White Paper

Virtual Route Reflector Using Cisco IOS XRv

Last Updated: February 2014

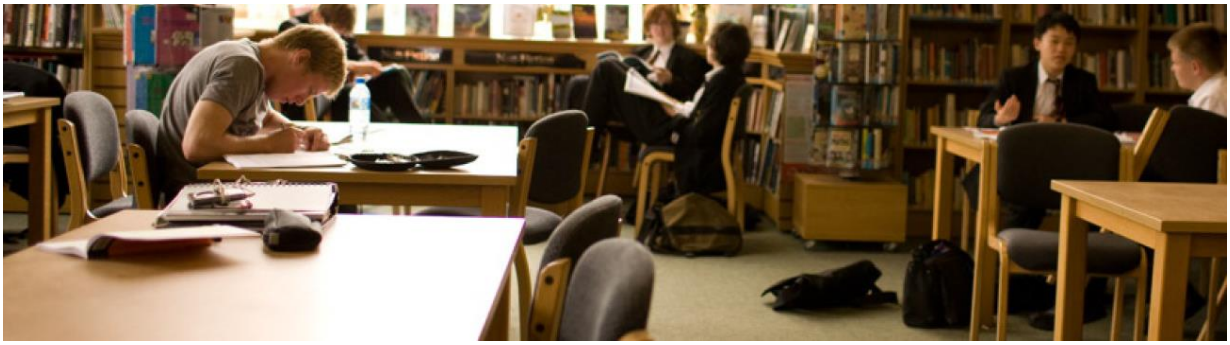

Service providers and enterprises worldwide have deployed and utilized Border Gateway Protocol (BGP) for an increasing number of services like Internet peering, Layer 2/Layer 3 VPN, BGP Link State harvesting, and security policy propagation by BGP flow specification. However, scaling a large iBGP full-mesh peering requires optimizing the control plane architecture. BGP route reflection as defined in RFC 4456 has emerged to become the de facto choice for scaling full-mesh iBGP deployments.

Traditional route reflectors were deployed using either core routers or dedicated physical hardware solely for control plane route-reflection purposes. While this has proved to be a viable solution that meets the demands of BGP with all the CPU and memory required, it lacks the flexibility, elasticity, and agility to meet the frequently changing demands of the services.

Deploying a virtualized network function (VNF) in the form of a route reflector is proving to be the next logical evolution to address the shortcomings brought about by physical device limitations regarding elastic flexibility of device resources such as memory and CPU.

Introducing Cisco IOS® XRv—a virtual router that uses the same carrier-class Cisco IOS XR operating system powering the ASR 9000 and CRS series high-end routers since 2004. It provides the same key benefits as Cisco IOS XR while providing the elasticity, agility, and flexibility that a VNF brings.

Introduction

Service providers and enterprises worldwide rely on Border Gateway Protocol (BGP) as an integral part in the deployment of services that are offered to their customers such as Internet peering, Layer 2/Layer 3 VPN services, MVPN, and so on. BGP offers the robustness, stability, and flexibility, that other upcoming technologies make use of for loop-free transport of information within the Autonomous System (AS) and across AS boundaries. Examples are technologies like BGP Link State, BGP Flow-Spec.

Deploying BGP comes with a caveat, as fully meshed iBGP neighbors are required to ensure a properly working and stable network allowing awareness of all paths in a loop-free control plane transport. This poses a significant scalability issue especially when the number of iBGP nodes increase. To address this scalability problem, the use of route reflectors (RRs) (RFC 4456) in the network has emerged as one of the solutions that can be used to reduce the total number of BGP peers in the network.

Deployment of a route reflector may be done using either dedicated or non-dedicated routers. In smaller POPs, it is not uncommon to deploy a non-dedicated RR solution, utilizing the control plane of the core routers as a route reflector and at the same time deploying it for other services. Deploying a dedicated RR, however, has other advantages as it could provide better stability, scalability, and convergence.

