CISCO VALIDATED PROFILE

Routing MPLS L3VPN Service Provider 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. Cisco ASR 1000 Series Aggregation Services Routers transform the service provider and enterprise network edge by delivering industry-leading performance, instant-on service capabilities, and high availability in a compact form factor.

This new line is built to meet the specific needs of the aggregation edge and offers performance, scale, flexibility, security, and programmability while offering cost savings previously unachievable to benefit both service providers and enterprises. The Cisco ASR 1000 Series Router is designed as a carrier-class system for both service provider and enterprise environments with stringent high-availability requirements

The Cisco ASR 1000 Series Routers deliver:

- Highly secure high-performance and integrated software-enabled services
- A new price class for high-performance edge routers
- High resiliency with convenience and cost-saving, in-service software upgrades
- Software redundancy on nonredundant hardware

For service providers, the Cisco ASR 1000 Series facilitates more flexible, efficient, and cost-effective delivery of complex consumer and business services. And for enterprises, it delivers a highly reliable, high-performance WAN edge solution where information, communication, collaboration, and commerce converge.

Service providers (SPs) and Enterprises alike are migrating from existing ATM, frame relay, and time division multiplex infrastructures to an IP-based backbone. Current IP backbones can no longer be designed just to transport IP packets. Instead, next generation (NG) IP backbones must be capable of providing multiple IP services over a single physical infrastructure, using techniques such as differentiated quality of service (QoS) and secure transport layer. In addition, NG IP backbones should provide Layer 2/3 VPNs, IP multicast, IPv6, and granular traffic engineering capabilities. Ultimately, these IP backbones should be scalable and flexible enough to support the mission critical, time-sensitive applications that all modern networks require and to meet new demands for applications, services, and bandwidth.

Multiprotocol label switching (MPLS), when used on an IP backbone, provides the mechanism to offer rich IP services and transport capabilities to the routing infrastructure.

This profile is focused on validating an MPLS-L3VPN service on the SP core network, with MPLS being the prime transport mechanism for the L3VPN service. MPLS in the service provider core requires interior gateway protocol (IGP), label distribution protocol (LDP), multiprotocol-border gateway protocol (MBGP) that uses piggyback mechanisms for exchanging the VPN-labels, and interior border gateway protocol (iBGP) with route-reflector (RR) in the core for transit traffic. The core shall also have resource reservation protocol–traffic engineering (RSVP-TE) enabled for single-hop node/link redundancy.

Table 1 lists the key areas on which the MPLS-L3VPN profile focuses.

Table 1 *MPLS-L3VPN profile feature summary*

Deployment areas	Features
Service provider core	MPLS LDP, LDP-Sync, LDP-Autoconfigure IGP: OSPF (Single-Area) EGP: BGP, iBGP, RR, BGP-Constrained Route distribution (BGP-RTC), IP-SLA, Single and Multihop BFD, NTP, SNMP, Non-Stop Forwarding (NSF), Non-Stop Routing (NSR)
L3VPN and 6VPE	Dual-stack VRFs and MP-BGP on all PEs in the SP-Core CE-PE Routing: EIGRP, OSPF, eBGP, BGP-PIC-EDGE, BGP Multipath redistribution, Single-hop BFD, ACLs, FNF, HQoS, IP-SLA
Multicast services	Multicast-VPNs using SP-CORE native multicast (PIM-SSM) Multicast-LDP (MLDP in the CORE, PIM-SSM inside VRF) Unicast Reverse Path forwarding (uRPF)
Traffic-engineering	RSVP-TE single hop/link redundancy in the CORE
CORE physical interfaces	All Core facing interfaces: 10 Gb/1 Gb Port-channel: 10 Gb member links
EDGE physical Interfaces	Multilink PPP (MLPPP) Frame-relay circuits ATM PVCs GigE dot1q sub-interfaces
Network planning and troubleshooting	NetFlow, EPC, IP-SLA, NTP, SNMP, single and multi-hop BFD,
Efficient network management	Cisco Prime Infrastructure

Network Profile

Based on the research, customer feedback, and configuration samples, the MPLS-L3VPN 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 *MPLS-L3VPN Service Provider Profile: topology overview*

