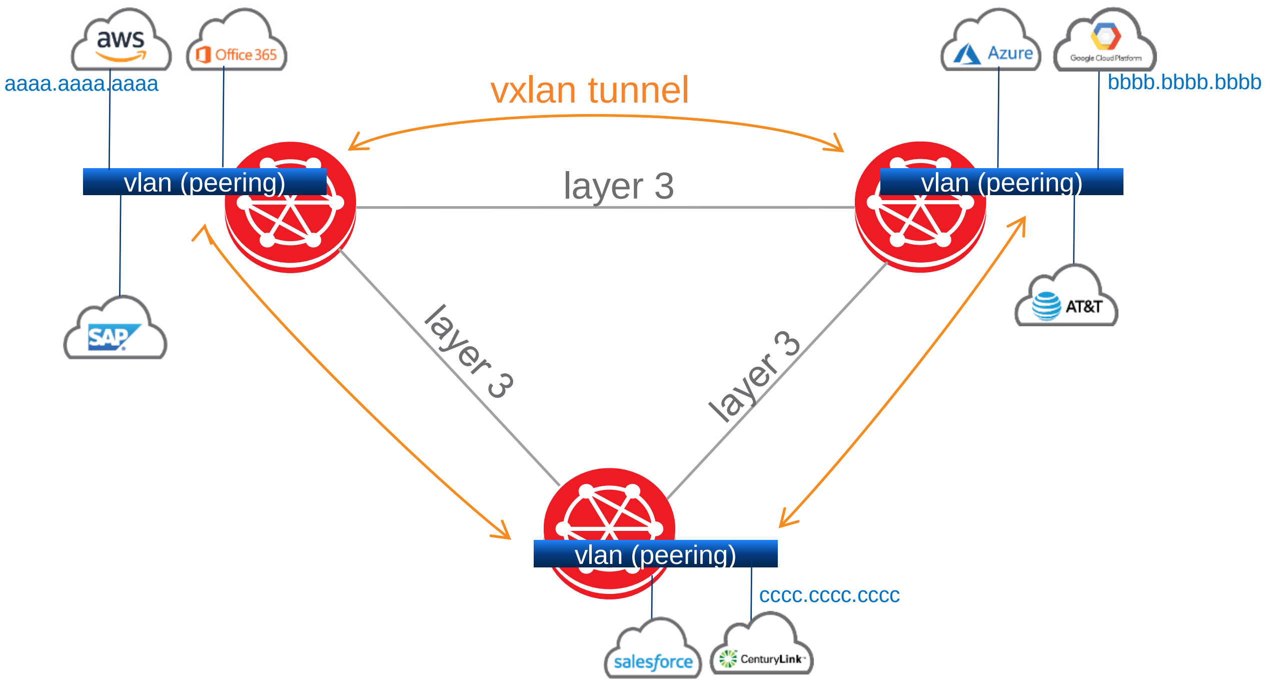

VXLAN Configuration Building Blocks:

- Loopback interface with IPv4 address
- Underlay routing (to reach loopback each other)
- VXLAN Configuration
  - ✓ Associate VLANs with VNI ID
  - ✓ VTEP destination (where BUM traffic to be flood)
  - ✓ Source interface and udp port (optional)

```
# LOOPBACK INTERFACE
interface Loopback0
   ip address 10.0.0.1/32
!
# VXLAN INTERFACE (CONFIGURATION)
interface Vxlan1
   vxlan source-interface Loopback0
   vxlan udp-port 4789
   vxlan vlan 100 vni 100
   vxlan vlan 200 vni 200
   vxlan vlan 100 flood vtep 10.0.0.2 10.0.0.3
   vxlan vlan 200 flood vtep 10.0.0.2 10.0.0.3
!
# UNDERLAY ROUTING
router ospf 1
   router-id 10.0.0.1
   network 10.0.0.0/16 area 0.0.0.0
!
interface Ethernet1/1
   description sw1.sg1, eth 1/1, layer-3 to sw2.sg1
   mtu 9214
   no switchport
   ip address 10.0.1.0/31
   ip ospf cost 1
   ip ospf network point-to-point
!
interface Ethernet2/1
   description sw1.sg1, eth 2/1, layer-3 to sw3.sg1
   mtu 9214
   no switchport
   ip address 10.0.1.2/31
   ip ospf cost 1
   ip ospf network point-to-point
!
```

IP Addressing

Switch Name	Loopback IP	Point-to-Point IP
sw1.sg1	10.0.0.1	10.0.1.x/31
sw2.sg1	10.0.0.2	10.0.1.x/31
sw3.sg1	10.0.0.3	10.0.1.x/31

sw1.sg1's mac-address (CAM) table

Mac-address	Outgoing interface
aaaa.aaaa.aaaa	Ethernet 1/1
bbbb.bbbb.bbbb	Vxlan1
cccc.cccc.cccc	Vxlan1

EIE Client Port Configuration & Onboarding

EIE Client Port Configuration

Key Components:

- Storm control (10Mbps of BUM traffic)
- Disable mac address learning, use static mac-address in CAM table
- Mac ACL (allow IPv4, IPv6 and ARP traffic with single mac-address)

Sample Config (lag, non tag port)

```
interface Port-Channel1
   description sw1.sg1, po 1, A, user1, circuit1
   load-interval 5
   switchport access vlan 599
   no switchport mac address learning
   mac access-group user1_circuit1 in
   storm-control broadcast level 0.01
   storm-control multicast level 0.01
   storm-control unknown-unicast level 0.01
!
interface Ethernet13/2/1
   description sw1.sg1, po 1, A, user1, circuit1, port1
   load-interval 5
   channel-group 1 mode active
   no lldp transmit
   no lldp receive
!
interface Ethernet14/23/1
   description sw1.sg1, po 1, A, user1, circuit1, port2
   load-interval 5
   channel-group 1 mode active
   no lldp transmit
   no lldp receive
!
mac access-list user1_circuit1
   10 permit ab:cd:ab:cd:ab:cd 00:00:00:00:00:00 any arp
   20 permit ab:cd:ab:cd:ab:cd 00:00:00:00:00:00 any ip
   30 permit ab:cd:ab:cd:ab:cd 00:00:00:00:00:00 any ipv6
   40 deny any any
!
```

