Cisco Evolved Programmable Network Services Design Guide, Release 5.0

First Published: 2017-06-22

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

CHAPTER 1

Overview

This chapter contains the following sections:

Evolved Programmable Network Overview

The Cisco Evolved Programmable Network (EPN) system is built on the successful EPN architecture framework to bring greater programmability and automation. This effort is part of a multi-year ongoing development program that is built towards a flexible, programmable, and cost-optimized network infrastructure that Cisco targets to deliver in-demand fixed and mobile services.

The Cisco EPN system follows a layered design aimed to simplify the end-to-end transport and service architecture. By decoupling the transport and service infrastructure layers of the network, it allows these two distinct entities to be provisioned and managed independently.

This guide focuses on the design aspects of the service infrastructure layer, with focus on Metro Ethernet Forum (MEF) and business services.

The essential features for the service implementation are as follows:

- Hierarchical quality of service (H-QoS) policy.
- SLA enforcement.
- Performance management.
- High availability.
- Operations, Administration, and Maintenance (OAM) tools.

Evolved Programmable Network Framework

The Cisco EPN framework is designed such that the network services are decoupled from the transport infrastructure and the network layers are interconnected through open APIs as shown in the below figure.

Figure 1: Evolved Programmable Network Framework

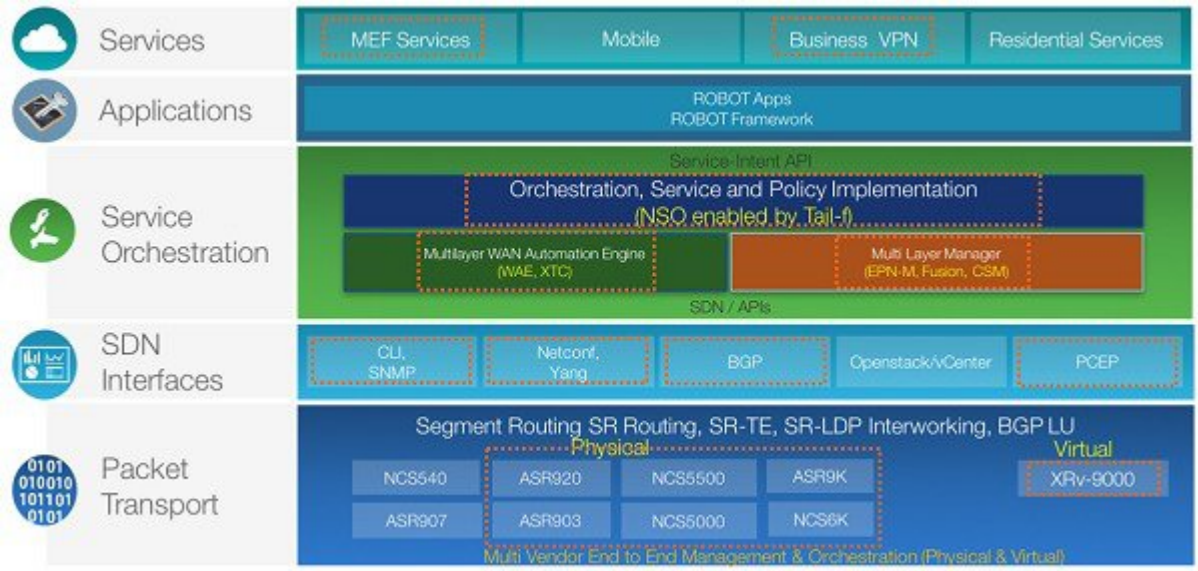

The goals of the Cisco EPN are:

- Increase the operation quality and capacity in service provider networks, by automating the instantiation of services.
- Increase the network utilization by intelligent engineering of traffic forwarding.

The Cisco EPN enables the activation of business and MEF services over a packet transport network that is based on segment routing. The standard based interfaces such as BGP link state (BGP-LS), Path Computation Element Protocol (PCEP), and Network Configuration Protocol (NETCONF) with Yet Another Next Generation (YANG) model, allow the elements in the service orchestration layer, to control all the aspects of network operations. The Cisco EPN also uses network programmability components such as Network Service Orchestrator (NSO), XR Transport Controller (XTC), Cisco Software Manager (CSM), and Evolved Programmable Network Manager (EPN-M), to further simplify the network operations in multi-vendor environment. For more details on the role of each controller and orchestrator, see Orchestration Layer Design and Device Management Tools, on page 29.

The Cisco EPN deployment models and services coverage are given in the below table. The implementation details of each deployment model can be found in the associated implementation guide.

Table 1: Deployment Models and Services Coverage

Use case	Transport Description	Access Origin Services	Aggregation Origin Services
1	Small Network - End to End Segment Routing	Label distribution Protocol (LDP) based Virtual Private Wire Service (VPWS)	NA
		Ethernet Virtual Private Network (EVPN) based VPWS	
		Static VPWS	
		LDP based Virtual Private LAN Service (VPLS)	
		BGP Auto-Discovery (BGP-AD) based VPLS	
		Layer-3 Virtual Private Network (L3VPN)	
2	Large Network - End to End Segment Routing	LDP based VPWS	LDP based Hierarchical VPLS (H-VPLS)
		EVPN based VPWS	BGP-AD based H-VPLS
		LDP based VPLS	H-VPLS using Provider Backbone Bridging Ethernet VPN (PBB-EVPN)
		BGP-AD based VPLS	
		L3VPN	
3	Large Network - End to End Programmable Segment Routing	LDP based VPWS	L3VPN with PW-HE
4	Large Network - Inter-AS End to End Segment Routing	LDP based VPWS	LDP based H-VPLS
		EVPN based VPWS.	BGP-AD based H-VPLS
		LDP based VPLS	H-VPLS using PBB-EVPN
		BGP based VPLS	
		L3VPN	

Use case	Transport Description	Access Origin Services	Aggregation Origin Services
5	Large Network - Segment Routing and LDP Interworking	LDP based VPWS	LDP based H-VPLS
		EVPN based VPWS	BGP-AD based H-VPLS.
		LDP based VPLS	H-VPLS using PBB-EVPN
		BGP-AD based VPLS	
		L3VPN	
6	Large Network - Layer2 Access to Segment Routing Transport		LDP based VPLS
			BGP based VPLS
			L3 VPN

CHAPTER 2

MEF Services

This chapter contains the following sections:

Design Overview

The Cisco EPN system supports the following MEF transport services on a single converged network:

- Point-to-Point E-Line services through Ethernet over Multiprotocol Label Switching (EoMPLS), pseudowire (PW), static pseudowire (PW), and EVPN.
- Multipoint E-LAN services through PBB-EVPN, VPLS or H-VPLS.

