

Next Gen iBGP Route Reflectors

Experiences and views from a Service Provider

Amit Dhamija

Network Engineering , Product & Technology - Infrastructure Design
Bengaluru, India

amit.dhamija@colt.net

Agenda



- Route Reflection Platforms
- Internet RR Behaviour with iBGP Multipath
- Add Path and BGP Optimal Route Reflection.
- COLT Route Reflection Deployment

Route Reflection Models In Use


- Hardware Based
- Either deployed in existing core routers
- Or deployed as dedicated routers

Generally Used Platforms :-

Juniper JCS/SRX

Cisco ASR-1001

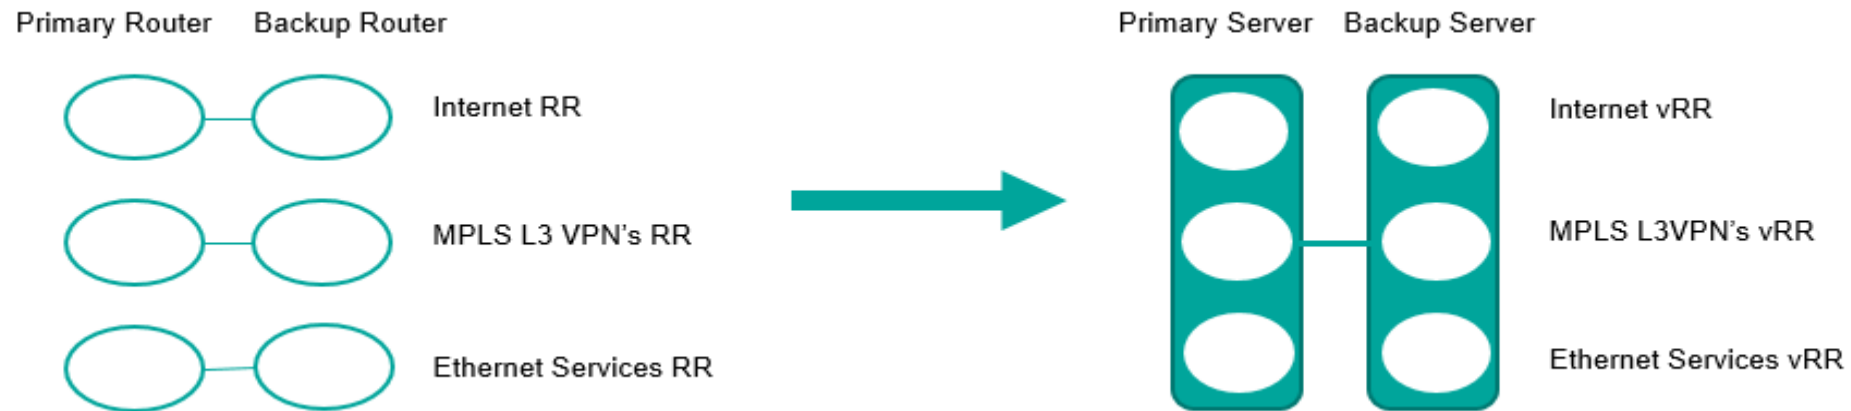
Cisco 7200



New Virtualized Route Reflector (vRR)

Operators prefer to keep separate RR for each service to address following concerns:-

- Scalability
- Convergence
- Security
- Resiliency



Now we are moving from using dedicated RR hardware for each control plane to virtualized RR Platform

Need for vRR

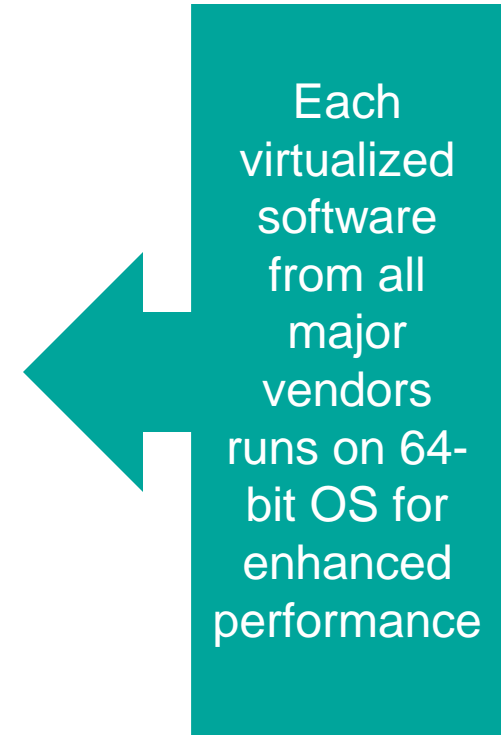
- To eliminate the Hardware limitations.
- To address control plane memory requirements.
- To save Infrastructure Space.
- To deploy the Out of path topologies RR for better convergence.
- To Leverage the Commodity hardware.
- Innovating through software.