Sample Config (tagged port)

```
interface Ethernet3/18
   description sw1.sg1, eth 3/18, user2, circuit1
   load-interval 5
   switchport mode trunk
   no switchport mac address learning
   switchport trunk group 100
   switchport trunk group 200
   mac access-group user2_circuit1 in
   no lldp transmit
   no lldp receive
   storm-control broadcast level 0.1
   storm-control multicast level 0.1
   storm-control unknown-unicast level 0.1
!
mac access-li`st user1_circuit1
   10 permit 01:23:01:23:01:23 00:00:00:00:00:00 any arp
   20 permit 01:23:01:23:01:23 00:00:00:00:00:00 any ip
   30 permit 01:23:01:23:01:23 00:00:00:00:00:00 any ipv6
   40 deny any any
!
```

EIE Client Port Protection

- Storm control mitigate incoming BUM traffic to 10Mbps.

Initial state

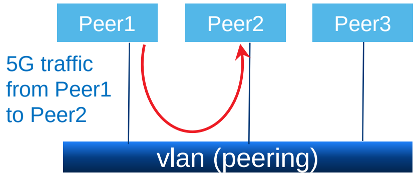

Without storm-control

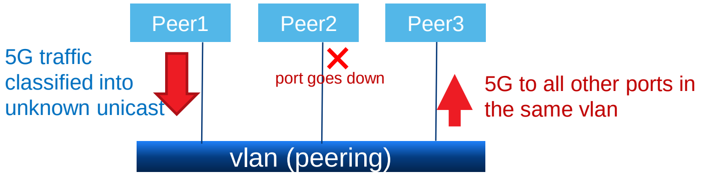
Peer1
Peer2
Peer3
5G traffic classified into unknown unicast
port goes down
5G to all other ports in the same vlan
vlan (peering)

With storm-control

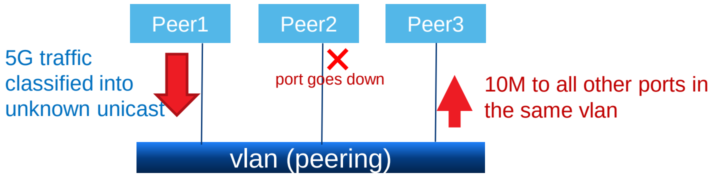
Peer1
Peer2
Peer3
5G traffic classified into unknown unicast
port goes down
10M to all other ports in the same vlan
vlan (peering)

- Static Mac-address mitigate BUM traffic caused by silent peer ports.

Initial state

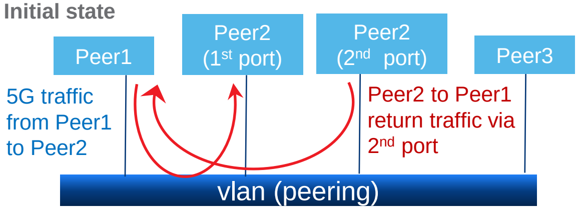
Peer1
Peer2 (1st port)
Peer2 (2nd port)
Peer3
5G traffic from Peer1 to Peer2
Peer2 to Peer1 return traffic via 2nd port
vlan (peering)

With dynamic mac learning, Peer2 1st port mac-address expires (no ingress traffic)

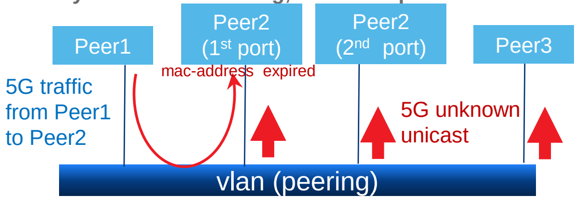

With static mac, no mac-address expires (permanent)

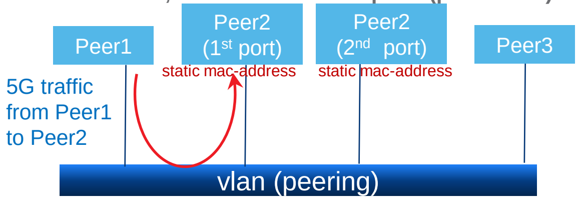

- Enforce single mac-address per port

Peer with 2 ports (causing switching loop

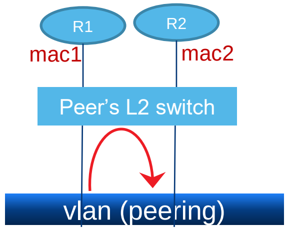

Peer with 2 x Ports (using mac-acl)

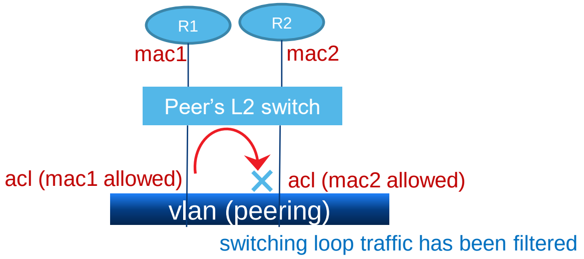

EIE Client Port Onboarding

Goal:

- The port configuration is correct at both sides (tag/untag, lag bundle, vlan membership, etc…).
- The correct mac-address is connected on the new port and being configured on the switch's mac-acl.
- Ensure only ARP/IPv4/IPv6 Ether types are sent by new port.