The Cisco EPN system also supports MEF service on following access networks:

- Segment routing/MPLS-based access network - To deploy a converged architecture to transport various types of services with a uniform control plane.
- Ethernet access network - To cap the investments in legacy network deployment and to facilitate the migration to a packet-switched layer 3 network architecture.

  EPN 5.0 supports only two types of ethernet access network, one based on the G.8032-enabled rings and the other on Resilient Ethernet Protocol (REP)-enabled rings for SP Carrier Ethernet Architecture.

The service transport reachability to remote nodes can be achieved in two ways:

- Unified MPLS approach – The service edge nodes, for example, access PE enable the MEF services by learning each other's loopbacks through BGP labeled-unicast (BGP-LU). The advertisement of end-to-end routes through BGP-LU is extended to the access network, by using the inter-domain nodes as inline route reflector.
- Path Computation Engine (PCE) - The PCE node computes and provides a Segment Routing Traffic Engineering (SR-TE) end-to-end path. The PCE function is a part of the EPN orchestration layer, and implemented by the XR Transport Controller (XTC). This scenario is covered only for EPN system business services from implementation perspective but the concept is equally applicable to MEF services.

Note For EPN 5.0 release, the above case is only valid for business services.

Note The BGP-LU is not required in PCE case. The XTC has the full visibility of the network topology and it is also capable of steering the traffic based on the user's constraints such as shortest path (IGP metrics), traffic-engineering metrics, and disjoint path. Therefore, there is no need for Service Edge (SE) nodes to learn the loopbacks of other domains and to store them into their forwarding table ahead of time.

The Cisco EPN uses the Unified MPLS method for MEF services transport reachability. Once the transport is established, the services are enabled. Due to the transport mechanism, the point-to-point E-Line service is supported end-to-end between service edges without requiring Multi-Segment PseudoWire (MS-PW) implementation on the intermediary nodes. The service edge node can be access provider edge (PE) or pre-aggregation node/aggregation node depending on the customer requirements as shown in below figure.

The Cisco EPN system design covers following types of E-Line services:

1. Traditional VPWS using LDP Signaling
2. EVPN based VPWS using BGP control plane
3. Static Pseudowire

Figure 2: Supported E-Line Services

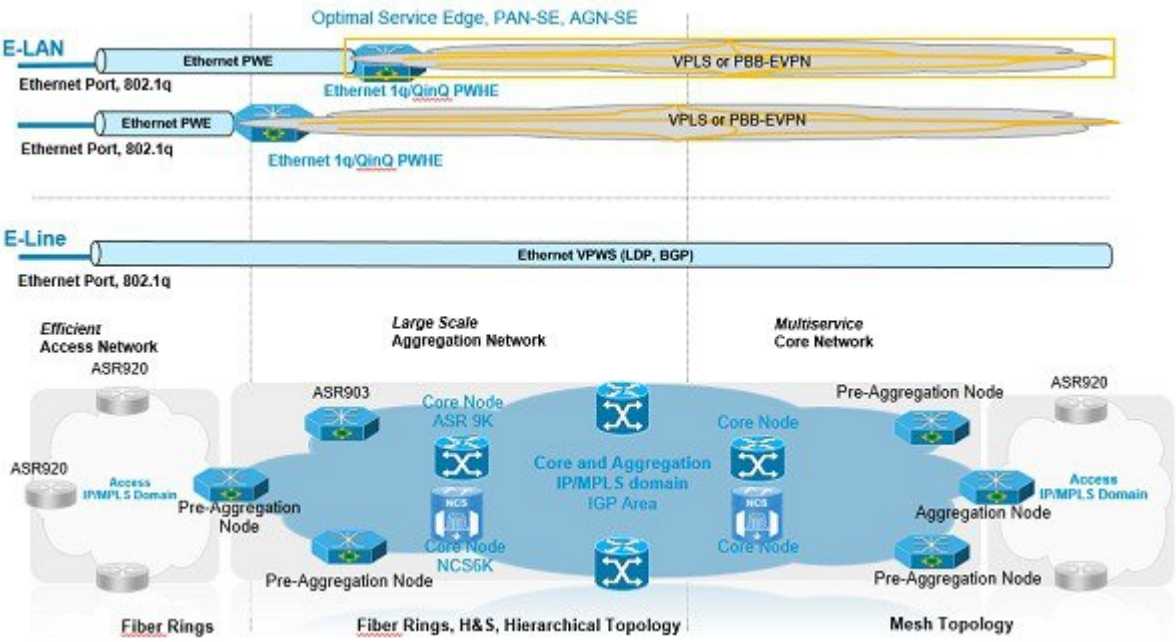

The Multipoint L2VPN Ethernet LAN (E-LAN) services are also part of Cisco EPN system as shown in the below figure. The Cisco EPN system supports the two mechanisms for providing L2VPN services at the

service edge: flat multipoint full mesh VPLS and traditional H-VPLS virtual forwarding instance (VFI) or PBB-EVPN.

The VPLS VFI automatically creates a full mesh of PWs to transport the L2VPN services between the service edge (SE) nodes. The VPLS VFI is configured on the SE nodes. The SE for multipoint services could be the access node (AN), pre-aggregation node (PAN), or aggregation (AGG) node.