Benefits

- Scalability (64b OS), Performance (Multi-core support), Independence of Operations and BGP Software consistency are not compromised.
- New virtualized RR can run on commodity x86 Hardware using VM Image/VMware ESXI, KVM, Citrix XenServer, Microsoft Hyper-V etc.

vRR - Products

- Cisco
 - CSR –1000v
 - runs standard IOS-XE Software
 - IOS-XRv 9000
 - runs standard IOS-XR Software
- Juniper
 - vRR and vMX
 - runs standard Junos Software
- Nokia
 - vSR-RR
 - runs standard SR-OS Software

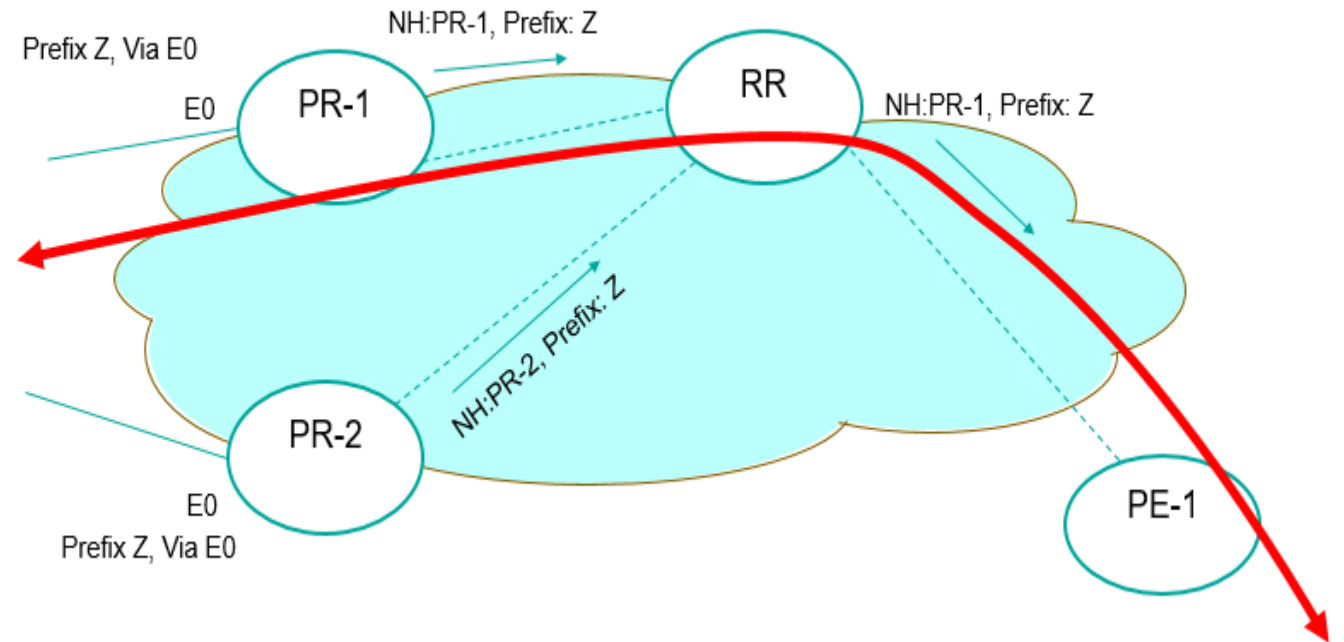


Each virtualized software from all major vendors runs on 64-bit OS for enhanced performance

Internet RR Behaviour with iBGP Multipath

Current Route Reflector Behaviour

- RR picks route that is considered best from it's point of view.
- RR does best path algorithm and advertise only one update to the client PE, which results in suboptimal routing.
- RR's are deployed based on exit points in network, either on multiple Core Routers or multiple out of path.