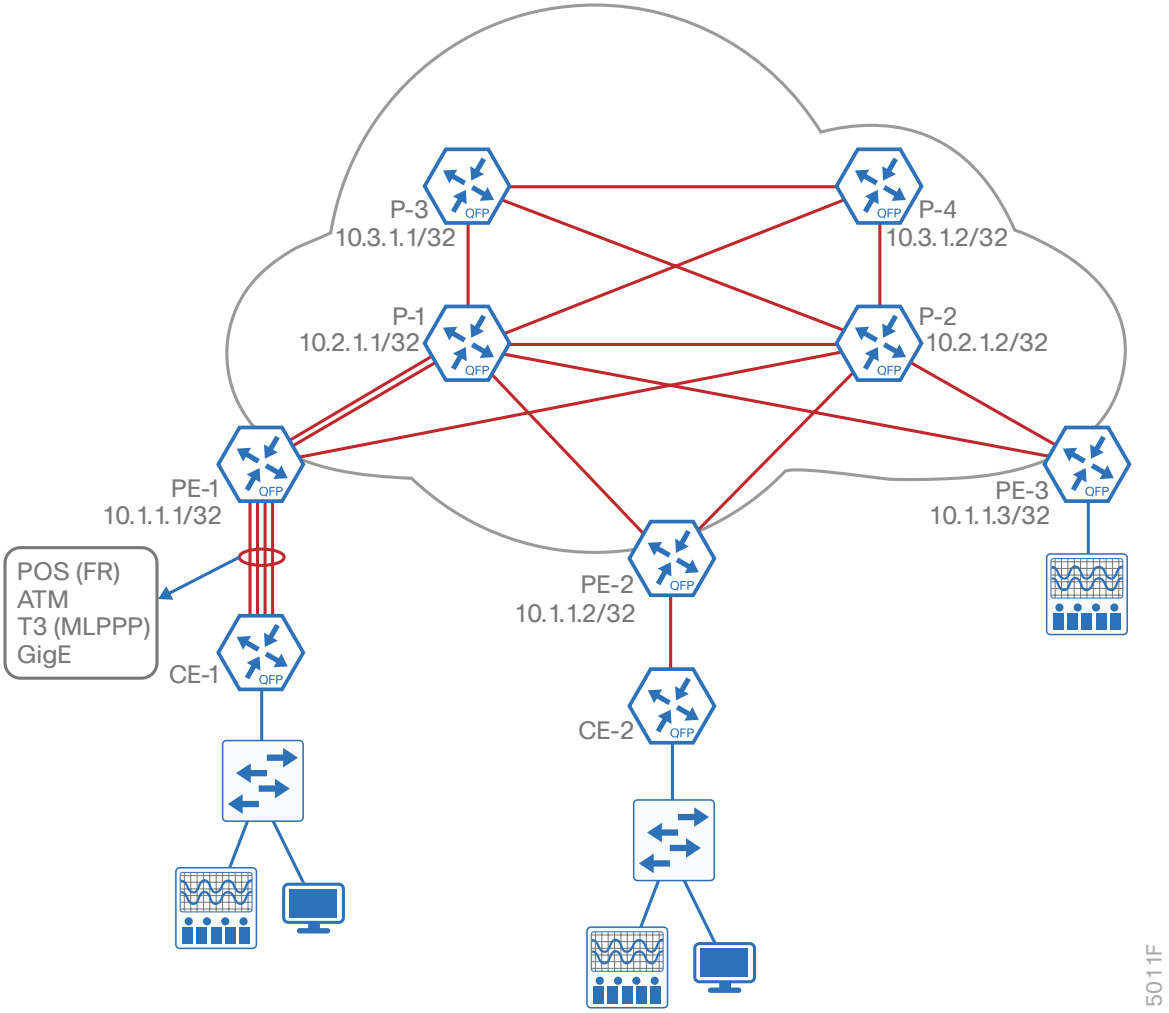

HARDWARE & FEATURE SPECIFICATIONS

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

Key Vertical Features

Table 2 defines the 3-D 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..

Disclaimer

Refer to appropriate CCO documentation for release/feature support across different platforms.