The E-LAN services are also implemented in a hierarchical model, using the pseudowire transport from the access node to the pre-aggregation node or aggregation node that provides the service edge functionality:

- VPLS VFI automatically creates a full mesh of PWs to transport L2VPN services between service edge (SE) nodes. The service edge (SE) could be access node, pre-aggregation or aggregation node but usually VPLS VFI is configured at the SE node in pre-aggregation or aggregation node. The VPLS VFI is configured at the SE node rather than at the access nodes to minimize the number of neighbors involved in the VPLS VFI and to avoid any potential MAC address scaling issues on the access node. The Ethernet Pseudowire is responsible for transporting the L2VPN service from the access node to the SE in the pre-aggregation node or aggregation node. The Pseudowire from the AN is then terminated at the SE node on the VPLS VFI that provides the L2VPN service.
- PBB-EVPN - To replace the full mesh of VPLS core PWs between the SE nodes. The EVPN uses the BGP control plane at the pre-aggregation node or aggregation node, to distribute the MAC address reachability information over the MPLS core and aggregation networks, while bringing the sophistication of policy control available with BGP-based L3VPNs into MAC address learning. The PBB aspect of PBB-EVPN controls the MAC address distribution. In contrast to E-VPN, the PBB-EVPN uses BGP to advertise only the SE nodes Backbone MAC (B-MAC) address reachability. The remote Customer MAC (C-MAC) to B-MAC binding still uses data plane learning. As a result, the number of MAC addresses learned by BGP speakers in the EPN core and aggregation domains is reduced to the number of SE nodes, which are usually in the hundreds (far fewer than the millions of C-MAC addresses typically found in the large SP networks).

Similar to E-Line services, the E-LAN multipoint L2VPN service provisioning occurs on the SE nodes and route reflectors without requiring the configuration on the intermediary nodes between the SE nodes.

Figure 3: Supported E-LAN services

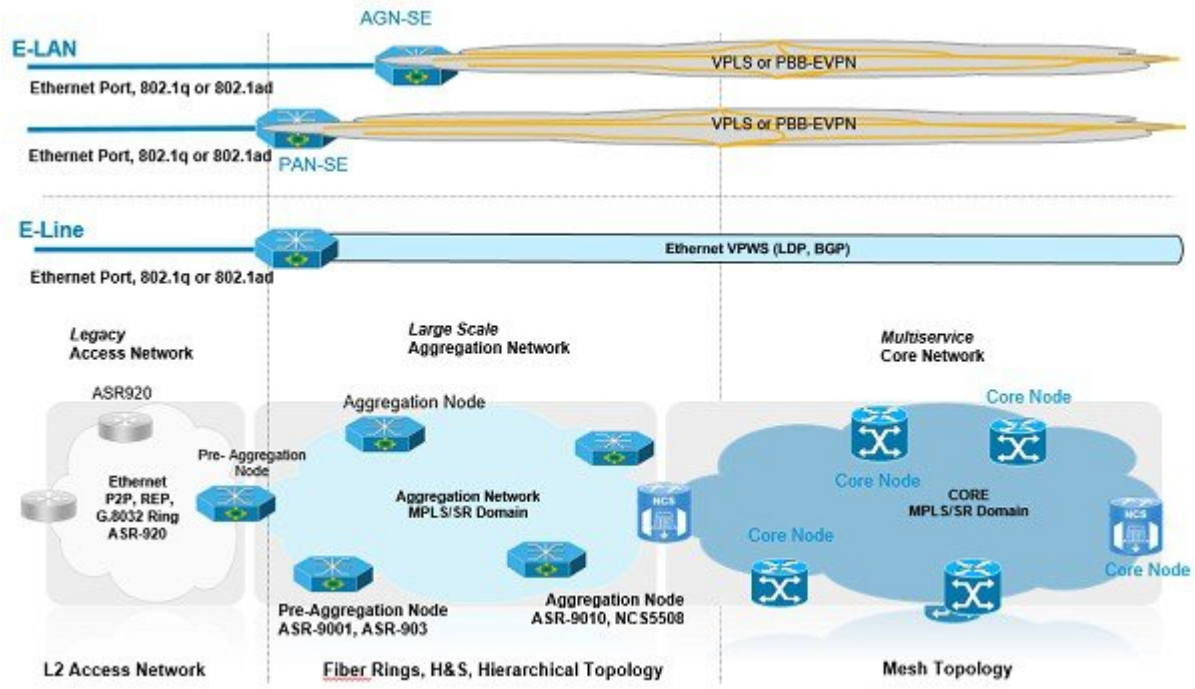

The E-Line (Point-to-Point) and E-LAN (Multipoint) MEF services use ingress filtering techniques, to reduce the number of routes in the access domain as shown in the below figure.

Figure 4: E-Line Service Example for Route Scale Ingress Filtering

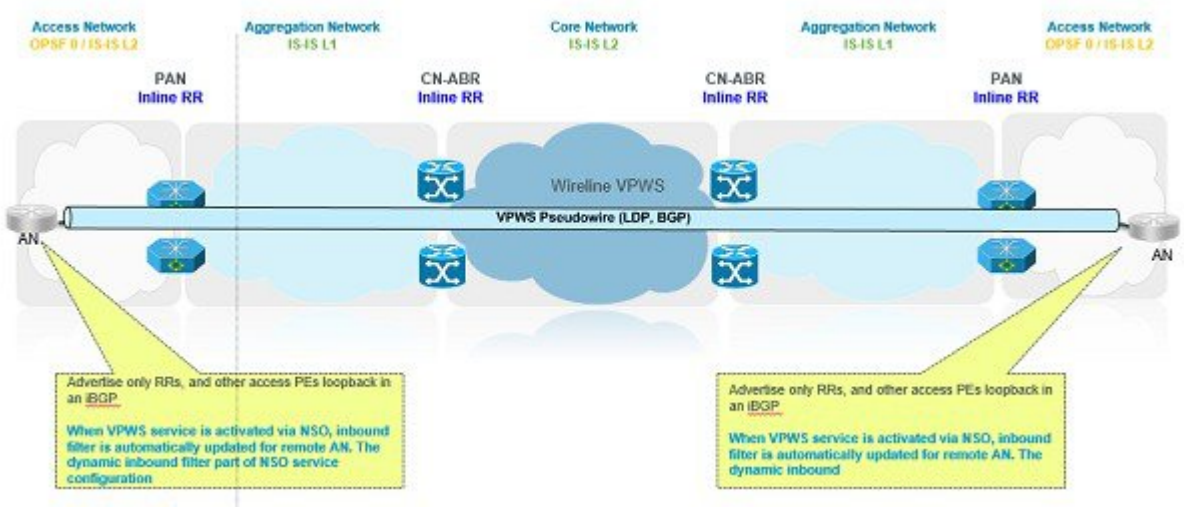

For IPv4 route scale, a dynamic IP prefix list is used for inbound filtering. The AN-SE transport configuration permits only the loopbacks of the service route reflector and other access PE. When a wireline service to a new destination is activated, the route-map used for inbound filtering needs to be updated. In case of the services provisioned through Cisco NSO, the inbound filters get dynamically updated as part of the orchestrated service provisioning phase. However, in case of services provisioned through CLI, the prefix list needs to be updated manually.

