CISCO VALIDATED PROFILE

Routing
Route Reflector Vertical

April 2016

Table of Contents

Profile Introduction

Cisco is transforming the network edge with Cisco ASR 1000 Series Aggregation Services Routers and Cisco 4000 Series Integrated Services Routers (ISRs), new lines of midrange routers that establish a new price-to-performance class offering, benefiting both enterprises and service providers. These routers provide a great opportunity for simplifying the WAN edge and significantly decreasing network operating expenses (OpEx). By efficiently integrating a critical set of WAN edge functions such as WAN aggregation, Internet edge services, firewall services, VPN termination, etc. into a single platform, enterprises can meet their business objectives by facilitating deployment of advanced services in a secure, scalable, and reliable manner while minimizing the total cost of ownership (TCO).

Cisco WAN aggregation solutions distinguish themselves from other solutions by offering multiservice routers with the highest performance, availability, and density for concurrent data, security, voice, and application-acceleration services with maximum headroom for growth. The solutions feature embedded security, performance, and memory enhancements, and high-performance interfaces featuring the latest WAN technologies can help enterprises meet the needs of the most demanding WAN network.

Route reflection is the operation of a border gateway protocol (BGP) speaker advertising a route that was learned through an internal BGP (iBGP) session to another iBGP peer. This practice is prohibited in a normal BGP operation because the traditional BGP routing protocol had no safeguards against routing loops within an autonomous system (thus requiring a full mesh of iBGP speakers). A BGP speaker that propagates iBGP routes to other iBGP peers is called a *route reflector* (RR), and such a route is called a *reflected route*.

Typically, all BGP speakers within a single autonomous system (AS) must be fully meshed and any external routing information must be re-distributed to all other routers within that AS. Expressed mathematically, an AS with *n* fully meshed BGP speakers requires you to create and maintain *n*(n-1)/2* unique iBGP sessions. This full mesh requirement clearly does not scale when there are a large number of iBGP speakers each exchanging a large volume of routing information, as is common in many networks. This scaling problem has been well documented, and a number of proposals have been made to alleviate this. This document represents another alternative in alleviating the need for a full mesh and is known as *route reflection*. This approach allows a BGP speaker (known as a *route reflector*) to advertise iBGP learned routes to certain iBGP peers. It represents a change in the commonly understood concept of iBGP and the addition of two new optional non-transitive BGP attributes to prevent loops in routing updates.

With the introduction of route reflectors, the scaling of the MP-iBGP sessions becomes easier because the full-mesh requirement is eliminated. However, there will always be a finite number of sessions that can be made to the route reflector, so a hierarchical structure of route reflectors may be desirable; the design of the network should be capable of catering to its introduction. The actual number of sessions that a route reflector can service without affecting its functionality or its capability to reflect routing information is difficult to judge. It depends on many factors, such as the number of routes per session, the use of peer groups, the CPU power, and the memory resources of the route reflector. Table 1 lists the key areas on which the profile focuses.

Table 1 *Profile feature summary*

Deployment areas	Features
Network services	MBGP, MPLS with LDP & OSPF
Address families	IPv4, IPv6, VPNV4 & VPNV6
Network planning	Failover method used with primary and secondary RR
Routing	EBGP & iBGP

Network Profile

Based on the research, customer feedback, and configuration samples, the profile is designed with a generic deployment topology that you can easily modify to fit any specific deployment scenario. Refer to the topology for further details.

TOPOLOGY DIAGRAM

Figure 1 *Route Reflector Profile: topology overview*

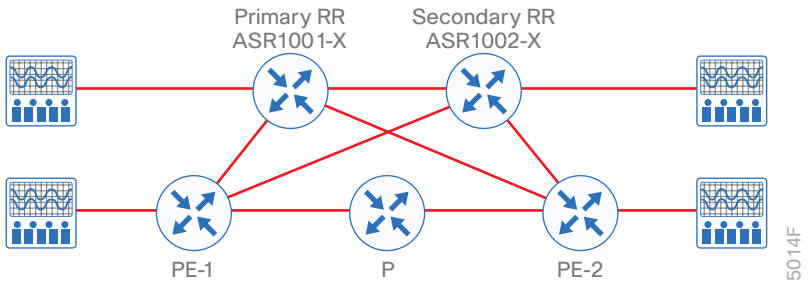

Disclaimer

The links between the different network layers in the topology are mainly to facilitate this profile validation across different platform combinations, and the actual deployment could vary based on specific requirements.

HARDWARE & FEATURE SPECIFICATIONS

This section describes the 3-D feature matrix where the hardware platforms are listed along with their place-in-network (PIN) and the relevant vertical deployed.