Table 2 3-D Feature Summary with Hardware and PIN

Deployment Layer (PIN)	Platforms	Critical vertical features
SP-Core	P1—ASR1004 (RP2/FP40) P2—ASR1004 (RP2/FP40) P3—CSR1000v (Route-Reflector) P4—ASR1002X (Route-Reflector)	MPLS LDP, LDP-Sync, LDP-Autoconfigure IGP: OSPF (Single-Area) EGP: BGP, iBGP, RR, BGP-Constrained Route distribution (BGP-RTC), Non-Stop Forwarding (NSF), Non-Stop Routing (NSR) All core facing interfaces: 10 Gb and 1 Gb Port channel: 10 Gb Member links RSVP-TE single hop/link redundancy NetFlow, EPC, IP-SLA, NTP, SNMP, Single and Multi-hop BFD
SP-EDGE	PE1—ASR1013 (RP2/ESP200, UUT1) PE2—ASR1006X (RP2/ESP100, UUT2) PE3—ASR1006 (RP2/ESP100, UUT3) CE1—ASR1006 (RP2/FP40) CE2—ASR1004 (RP2/FP40) CE-PE Links: Dot1Q Ethernet: SPA-5X1GE-V2 Multilink PPP: SPA-4XCT3/DS0 Frame-Relay: SPA-4XOC3-POS ATM PVC: SPA-3XOC3-ATM-V2	Dual-stack VRFs and MP-BGP on all PE's CE-PE Routing: EIGRP, OSPF, eBGP, BGP-PIC-EDGE, BGP Multipath redistribution, Single-hop BFD, ACLs, FNF, HQoS, IP-SLA Multicast-VPN's using SP-CORE native multicast (PIM-SSM) Multicast-LDP (MLDP in the CORE, PIM-SSM inside VRF) Unicast Reverse Path forwarding (uRPF)

Hardware Profile

Table 3 defines the set of relevant hardware, servers, test equipment and endpoints that are used to complete the end-to-end deployment.

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.

Table 3 *Hardware profile of servers and endpoints*

VM and HW	Software versions	Description
Ixia IxNetwork-FT IxExplorer	IxNetwork-FT–6.20.601.31, Protocols : 6.20.335.15 IxOS version–6.20.800.12 EA-SP1	Generate eBGP neighbors Generate dot1q connected interfaces (V4+V6) Simulate IGMP multicast clients Simulate Multicast sources Generate stateless traffic with QoS markings (Diff-Serv)

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 4 *MPLS-L3VPN Service-Provider Profile: feature scale*

Feature	Scale
L3VPN Scale	2500 Dot1Q VRFs 42 MLPPP VRFs 42 ATM PVC VRFs 42 Frame-Relay VRFs Each-VRF: 200 routes (V4), 10 routes (V6) Total L3VPN routes on the RR: approximately around 1 million routes
MLDP	500
FNF	100 CE-PE interfaces
BFD	Single-Hop: 2000 (v4), 2000(v6) Multi-Hop: 2 BFD template timers: bfd intervals 999 x 3 and BFD-multihop template 333 x 3, BFD-Singlehop-template 300 x 3
HQoS	3-Level HQoS applied to 2500 CE-PE interfaces
MLPPP	42 MLPPP interfaces with 4 member-links
Frame-Relay PVCs	42
ATM PVCs	42

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.

Images are loaded on the devices under test (UUT) via the tftp server using the Management interface.

To validate a new release, the network topology is upgraded with the new software image with existing configuration composed of the use cases and relevant traffic profiles. Addition of new use cases acquired from the field or customer deployments are added on top of the existing configuration.

With respect to longevity for this profile setup, CPU and memory use/leaks are monitored during the validation phase. Furthermore, to test the robustness of the software release and platform under test, negative events are triggered during the use case execution process.

USE CASES

Table 5 describes the use cases that were executed on the MPLS-L3VPN Service-Provider 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 CORE-Bringup, EDGE-Bringup, traffic engineering, route-reflector tests, multicast services, monitoring and troubleshooting, simplified management, and system health monitoring, along with system and network resiliency.

Table 5 *List of use case scenarios*

No.	Focus area	Use cases
CORE Bringup		
1	Configure and validate the SP-CORE features	Bringup the SP-CORE with all features and protocols required for enablement of services such as MPLS-L3VPN ▪ Configure OSPF single-area with key-chain authentication ▪ Configure MPLS with loopbacks, LDP autoconfig and LDP-IGP Sync, LDP authentication ▪ Configure NSF/NSR for OSPF, LDP and BGP ▪ Configure RR in the core and establish the iBGP neighborships with the PE routers ▪ Configure NTP in the core ▪ Configure single-hop BFD in the core ▪ Configure multi-hop BFD between the PE and the RR ▪ Enable all the core routers for Syslog, SNMP ▪ Enable native multicasting in the SP-Core using PIM-Sparse mode ▪ Enable RSVP-TE for single-hop/link backup in the core ▪ Configure and validate L3-port-channel with 10GigE member links between PE and P routers