Note In case of service provisioning through CLI, the prefix list can also be updated automatically using Embedded Event Manager (EEM) script.

The customer premises equipment (CPE) devices for E-Line and E-LAN services can be connected through an Ethernet port UNI, 802.1Q-tagged UNI, or 802.1ad double-tagged UNI directly to the access nodes. The access node may translate the VLAN tag of the customer UNI to a unique VLAN tag on the SP network, or may push a Service VLAN (S-VLAN) tag on the Customer VLAN (C-VLAN) to create a double-tagged (Q-in-Q or 802.1ad) NNI. For single or double-tagged NNI, the Ethernet NNI is terminated on the service edge node. The SE node maps the VLAN to a pseudowire to be transported across the aggregation domains and core network. Per-subscriber H-QoS and required subscriber ACLs are applied to the Ethernet NNI interface of the SE node.

Note In EPN system implementation, the subscriber CPE device is connected to the service provider (SP) network at the service edge through 802.1Q-tagged UNI for both E-Line and E-LAN services.

MEF Transport Service Architecture Design

E-Line Service - Traditional L2VPN VPWS

This section describes how MEF point-to-point wireline services are deployed using a PW-based model. The Customer Premise Equipment (CPE) is connected to the service edge node through ANs located in an urban area with fiber access network.

In the figure shown below, the AN is connected to a SE node through an Ethernet NNI.

In the case of Unified MPLS access, the PW is initiated to transport the service across the core network and to the remote SE node.

The figure below shows the traditional VPWS with an Ethernet Private Line (EPL) or Ethernet Virtual Private Line (EVPL) service enabled by using a PW between the service edge access nodes in a local and remote access network across the core network.

For EPL services, the AN ensures that customer BPDUs are forwarded end-to-end over the PW. For information on the implementation of Network Service Orchestrator (NSO) service Yet Another Next Generation (YANG) model, see *Cisco EPN Service Orchestration User Guide, Release 5.0*.

Figure 5: Traditional e-line L2VPN VPWS

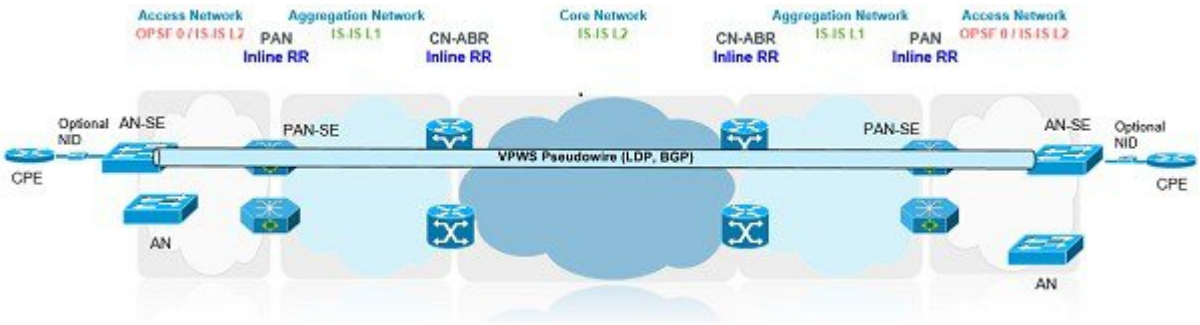

The ANs enable the VPWS to learn each other's loopbacks through BGP labeled-unicast that is extended to the access network by using the PANs as inline route reflector. For more information, see *Cisco EPN Transport Design Guide, Release 5.0*. The ingress filtering of BGP routes at the ANs, keep the route scale in the access network to a minimum.

The figure below shows a variation of VPWS deployment with native Ethernet (in G.8032/REP-enabled rings or hub-and-spoke topologies) access where the service utilizes a PW between the PAN-SEs.

Note Although, there are multiple access options shown in the figure below from design perspective, the EPN 5.0 system implementation guide only covers Resilient Ethernet Protocol (REP) and G.8032.

Figure 6: E-line Service between SEs with Non-Unified MPLS Access

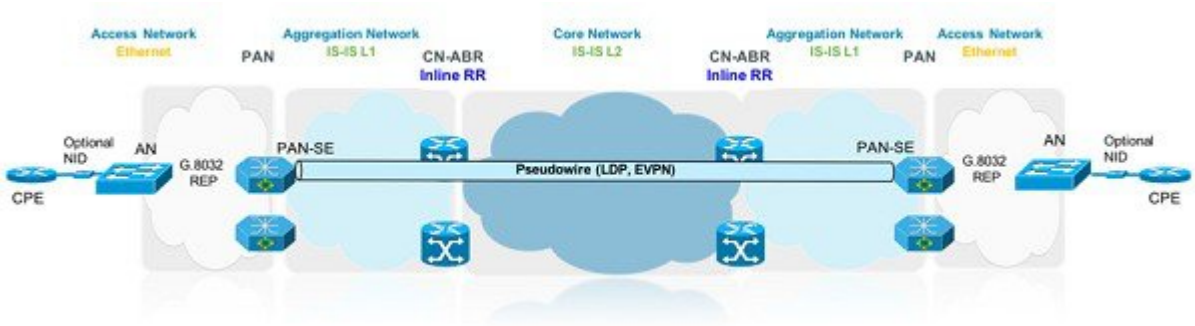

For EPL services design perspective, the access node ensures customer BPDUs are carried transparently across the shared Ethernet access network (G.8032-enabled rings and REP topologies) by implementing L2TP. The tunneling of L2 control protocol frames is done selectively depending on protocol and customer preference. The supported protocols include Spanning Tree Protocol (STP), Link Aggregation Control Protocol (LACP), Link Layer Discovery Protocol (LLDP), Cisco Discovery Protocol (CDP), VLAN Trunking Protocol (VTP), and Dynamic Trunking Protocol (DTP).

Note For information on the L2TP implementation, see *EPN 4.0 Design and Implementation Guide*.

E-Line Service – EVPN based Virtual Private Wire Service

This section describes the MEF point-to-point wireline service. The Cisco EPN uses EVPN as the next generation technology umbrella for the deployment of Ethernet services. The EVPN uses a BGP control plane for Ethernet segment and MAC distribution and learning over MPLS core. For E-LINE services, the benefits of using EVPN based VPWS over traditional LDP based VPWS are shown in the below table.