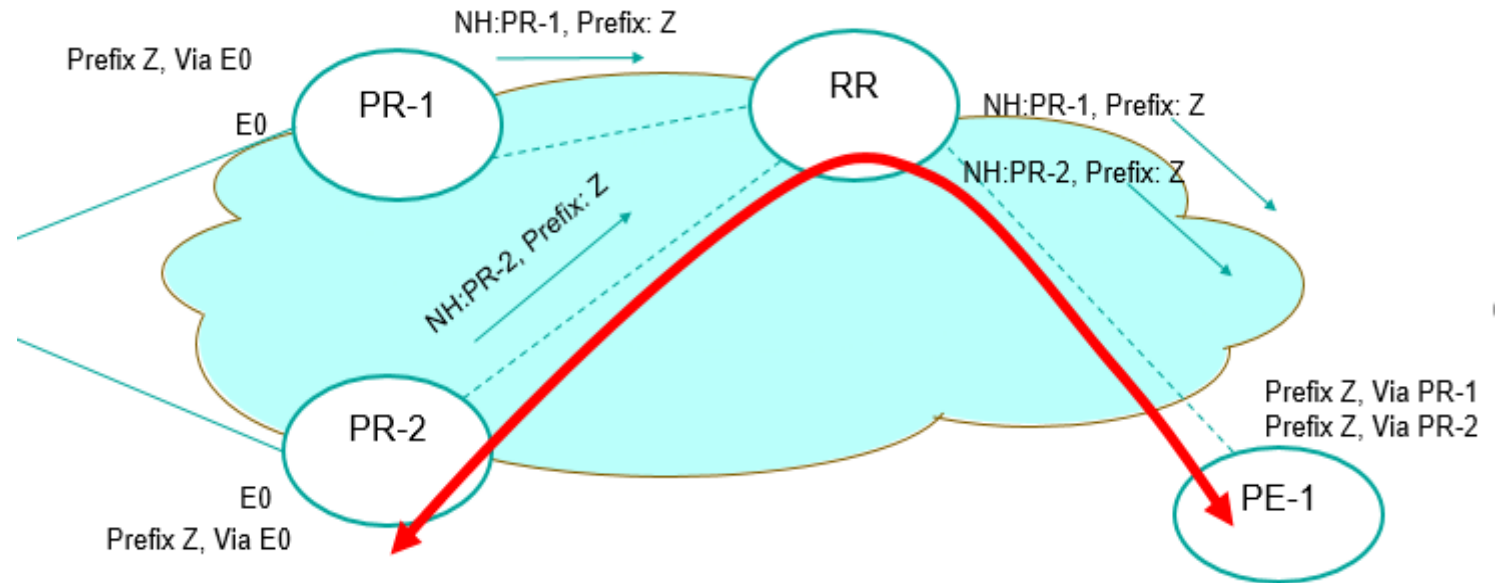


Solutions

Option 1: (BGP Add- Path)

<https://datatracker.ietf.org/doc/rfc7911/>

- Add-Path will signal the diverse path.
- Require support from both the RR and the clients as this is additional capability to be negotiated between two devices.



Advantages

- ✓ Reduce Routing Churns
- ✓ Faster Convergence
- ✓ Better Load Sharing
- ✓ Support & Availability

Caveats

More Load on the Control Plane

Increase Memory Requirements on the end devices

BGP Add-Path Support



- Cisco
 - IOS-XE - supported from 3.7S onwards.
 - IOS-XR - supported in 12K/CRS/ASR-9K started from Rel 4.0 onwards.
 - NX-OS - supported from 6.0(2) onwards.