Table 5 continued

Edge Bringup		
2	Configure and validate the SP-EDGE features	Bringup the SP-EDGE by configuring all the PE-Routers with necessary MP-BGP configurations required to support services like L3VPN/MLDP-MVPN, 6PE, etc. ▪ Configure dot1Q sub-interfaces, Frame-Relay Point-to-Point PVCs, ATM PVCs and Multilink-PPP bundles between CE-PE ▪ Configure and validate a combination of OSPF/EIGRP/eBGP over these links between CE-PE ▪ Configure dual-stack (V4+V6) on the CE-PE interfaces ▪ Configure and validate BGP-PIC-EDGE and BGP multipath redistribution ▪ Configure and validate route redistribution to and from OSPF/EIGRP into BGP ▪ Configure and validate the 6VPE feature on the PE routers ▪ Set up single-hop BFD towards edge and multi-hop BFD on the PE routers towards the RRs ▪ Configure and validate PE-to-CE features such as 3-Level HQoS, FNF, and input and output access-lists
Traffic engineering		
3	RSVP-Traffic engineering	Configure and Validate RSVP-Traffic engineering. ▪ Enable MPLS traffic-engineering auto-tunnels for single-hop backups in the core ▪ Validate the traffic through the tunnels.
Route-reflector tests		
4	RR-Bringup and validation	Configure and bringup RRs in a same cluster-id for redundancy purposes ▪ Validate that the RR is able to properly deflect the VPNv4 and VPNv6 routes between the PE routers in the core ▪ RR should be able to handle a minimum of 1 million V4 routes
5	RR with BGP-RTC	Configure constrained BGP Route import/export by configuring RTC address-family on the RR and PEs in the network ▪ Validate the RTC functionality by adding/removing RT-Filters into the VRF of PEs.

Table 5 continued

Multicast services		
6	Validate multicast VPN (MVPN) services	Configure and validate multicast services in the L3VPN VRFs ▪ Setup the SP-CORE to use the native multicasting to carry the customer VRFs multicast traffic. ▪ Validate PIM Sparse-mode with static and auto-rp and PIM-SSM.
7	Validate Multicast LDP (MLDP) based MVPN	In this use case, the SP-Core need not be configured with native-multicast. Instead, multicast-LDP is used to signal and build the distribution trees to carry the customer VRF's multicast.
Monitoring & troubleshooting		
8	Embedded packet capture (EPC)	Network admin should be able to troubleshoot the network by capturing and analyzing the traffic. ▪ Embedded packet capture ▪ Wireshark
9	Flexible NetFlow (FNF)	Enable IT admins to determine network resource use and capacity planning by monitoring IP traffic flows using Flexible NetFlow. ▪ Traffic types: IPv4, IPv6
10	Syslog, NTP, SNMP and Show commands	Enable IT admins to remotely monitor, troubleshoot and take corrective measures by enabling network nodes to send Syslog alarms, error messages, tracebacks, interface events, etc. ▪ Enable SNMP traps and execute an SNMP MIB-Walk from a remote SNMP-Server ▪ Configure and validate NTP on all the SP-CORE nodes, so that event-correlation becomes standardized with Synchronized time across the network ▪ Validate all the relevant feature and platform level show commands for appropriate and accurate field values
11	IPSLA	Network admin should be able to troubleshoot the network by enabling the IPSLA. ▪ Configure and validate IPSLA probes from PE to RR, CE to PE, CE to CE.

Table 5 continued

System health monitoring		
12	System health	Monitor system health for CPU use, memory consumption, and memory leaks during longevity. ▪ Capture the outputs on all the UUTs (PE routers)
System & network resiliency, robustness		
13	Network resiliency	Verify system-level resiliency during the following events on all the UUTs (PE routers) ▪ Active RP failure/RP switchover ▪ Active ESP failure/ESP switchover ▪ Port-channel member link flaps ▪ Interface flaps ▪ Flapping of member links in multilink-PPP bundles ▪ SIP/SPA reloads and OIR
14	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 ▪ Config/Unconfig VRFs ▪ Remove/Add BGP-address-families ▪ Clear IGP/BGP processes ▪ Clear counters ▪ IP-Fragmentation ▪ On the fly modification of HQoS ▪ Addition/deletion of new objects in OBACL

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.