Table 2: E-Line – Additional Service Capabilities Using EVPN based VPWS

Service	Additional Capabilities
E-Line	• All-active & single-active multi-homing support. • Both single-segment and multi-segment support. • Discovery and signaling through single protocol, for example, BGP.

This transition in core technology from LDP to BGP for MEF services does not affect the connectivity model between the CPE and the access node.

In the figure shown below, the access node is connected to a SE node through an Ethernet NN. In case of Unified MPLS access, the access node can directly initiate the PW to transport the service across the core network and to the remote SE node.

The service is enabled across the core network by using a PW between service edge access nodes.

For information on the implementation of the EVPN based VPWS using NSO YANG model, see *Cisco EPN Service Orchestration User Guide, Release 5.0*.

Figure 7: EVPN based VPWS Operations

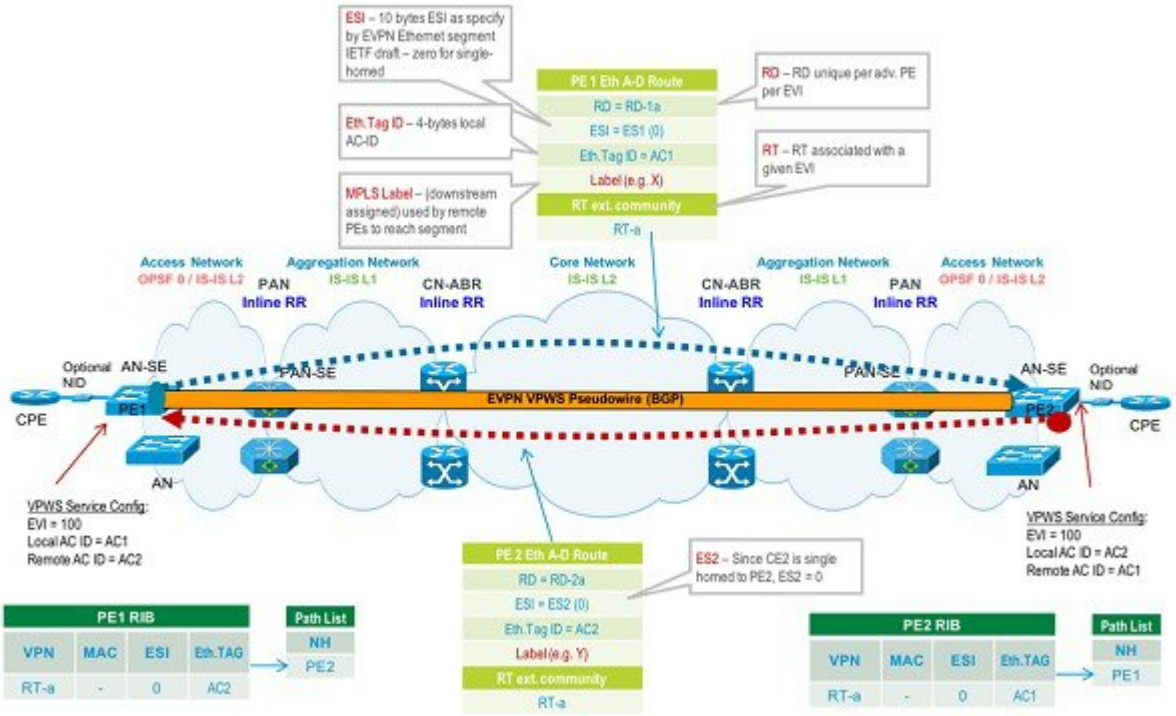

E-Line Service with Static Pseudowire

The E-Line service based on static pseudowire provides a migration path to the operators transitioned to segment routing at the transport layer. The migration path is provided to eliminate the LDP protocol from the service signaling layer without enabling BGP based EVPNs.

The elimination of LDP and associated TCP sessions, enables to provide:

- Massively scalable solution for Hierarchical L2VPN service while allowing the access aggregation device to aggregate large number of PWs.

  Note The static PW with SR Anycast Segment Identifier (SID) provides the Active-Active Point to Multipoint service, in case of failure. This helps to simplify High Availability (HA) configuration in the whole network.

- Adverse effect of inhibiting the end-to-end signaling of the attachment circuit status. This precludes remote port shutdown behavior upon local UNI failure. However, this limitation can be addressed by static Operations, Administration, and Maintenance (OAM). For details on OAM, see OAM, on page 25 section.

E-LAN Services with VPLS Core

This section explains the deployment of MEF multipoint services using a hierarchical VPLS model. The aggregation service edge or pre-aggregation service edge node implements the VPLS VFI for the E-LAN service. The business customer's CPE equipment is connected either directly to this service edge node or through access node located in an urban area with fiber access between two business access nodes. The access node is connected to the SE node through an Ethernet NNI or spoke PWs, that transport the service between the SE nodes and any MPLS-enabled access node subtending from the SE node.

E-LAN Services with VPLS Core and Unified MPLS Access (H-VPLS)

An EoMPLS PW is used between each pre-aggregation node with service edge (PAN-SE) and the respective access node. The figure below shows a wireline VPLS service. The VPLS service can be an Ethernet Private LAN (EPLAN) or Ethernet Virtual Private LAN (EVPLAN) business service that is enabled by using a VPLS VFI between two PAN-SEs across the core network.

Note The EVPLAN implementation is considered as an example in this section and can be used to deploy the EP-LAN (multipoint service at port level).