Steps:

1) Initially, the new port is put under quarantine VLAN and unshut.
2) The necessary switch port configuration (lag bundle, trunk/access mode, vlan membership, stormcontrol) has been added by provisioning script or turn-up engineer.
3) When the port comes up, verify physical connection (optical transceiver reading) and IP (v4, v6) connectivity (ping test from quarantine route server).
4) Capture the ingress traffic in new port and verify if there is any prohibited traffic, such as: STP, LLDP, CDP protocols.
5) Establish BGP session with quarantine route server and verify the advertised routes.
6) After verification steps are completed, change the new port from quarantine VLAN to production VLAN.
7) The new client can establish peering with production route servers (MLPA) as well as other clients (BLPA).

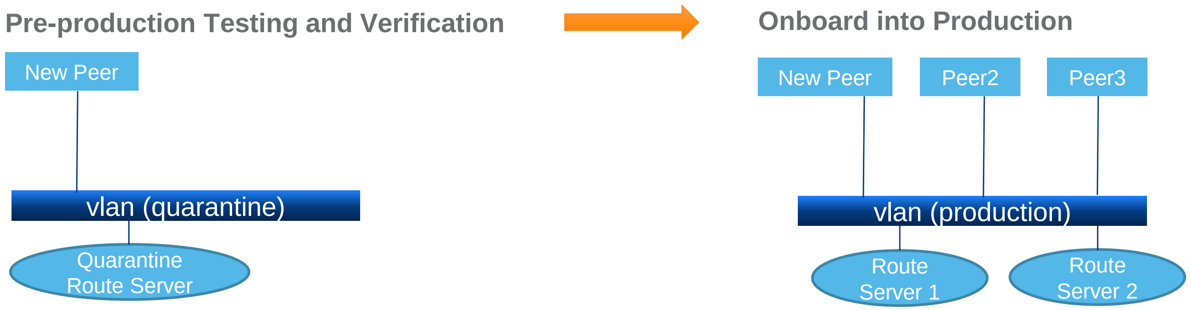

EIE RPKI Integration

What is RPKI ?

Resource Public Key Infrastructure

Public key infrastructure framework designed to secure the Internet’s routing infrastructure specifically the BGP & reduce the risk of BGP hijacking.

Security framework provides a way to verify the association between the resource holder and their Internet Number resources (IP Addresses, ASNs)

- ROA: *Route Origin Authorization* is an attestation of a BGP route announcement
- ROV: *Route Origin Validation* is the application of RPKI to validate the Origin AS

Benefits:

- Validate the prefix announcements are coming from the legitimate Internet Number Resource holder
- Prevent route hijacking (malicious intent or accidental)
- Minimize common routing errors