- Juniper
 - MX/M/T/SRX Series-Platforms - supported from Junos 11.4 or later

- Nokia
 - 7750 SR-OS - supported from 10.0 R1 onwards

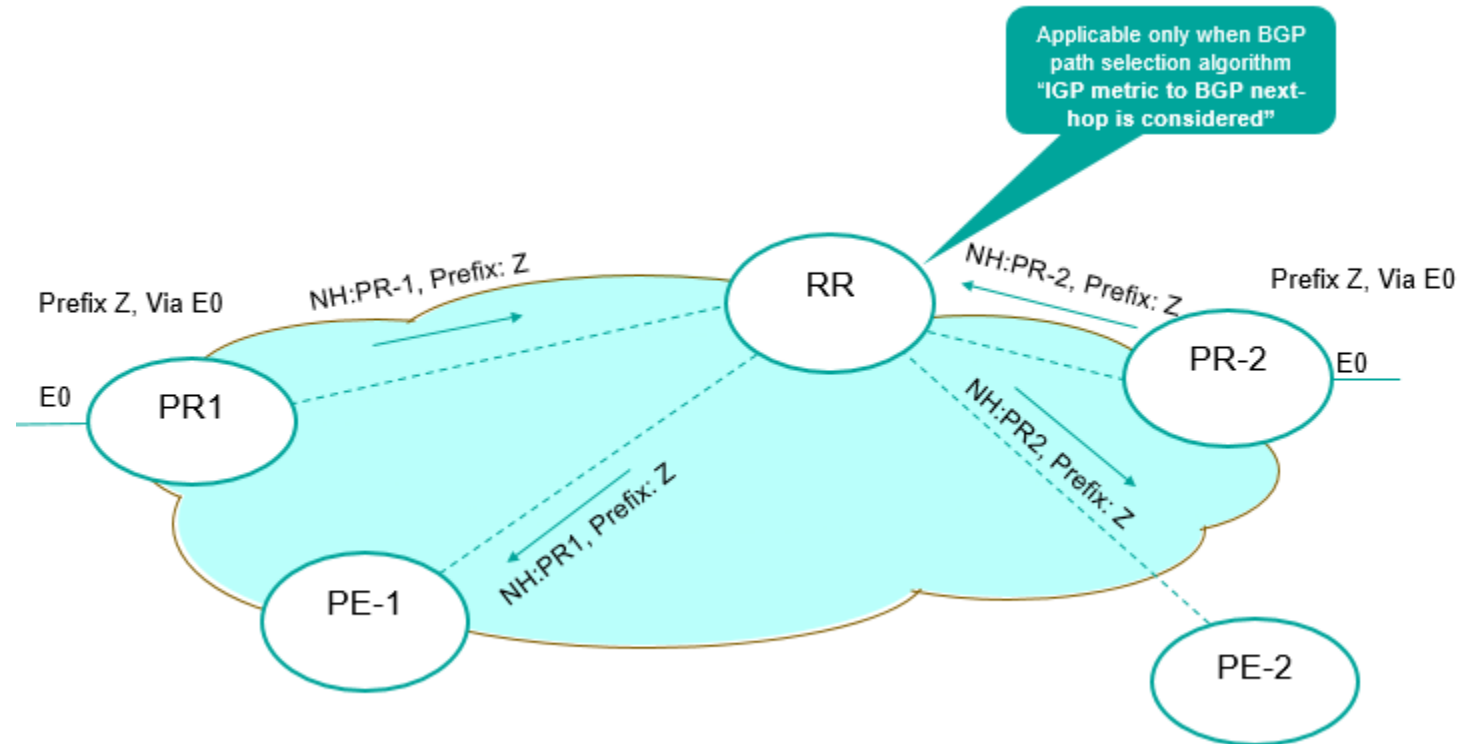
Supported for all the address families

- IPV4 unicast and IPV6 unicast (including labelled IPV4 and labelled IPV6 routes)
- VPN-V4 and VPN-V6

Option 2: (BGP Optimal Route Reflection)

<https://datatracker.ietf.org/doc/draft-ietf-idr-bgp-optimal-route-reflection/>

- Allow Client Specific best path.
- RR runs SPF multiple times, one per each RR client BGP Speaker.
- BGP best path mechanism modified to compute best path per RR Client.
- BGP Route advertisement modifies to announce best path to client.