Figure 8: Wireline VPLS Services with Unified MPLS Access

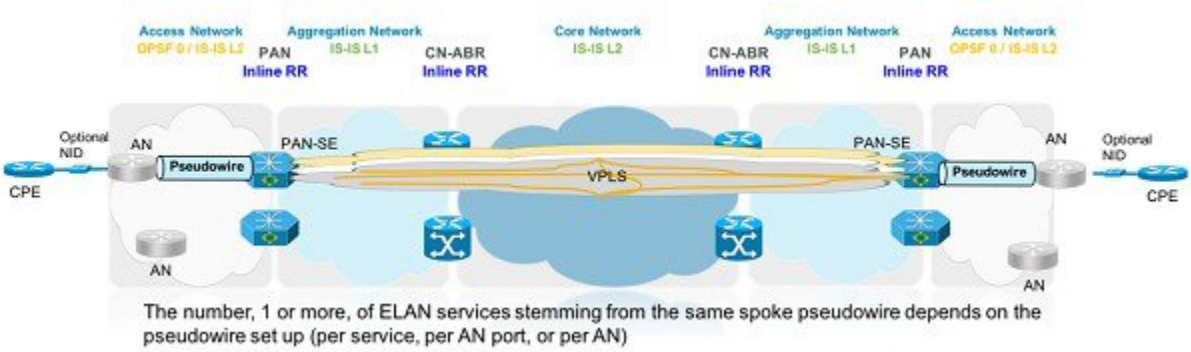

You can set up the spoke PW per access node, per access node port, or per service, depending on the platform support and operator preferences.

- To set spoke PW per service, the PW must be unique per service and transport the individual E-LAN (EVPLAN) services that originate from the UNI port. The pre-aggregation node with service edge simply maps the access PW) to the VPLS VFI.
- To set spoke PW per AN, a single PW must connect the AN to the SE device and transport all the E-LAN (EPLAN, EVPLAN) services that originate from that node.
- To set spoke PW per AN port, the PW must be unique per AN port and transport all the E-LAN (EVP-LAN) services that originate from that UNI port.

Within the spoke PW, different services are identified through a unique service-provider VLAN (S-VLAN) which is imposed on the access node. At the SE node, the spoke PW terminates on the PWHE main interface,

while services carried within the PW are de-multiplexed on individual sub-interfaces using the S-VLAN information. The sub-interfaces dispose the S-VLAN tag and act as the access circuits for the VPLS core transport service. For more details, see *EPN 4.0 MEF Transport Design and Implementation Guide.*

The setup of end-to-end connectivity through the core requires the following:

- PAN-SEs or AGG-SEs use inbound route filtering to limit the number of routes they learn.
- All the PAN-SEs or AGG-SEs must import and export the respective route targets associated with the service, in case of BGP-AD.

The SE node implements service SLA enforcement through per-subscriber Quality of Service (QoS) policies and all required ACLs. The AN provides the aggregate class enforcement through QoS.

Note For EP-LAN services, the access node needs to ensure that customer's L2 control protocols BPDUs are tunneled inside the spoke PW and carried through the VPLS core to the remote customer site. For more details on implementation, see *EPN 4.0 MEF Transport Design and Implementation Guide.*

E-LAN Services with VPLS Core and Ethernet Access

The figure shows a wireline VPLS service such as an EP-LAN or EVP-LAN business service. The E-LAN service is enabled by using a VPLS VFI between two PAN-SEs across the core network with an 802.1q, or a tagged (Q-in-Q or 802.1ad) Ethernet NNI between each PAN-SE and the respective AN.

Figure 9: VPLS Core and Legacy Ethernet Access

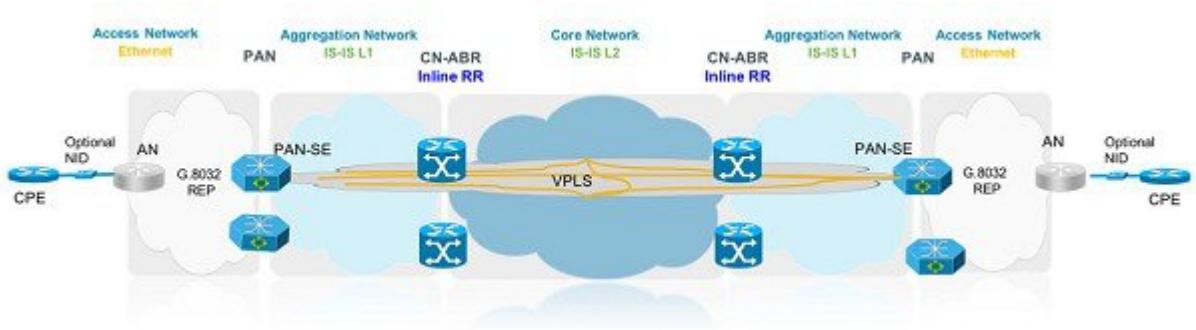

The AN is responsible for mapping the fixed access UNI to the proper 802.1q Q-in-Q or 802.1ad Ethernet NNI for transport to the PAN-SE. The PAN-SE then maps the S-VLAN and/or C-VLAN from the UNI or Ethernet NNI, to the VPLS VFI.

To enhance MAC address scalability, the SE nodes may implement 802.1ah encapsulation of C-MAC addresses in a single B-MAC address.

The transport to the PAN-SE is accomplished over a G.8032-enabled ring or a REP rings of AN nodes. The AN ensures that customer BPDUs are carried transparently across the shared Ethernet access network by implementing L2TP. The tunneling of L2 control protocol frames is done selectively, based on the protocol and customer preference. The supported protocols include STP, LACP, LLDP, CDP, VTP), and Dynamic Trunking Protocol (DTP).

The setup of end-to-end connectivity through the core requires the following:

- PAN-SEs or AGG-SEs use inbound route filtering to limit the number of routes they learn..
- All the PAN-SEs or AGG-SEs must import and export the respective route targets associated with the service, in case of BGP-AD.

The SE node implements service SLA enforcement through per-subscriber QoS policies and all required ACLs. The AN provides aggregate class enforcement through QoS.

E-LAN Services with PBB-EVPN Core

This section explains the deployment of MEF multipoint services using an EVPN with PBB model. The PBB-EVPN brings enhanced functionality, scalability, flexibility, and greater operational simplification to L2VPN services versus the VPLS VFI model.

In this model, the AGN-SE or PAN-SE node implements the E-LAN service through EVPN Instance (EVI), which is then routed between PE nodes by utilizing the address-family L2VPN EVPN construct in BGP. The CPE equipment is either connected directly to the SE node or through an AN located in an urban area. The AN may be connected to the SE node either through an Ethernet NNI, an Ethernet ring network, or an Unified MPLS access network.

E-LAN Services with IP/Unified MPLS Access

The figure given below shows a wireline L2VPN service. The L2VPN service can be an EP-LAN or EVP-LAN business service enabled by using an EVI with PBB configured between two PAN-SEs across the core network. The EoMPLS PW is used to transport traffic from the CPE UNI to the PAN-SE.