Key Vertical Features

Table 2 defines the hardware, PIN, and the features deployed. The scale of these configured features, the test environment, the list of endpoints, and the hardware/software versions of the network topology are defined later.

The following physical topology is used for route reflector profile testing. Scale and performance numbers were taken for the following devices, which were used to test the Active and Standby UUT as depicted in the topology:

- ASR1001-X (8GB Memory) as Primary RouteReflector
- ASR1002-X (8GB Memory) as Secondary RouteReflector

The following device simulate the ISP with which the edge router peers:

- ASR1004 with RP2/ESP40 simulating the ISP1
- ASR1004 with RP2/ESP40 simulating the ISP2

Hardware Profile

Table 3 defines the set of relevant hardware, servers, test equipment, and endpoints that are used to complete the Route Reflector Vertical Profile deployment.

This list of hardware, along with the relevant software versions and the role of these devices, complement the actual physical topology that is defined in Figure 1 of the previous section.

IXIA used for scaling route reflector clients to push the routes.

Table 2 Hardware profile of servers and endpoints

VM and HW	Software versions	Description
Ixia	IxNetwork,IxLoad and IxExplorer version 6.20	Generate traffic streams

TEST ENVIRONMENT

This section describes the features and the relevant scales at which the features are deployed across the physical topology. Table 4 lists the scale for each feature.

Disclaimer

The table below captures a sample set of scale values used in one of the use cases. Refer to appropriate CCO documentation/datasheets for comprehensive scale data.

Table 3 Route Reflector Profile: features

Feature	Scale
IPv4 unidimensional scaling	5.25 M BGP neighbors, 350 seconds convergence time
IPv6 unidimensional scaling	5.25 M BGP neighbors, 624 seconds convergence time
VPNV4 unidimensional scaling	4.25 M BGP neighbors, 1212 seconds convergence time
VPNV6 unidimensional scaling	4.25 M BGP neighbors, 1253 seconds convergence time

Use Case Scenarios

TEST METHODOLOGY

The use cases listed in Table 5 are executed using the topology shown in Figure 1, along with the test environment shown in Table 4.

With respect to the longevity for this profile setup, the CPU and memory use are monitored overnight, as well as during the weekends, along with any mem-leak checks. In order to test the robustness, certain negative events are triggered during the use case testing.

USE CASES

Table 4 describes the use cases that are executed on the Route Reflector Profile. These use cases are divided into buckets of technology areas to show the complete coverage of the deployment scenarios.

These technology buckets are composed of system upgrade, security, network services, monitoring & troubleshooting, simplified management, and system health monitoring, along with system and network resiliency.

Table 4 *List of use case scenarios*

No.	Focus area	Use cases
System upgrade		
1	Route reflector upgrade	Network administrator should be able to perform route reflector upgrade and downgrade between releases seamlessly. ▪ All of the configuration should be migrated seamlessly during the upgrade/downgrade operation. ▪ In Service Software Upgrade (ISSU) ▪ Upgrade secondary RR to the new image and then the primary RR
Security		
2	RR failure (primary to secondary SWO and back to primary)	Network admin should bring secondary RR as active and primary RR as backup and bring back the primary RR as active and secondary RR as backup. ▪ RR failure scenario ▪ All of the configuration should be migrated seamlessly during the failover operation.
Network services		
3	LDP	Network admin to bring the CORE MPLS with LDP config ▪ Verifying the RR feature works fine with core LDP
4	iBGP	Network admin to bring the CORE MPLS with iBGP config ▪ Verifying the RR feature works fine with core iBGP

Table 4 continued

5	OSPF	Network admin to bring the CORE MPLS with OSPF config ▪ Verifying the RR feature works fine with core OSPF
Monitoring & troubleshooting		
6	SNMP	Network admin should be able to use SNMP for monitoring ▪ SNMP mibwalk
Simplified management		
7	Troubleshooting & monitoring	Simplify network troubleshooting and debugging for IT admins ▪ Monitor network for alarms, syslogs, and traps
System health monitoring		
8	System health	Monitor system health for CPU use, memory consumption, and memory leaks during longevity
System & network resiliency, robustness		
9	System resiliency	Verify system level resiliency during the following events: ▪ Active RP failure/RP switchover ▪ Active/standby ESP failure ▪ WAN/LAN interface flaps ▪ SIP/SPA reload/OIR
10	Negative events, triggers	Verify that the system holds well and recovers to working condition after the following negative events are triggered: ▪ Config changes–add/remove config snippets, config replace ▪ Routing protocol interface flaps ▪ PE router failure scenario

Appendix A

You can find example configurations at the following location:

http://cvddocs.com/fw/cvpconfig-routing

Please use the feedback form to send comments and suggestions about this guide.

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.