RPKI integration into Internet Exchange

Architecture

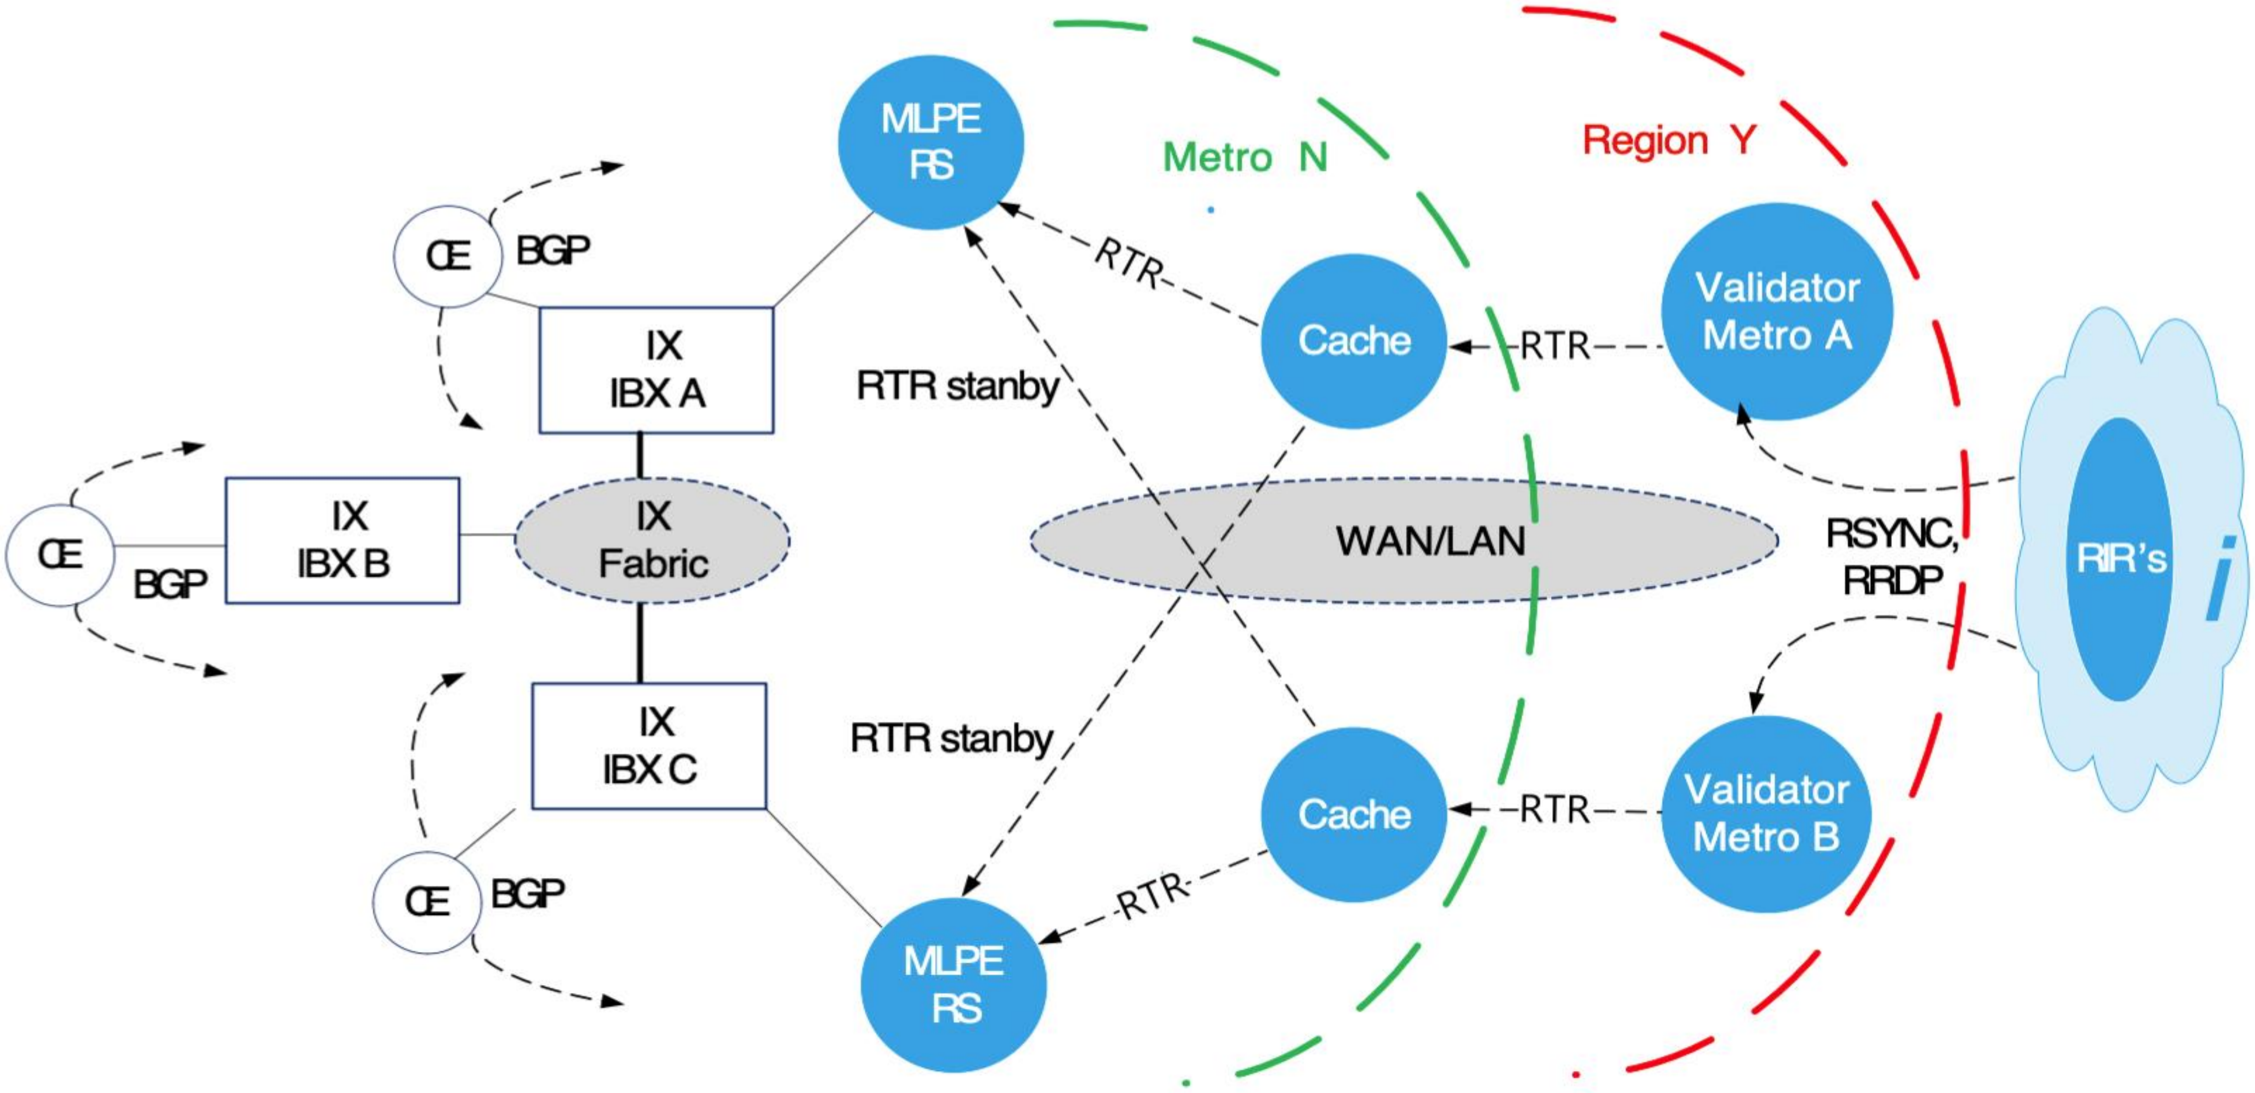

Route validation workflow

ROV (Route Origin Validation)