Figure 10: PBB-EVPN Core and Unified MPLS Access

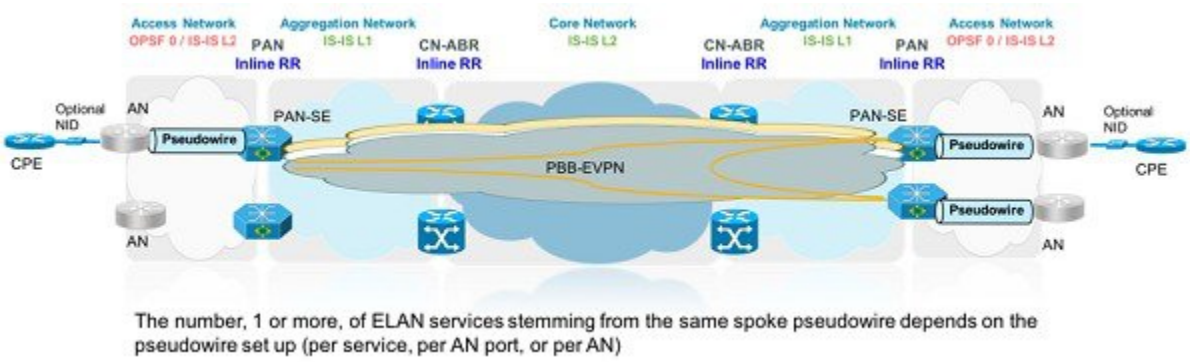

The access PW can be set up per access node, per access node port, or per service, based on the platform support and operator preferences.

Note The current release of the EPN System only validates E-LAN services with per-service access PW.

The PAN-SE packages the access PW into a bridge domain. This is referred to as the PBB-Edge bridge domain (BD).

With the PBB functionality of the SE node, the PBB-Edge BD is connected to a another PBB-Core BD on which the EVPN service is configured with an EVI. This connection is done to ensure that all C-MAC addresses are encapsulated in a B-MAC address and only B-MAC information is shared between the SE nodes. The EVI is configured on all SE nodes participating in this PBB-EVPN. The service traffic between the SE nodes is routed through BGP "address-family l2vpn evpn" information.

To ensure end-to-end connectivity through the core, the PAN-SEs or AGG-SEs use inbound filtering to limit the route at service edge.

To achieve the multipoint type of connectivity in case of BGP-AD, the E-LAN service requires all the PAN-SEs or AGG-SEs to export and import the respective route targets.

The SE nodes implement the service SLA enforcement through per-subscriber QoS policies and required ACLs, and learn the MAC addresses in the service. The ANl provides aggregate class enforcement through QoS

Note For EPLAN services, the access node needs to ensure that customer's L2 control protocol BPDUs are tunneled inside the spoke PW and carried through the VPLS core to the remote customer site. For more details on design implementation, see *EPN 4.0 MEF Transport Design and Implementation Guide.*

CHAPTER 3

Business Services

This chapter contains the following sections:

Design Overview

- L3 VPN services through Legacy Ethernet access, including G.8032 and REP-enabled ring topologies.
- L3 VPN services over segment routing and MPLS PW, for example, EoMPLS PW with Pseudowire Headend (PWHE) connectivity to MPLS VPN Virtual Routing and Forwarding (VRF) at the service edge (SE) node.
- L3 VPN services with programmable Segment Routing On-Demand Next-Hop (SR-ODN).

The Cisco EPN system also supports business services on the following access networks:

- Segmenting Routing and LDP based MPLS access network – To deploy a converged architecture to transport all service types with a uniform control and forwarding plane.
- Layer 2 access network – To cap investments in legacy network deployments and facilitate migration to a packet-switched network architecture. Ethernet access networks are supported in G.8032-enabled rings and REP enabled rings for SP Carrier Ethernet Architecture.

The business service transport reachability to remote nodes can be achieved in two ways:

- By using the Unified MPLS approach – The service edge nodes, for example, access PE nodes enable business services by learning each other's loopbacks through BGP-LU. The advertisement of end-to-end routes through BGP-LU is extended to the access network by using the inter-domain nodes as inline RR.
- By using a PCE – The PCE node computes and provides a SR-TE end-to-end path. The PCE function is part of the EPN orchestration Layer , and implemented by the XTC.

Note BGP-LU is not required for PCE. The XTC having the full visibility of the network topology is capable of steering the traffic, based on the user's constraints. Therefore, there is no need for the service edge nodes to learn the loopbacks of other domains and to store them into their forwarding table ahead of time.

Once the transport reachability is achieved, the services are enabled.

In case of Unified MPLS access, the scalability characteristics of a business service L3VPN can vary dramatically from one use case to another, including the number of VRFs involved and number of prefixes in each VRF. The service edge function implementing the L3VPN service can be placed at the access node AN or at the pre-aggregation or aggregation ABRs.

If the SE is not placed at the access node due to scalability reasons, the Ethernet PW enables the transport of the L3VPN service from the access node to the pre-aggregation or aggregation ABR. The service is then enabled at the PW termination point or Pseudowire Headend. In case of PWHE, the L2 and L3 transport converge at the service edge node.

In case of L3VPN SR-ODN scenario, the XTC acts as a PCE to provide the end-to-end path between aggregation service edges. The SR ODN enables faster L3VPN service and brings programmability aspect to the network, by automating the endpoint path calculation and programming the SRTE policy.

Note The current release of EPN system only supports SR-ODN at the aggregation node.

The business services implement inbound filtering for route-scale at the access node.

For a unified MPLS access, the CPE device is connected to the service provider network through an Ethernet 802.1Q-tagged UNI at the access node.