Dedicated RRs, often also called off-path RRs or RR-on-a-stick, demand that the node be capable of CPU-intensive calculation for path computation and have ample memory to store all the learned routes. Bandwidth is a key factor for communication of BGP path updates; however, throughput is not a key factor since the RR nodes in this case are not in the forwarding path.

Virtual Route Reflector

Virtualization technologies have been widely deployed in enterprise data centers. It has proven to be a reliable technology that can be extended and deployed onto critical and highly available network infrastructure.

Virtualization helps to achieve reduced time to market, lower capital expenditures (CapEx), and lower operating expenses (OpEx) without sacrificing reliability and stability. Virtualizing the route reflector function is a great application for network function virtualization (NFV) implementation.

A virtualized route reflector (vRR) is an excellent solution for dedicated off-path RR application. In order to meet the intensive CPU and memory demand to handle path calculations on a large number of prefixes, CPU and memory resources can easily be carved out and provisioned according to requirements.

Another aspect that is increasingly becoming an important consideration in data centers or collocation facilities is power and cooling. Scaling up the number of vRR instance doesn’t increase the power consumption significantly as the common physical components like fans, processors, remain the same.

A virtual route reflector also allows us to position the RR within the highly-available data center and still have logical placement at those locations that are deemed most useful for the customer routing architecture.

Cisco IOS XRv Virtual Route Reflector

With Cisco IOS XR's proven track record of scalability, stability, and longevity, Cisco IOS XRv uses the same carrier-class operating system that has powered the ASR 9000 and CRS series high-end routers since 2004. Cisco IOS XRv is based on a virtualized Cisco IOS XR router that runs in a VM environment on top of hypervisors such as VMware ESXi and QEMU/KVM.

Cisco IOS XRv provides many of these key Cisco IOS XR benefits, including but not limited to:

- Comprehensive BGP address family (AFI) and subsequent address family (SAFI) support (see Table 1 for a full list)
- Modular operating system provides fault isolation not possible with monolithic network operation system
- Process restartability allows restarting a process and not affecting other processes and BGP process alike
- High-performance multithreaded operations—the BGP process is multithreaded allowing Cisco IOS XRv to be highly scalable
- Consistent command-line interface (CLI) with existing Cisco IOS XR Platforms—ASR 9000 and CRS
- Rich route policy enforcement using the simple yet powerful Route Policy Language (RPL)
- Live software patch through software maintenance updates (SMU)
- Support for multi-instance BGP, providing process-level isolation and independent scale for different BGP address families.
- Twenty-four hour access to the Cisco experts at Cisco's Technical Assistance Center (TAC)[1]

The Cisco IOS XRv can be deployed with existing network infrastructure allowing it to be deployable from day one. As the service scale requirement of the network increases, additional Cisco IOS XRv instances can easily be spun up on-demand manually or through means of orchestration within a matter of minutes proving it to be a highly elastic and scalable solution.

Table 1. SAFI Supported (for IPv4 and IPv6 AFI)[2]

SAFI	Description
1	Network layer reachability information (NLRI) used for unicast forwarding
2	Network layer reachability information used for multicast forwarding
4	Network layer reachability information with Multiprotocol Label Switching (MPLS) labels
5	MCAST-VPN (NG mVPN)
8	MCAST-VPLS
65	Virtual Private LAN Service and Virtual Private Wire Services (VPLS and VPWS)
66	BGP MDT SAFI (mVPN)
70	BGP EVPNs (including PBB-eVPN and EVPN)
71	BGP-LS
128	MPLS-labeled VPN address (Layer 3 VPN)
132	Route Target constraints

[1] Subject to support contract deliverables
[2] As of February 2014

Where Can I Get Cisco IOS XRv?

Cisco IOS XRv comes in two versions—a Demo version and a Simulation/Production version.

- Demo version is available for download from Cisco.com.
  - Free version with limited throughput ideal for Cisco IOS XR training and learning environments (no support).
- Simulation/Production version can be purchased through Cisco Commerce Workspace.
  - Licensed version for higher/full throughput and vRR deployment (supported).

For More Information on Cisco IOS XRv

http://www.cisco.com/go/iosxrv

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-731132-00 02/14