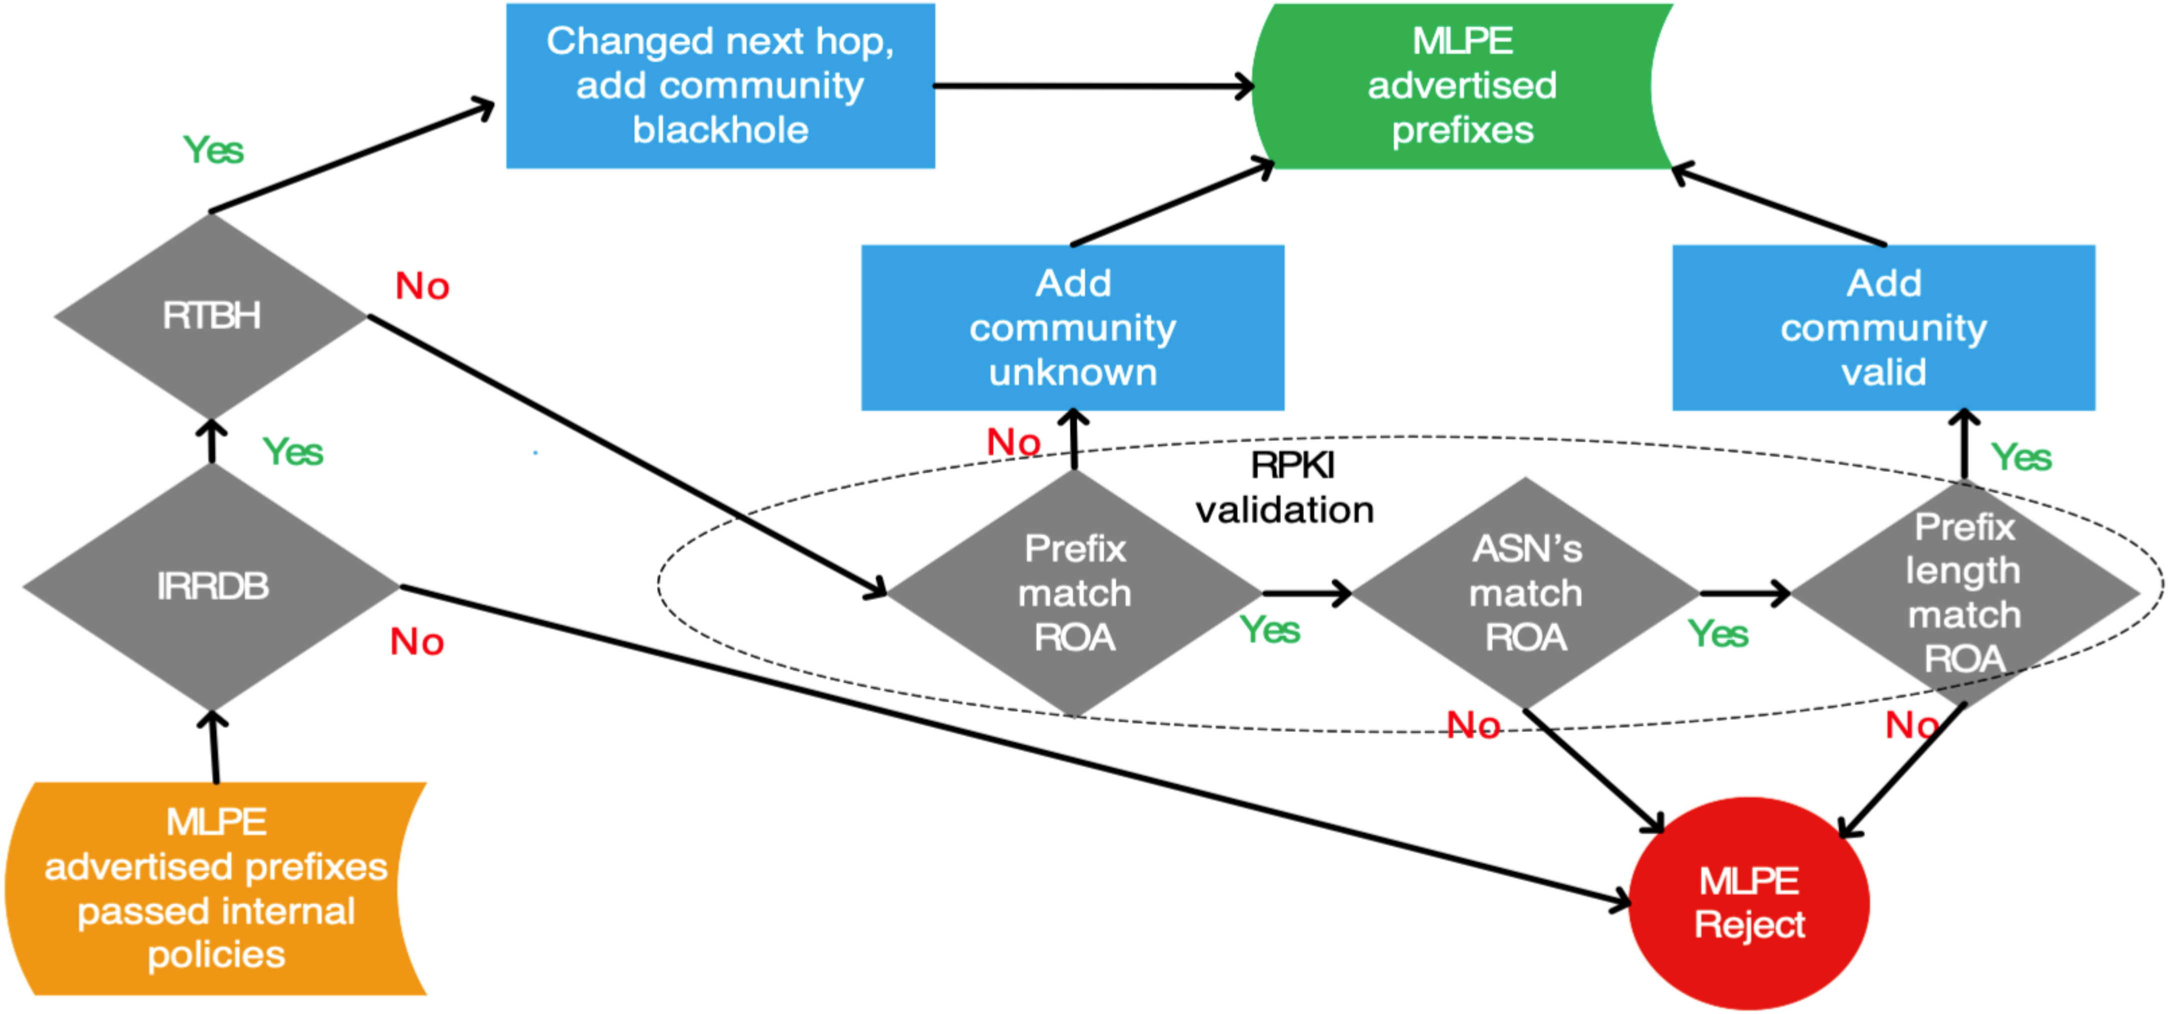

EIE RTBH

EIE RTBH

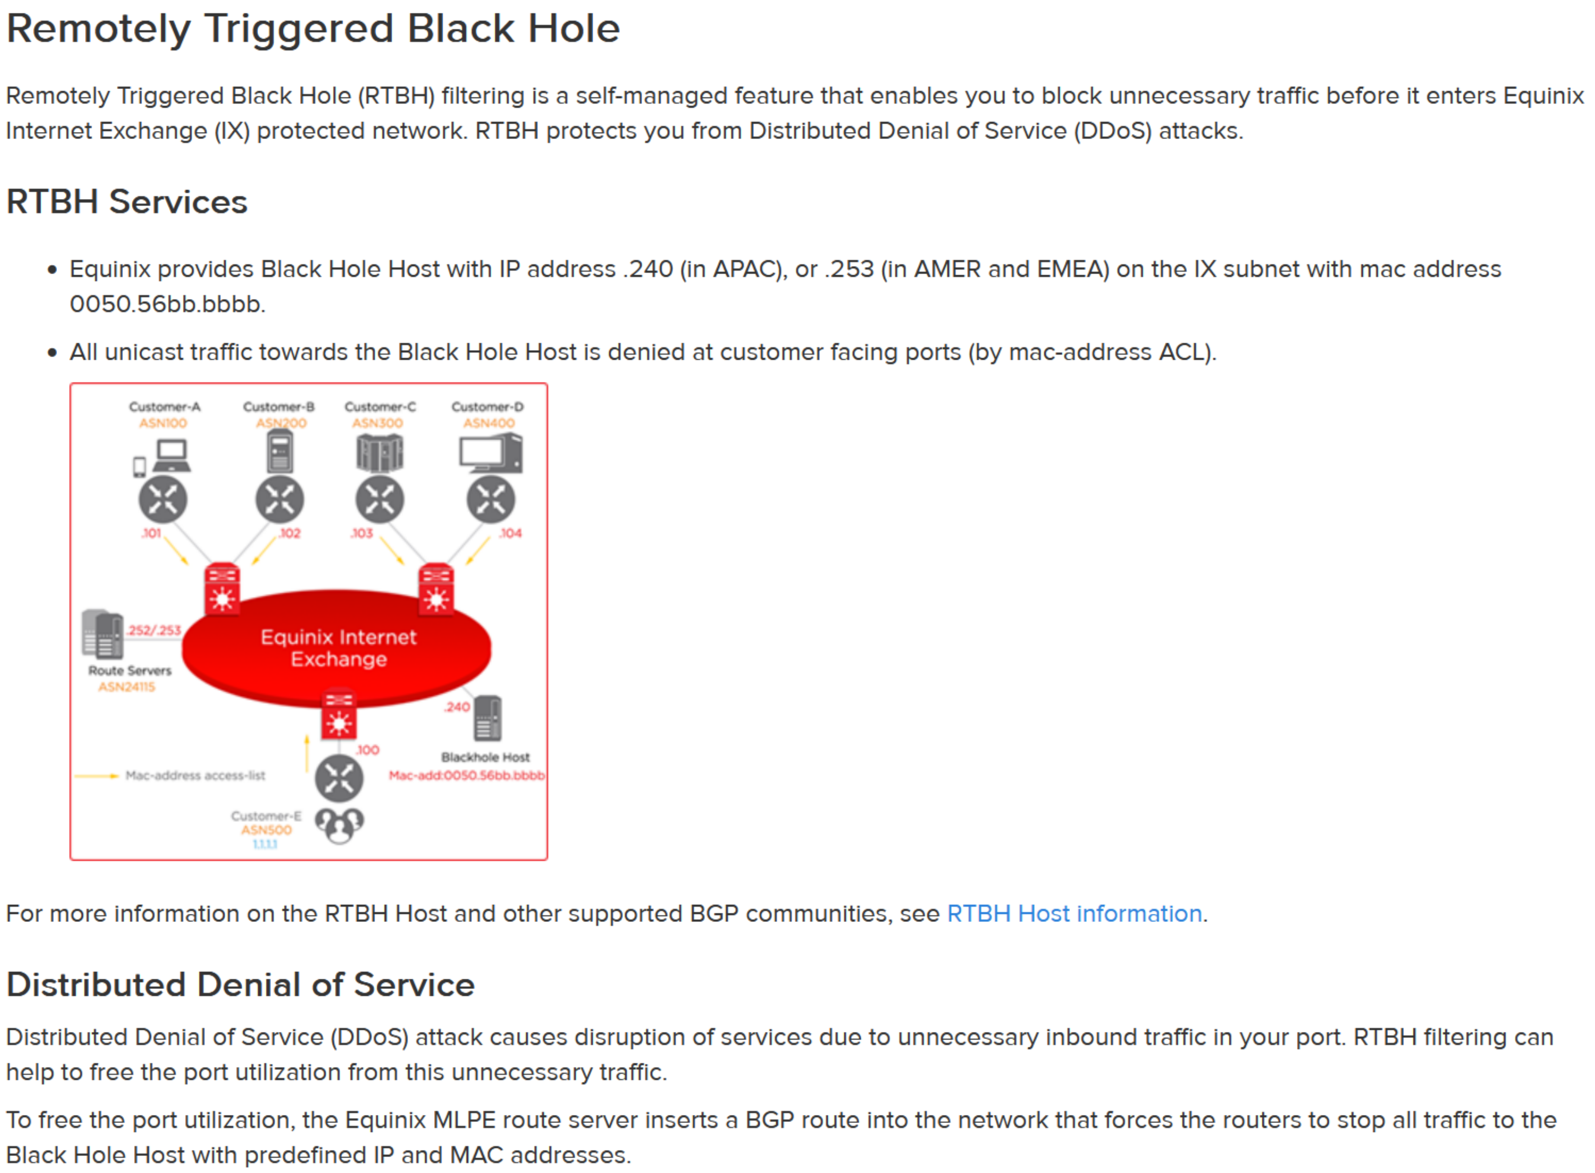

Remotely Triggered Black Hole

Remotely Triggered Black Hole (RTBH) filtering is a self-managed feature that enables you to block unnecessary traffic before it enters Equinix Internet Exchange (IX) protected network. RTBH protects you from Distributed Denial of Service (DDoS) attacks.