All control on RR no help is required from client routers

Topology free route reflectors

- Cisco
 - Cisco -IOS-XR – 6.1.1 onwards (Supported in Hardware and Virtualized platforms)
 - Supported for both IGP's OSPF and ISIS
 - Not supported in any IOS-XE and NX-OS as of now.
- Juniper
 - Juniper JunOS 15.1F4.15 onwards
 - Supported for both IGP's OSPF and ISIS
- Nokia
 - Nokia SR-OS 15.0 onwards
 - Supported for both IGP's OSPF and ISIS

BGP ORR – Advantages & Caveats



Advantages

- ✓ Gives flexibility to place RR any where in topology.
- ✓ Solves Hot potato Routing.
- ✓ Supports Resiliency for ORR Groups.
- ✓ Require no support from clients.
- ✓ Even better with ADD-PATH.

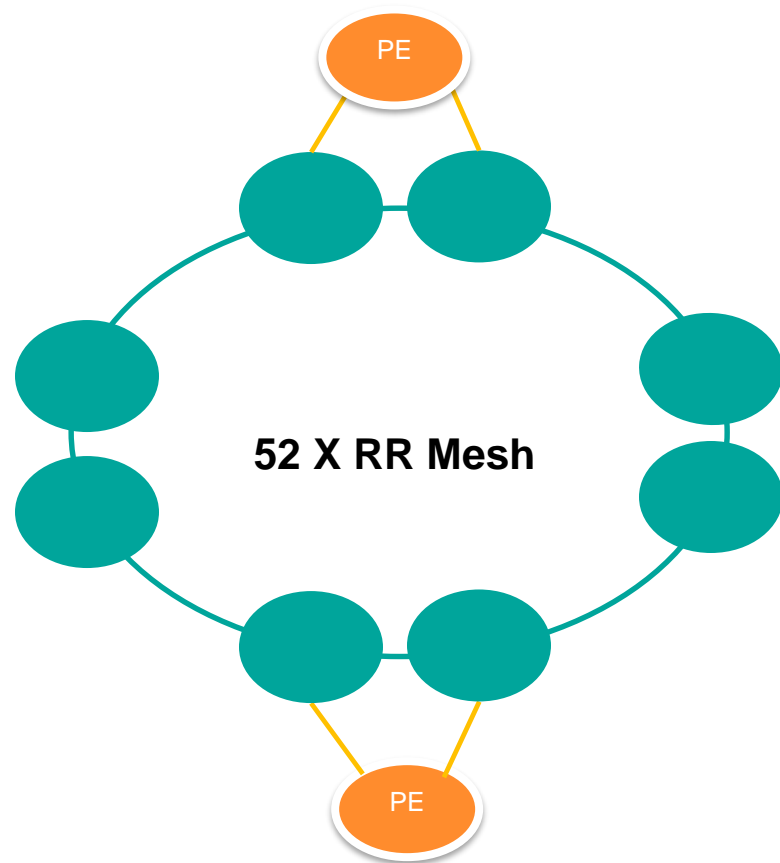
Caveats

No Support for MPLS

No Support for 6PE and 6VPE Technology

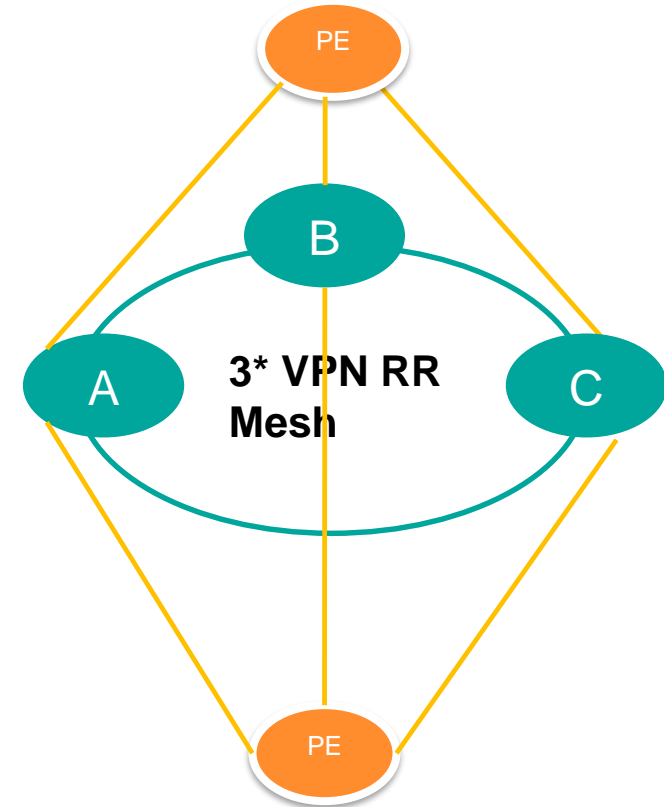
COLT Route Reflector Deployment

Current COLT RR Design



Internet RR Design:

- 52 IRR, 26 Cities/ Region
- Hierarchal design.
- Core acts as IRR
- Full mesh between all the Cores
- PE BGP to nearest RR region pair
- Bigger Scaling Issues

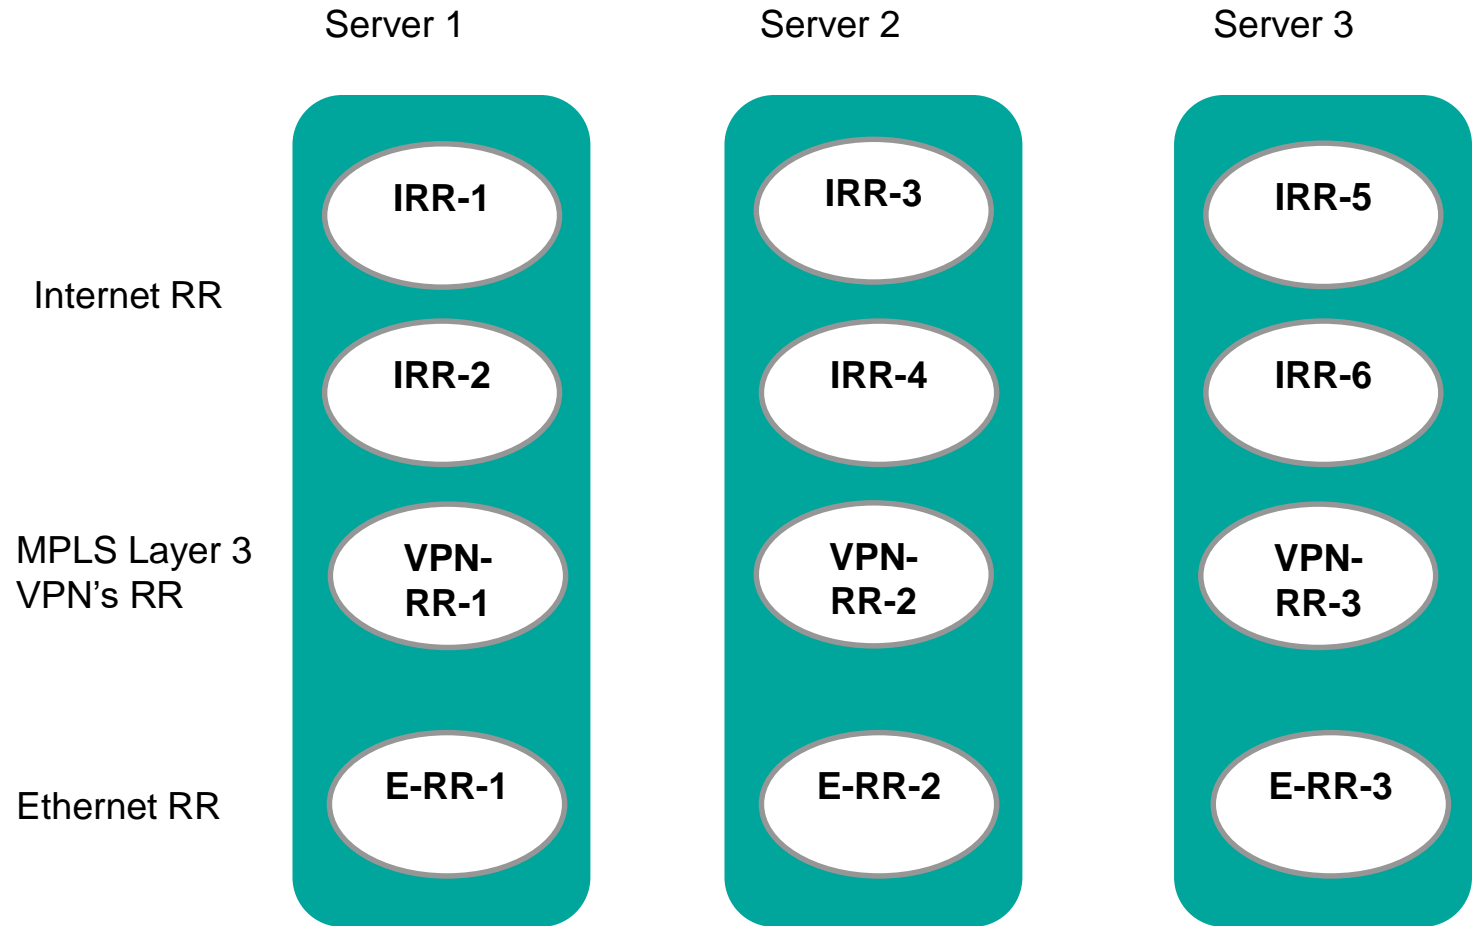


IP-VPN RR Design:

- Centralized with 3 dedicated H/W
- Full mesh between the 3 RR's.
- PE BGP to all three RR's.
- Eox, Scale and Feature Limitation

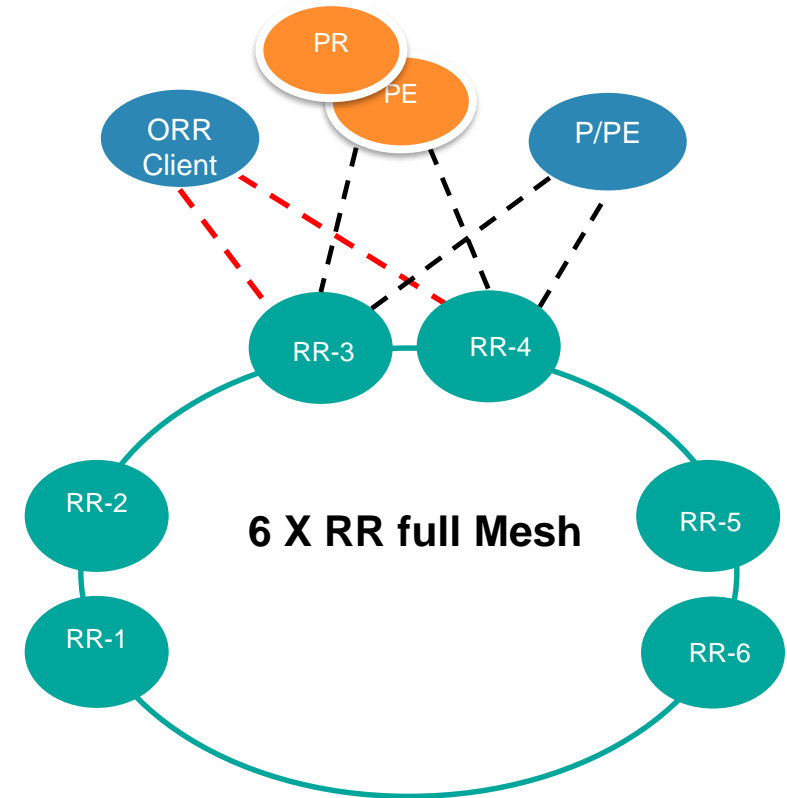
Colt New vRR Design

- Virtualizing all the three RR's on KVM Hypervisor.
- Three Servers in tier 1 locations.
- Dedicated VM's for each RR per Server.
- BGP-ORR feature on Internet vRR for Optimal Routing.



Design Consideration for BGP-ORR

- Each ORR Client is associated with Single BGP Update Group.
- Ideal design is to have maximum 15-20 ORR Groups per RR.
- Choose the ORR points in network to achieve the “Hot Potato Routing”.
- 6PE Hot potato routing can be achieved using the “BGP Add-Path feature”.



Thank you

amit.dhamija@colt.net