RTBH Services

- Equinix provides Black Hole Host with IP address .240 (in APAC), or .253 (in AMER and EMEA) on the IX subnet with mac address 0050.56bb.bbbb.
- All unicast traffic towards the Black Hole Host is denied at customer facing ports (by mac-address ACL).

For more information on the RTBH Host and other supported BGP communities, see RTBH Host information.

Distributed Denial of Service

Distributed Denial of Service (DDoS) attack causes disruption of services due to unnecessary inbound traffic in your port. RTBH filtering can help to free the port utilization from this unnecessary traffic.

To free the port utilization, the Equinix MLPE route server inserts a BGP route into the network that forces the routers to stop all traffic to the Black Hole Host with predefined IP and MAC addresses.

Mitigation Stages

To mitigate the risk of DDoS attacks, RTBH involves the following stages:

▼ Mitigation Stage 1

1. You announce 1.1.1.1/32 with Black Hole BGP community 65535:666.
2. MLPE route servers modify these prefix announcements (tagged with 65535:666) with next-hop to .240 (in APAC) or .253 (in AMER and EMEA), and re-announce the same prefix to other peering participants.

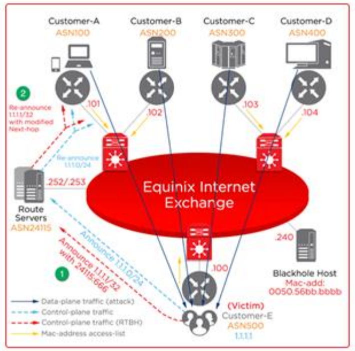

▼ Mitigation Stage 2

1. Peering partners start to resolve next-hop IP address .240 (in APAC) or .253 (in AMER and EMEA) to reach 1.1.1.1.
2. Black Hole Host replies with an ARP with mac-address 0050.56bb.bbbb.

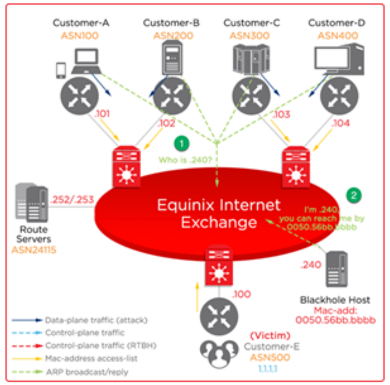

▼ Mitigation Success

1. The attack traffic with next-hop .240 (in APAC) or .253 (in AMER and EMEA) is stopped by Equinix IX switch inbound access list.
2. The DDoS attack going through your switch port is mitigated.

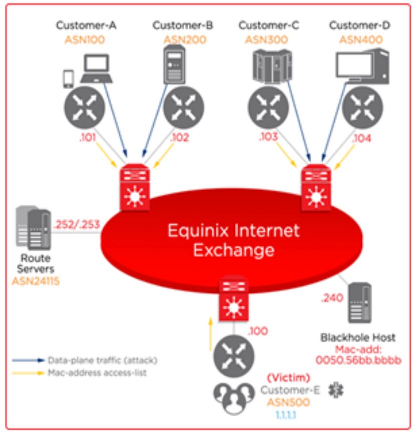

EIE with EVPN Control-plane (Future Roadmap)

EIE with EVPN Control-plane

Benefits and Expectation:

- No MAC learning on customers ports, MAC static configured, MAC ACL’s
- BUM storm control
- EVPN/VXLAN loop prevention, detection and black-hole looped MAC’s
- Under 50 ms advertise or withdraw MAC in metro
- Simple BGP peering schema preferred with no extra hardware or RR
- Support tools RTBH, SFLOW
- EVPN Ethernet segment support, EVPN-multihoming (optional).
- EVPN proxy ARP/NB Dynamic and configured, anti-spoofing, IP duplicate detection/black-hole (optional)
- Interconnect to other EVPN domains and classic L2 domains
- IP subnet scaling

EIE with EVPN Control-plane

Design Concept and Implementation Plan:

- Spine - leaf topology, clients on leaf switches only
- Underlay network ISIS-SR (SR is optional)
- Overlay network eBGP, unique ASN# for every switch
- EVPN control plane, VXLAN encapsulation, VTEP on leaf switches only
- MAC and IP (optional) advertisements over eBGP control plane
- 1,2,4 types EVPN
- L2 connectivity between leaf switches utilized via L3 spine-leaf topology
- Spines switches re BGP “listeners” dynamically accepting peering from leaf switches loopbacks
- Loopback0 for VTEP, Loopback 1 for eBGP peering
- Underlay and overlay network peering sessions are password protected with SHA256

EIE with EVPN Control-plane

Topology Diagram:

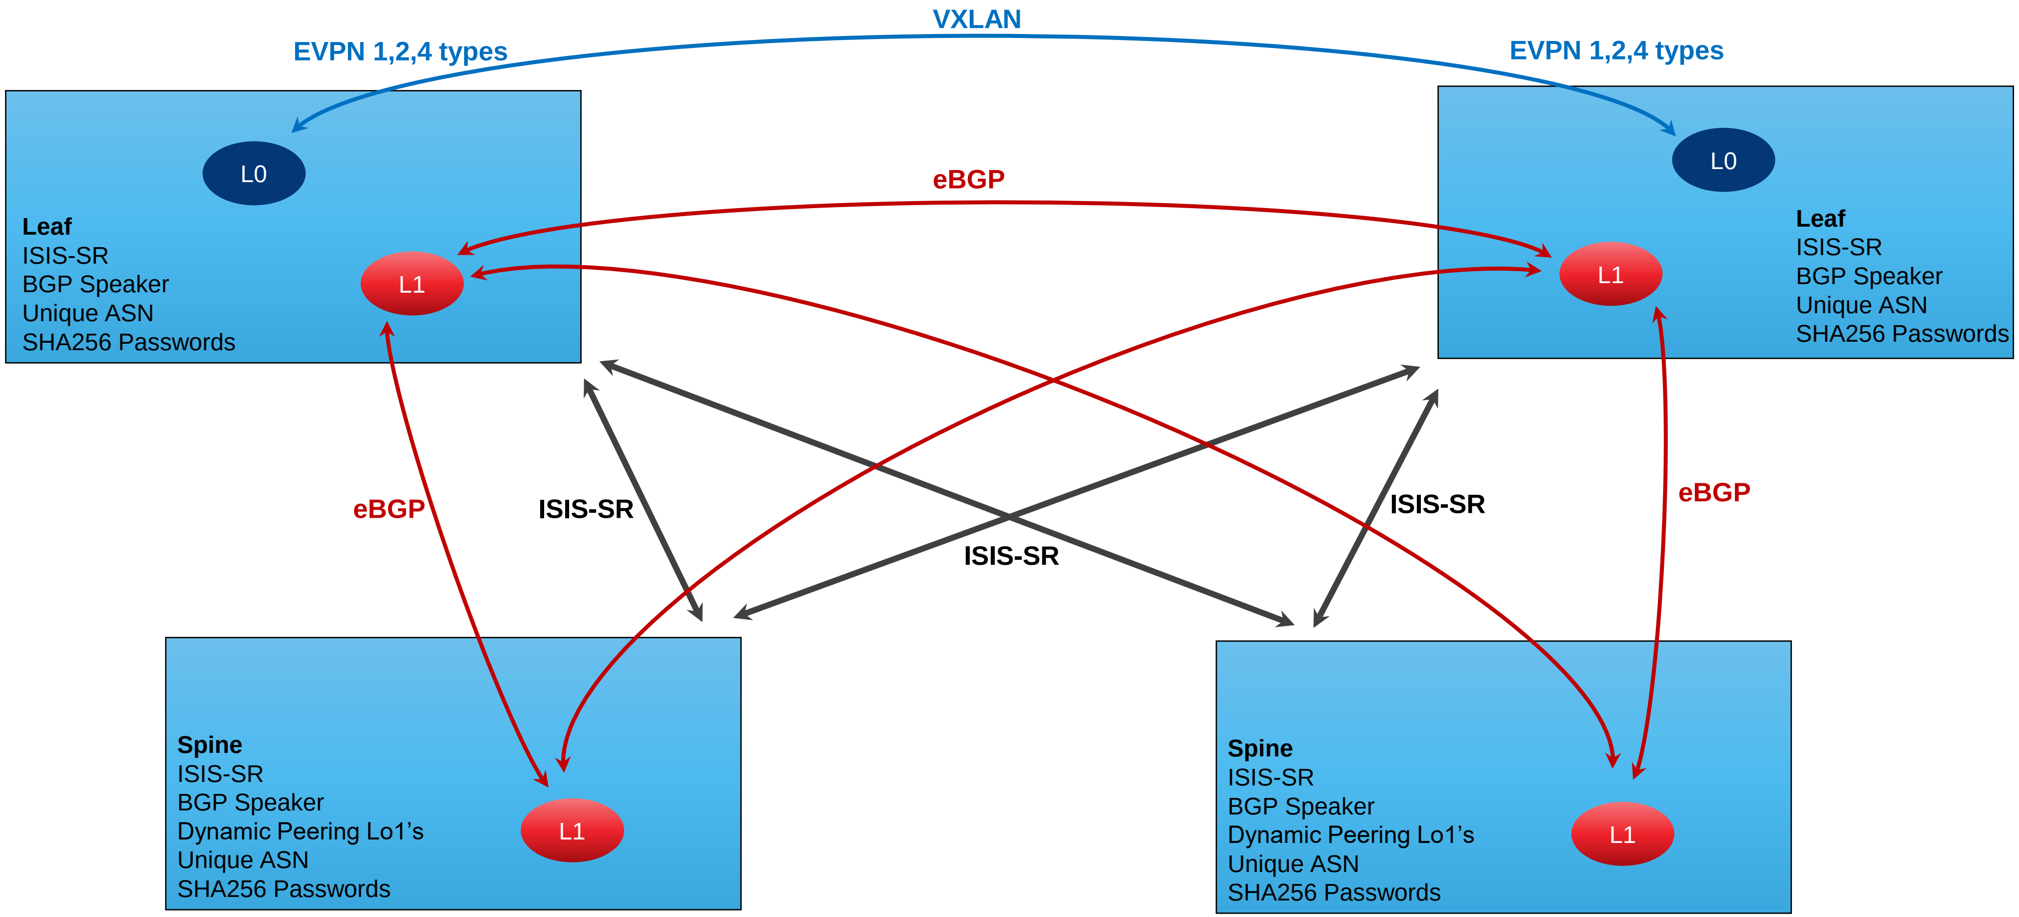

?

INTERCONNECT EVERYONE
INTEGRATE EVERYTHING
REACH EVERYWHERE
PLATFORM EQUINIX®

Thank You