Figure 11: Business Services

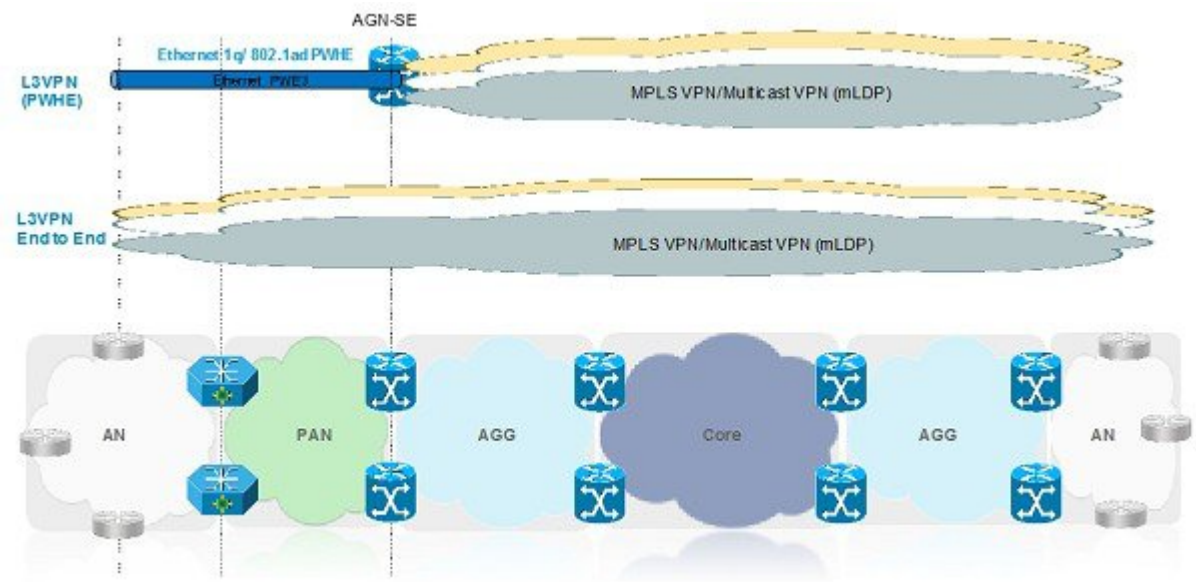

For the L3VPN in Layer 2 Native Ethernet Access scenario, the CPE device is connected to the SP network through an Ethernet 802.1Q-tagged UNI on the access node. The transport of the L3VPN service from the

access node to the service edge in the pre-aggregation node or aggregation node, is accomplished through direct Ethernet connections over a G.8032 or REP enabled Ethernet ring of access nodes.

The access node either translates the VLAN tag of the customer UNI to a unique VLAN tag on the SP network, or pushes a S-VLAN tag on the C-VLAN, to create a QinQ NNI. The Ethernet NNI is terminated on the SE node regardless of VLAN tagging configuration. The VLANs carrying the L3VPN service are mapped to an MPLS VPN VRF, which is then transported over the Unified MPLS transport network. The H-QoS and required subscriber ACLs are applied to the Ethernet NNI interface.

Note The current release of EPN system only supports single-tagged VLAN scenario.

Business Service Architecture Design

This section describes the transport and service edge functions required for the deployment for business L3VPN wireline services. The MPLS service virtualization provides an abstraction to deploy different types of services over a shared transport.

The L3VPN service can be implemented at the AN-SE, AGN-SE, or PAN-SE node. Depending on the node that implements the SE function, the CPE equipment is connected to the SE in one of the following ways:

- Directly, when the SE role is implemented at the AN.
- Through the AN subtending from the SE, when the SE role is implemented at the PAN or AGN ABR.

Wireline L3VPN Service with Unified Multiprotocol Label Switching Access

In this scenario, a PW transports the L3VPN service between the MPLS-enabled AN and SE node. The PW can be set up per access node, per access node port, or per service, based on the platform support and operator preferences. The setup of PW can be done as given below:

- To set PW per service, the PW must be unique per service and carry individual services that originate from the UNI port.
- To set PW per access node, a single PW is connected between the access node and SE device to carry all the services that originate from that node. Within the multiplexed PW, different services are identified through a unique S-VLAN which is imposed on the access node UNIs.

- To set PW per access node port, the PW must be unique per AN port to carry all the E-LAN and E-Tree services that originate from that UNI port.

Figure 12: Wireline L3VPN Service with Unified MPLS Access

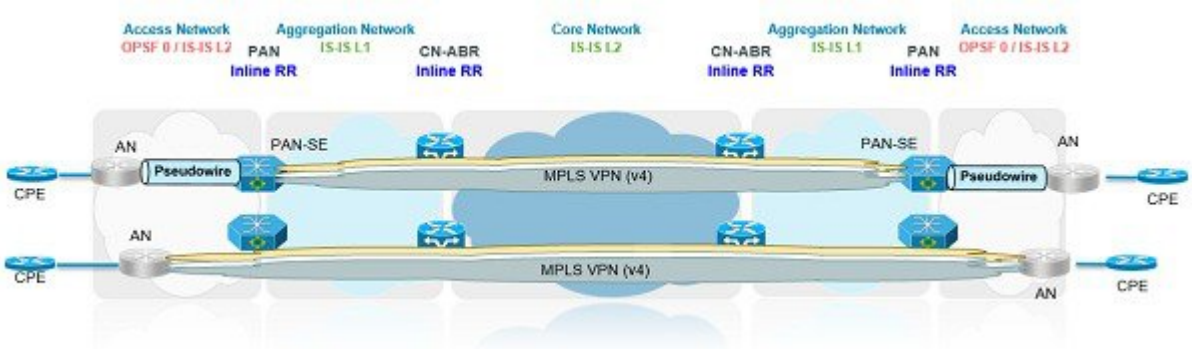

Note The current release of the EPN System only validates L3VPN services with per-service access PW.

At the SE node, a multiplexed PW terminates on the PWHE main interface, while the services carried within the PW are demultiplexed on the individual subinterfaces using the S-VLAN or C-VLAN information. The subinterfaces dispose the VLAN tags and become the service UNI for mapping of the L3VPN service directly to an MPLS VPN VRF. The H-QoS and security access lists are applied on these subinterfaces.

To ensure the end-to-end connectivity through the core:

- PAN-SEs or AGG-SEs implement the PWHE interfaces, to terminate the PW from the access node and map it to the MPLS VRF for the L3VPN service.
- Nodes with fixed access UNIs, transport all the services to the PAN-SE or AGG-SE.

The PAN-SE or AGG-SE implements the SLA enforcement through per-subscriber QoS policies and required ACLs. The access node provides the aggregate class enforcement through QoS.

Wireline L3VPN Service with Ethernet Access

The transport of the L3VPN service from the access node to the PAN-SE is accomplished over direct connection in a hub-and-spoke topology, G.8032-enabled or REP ring of access nodes, or over network virtualization (nV) based architecture. In the nV scenario, the access nodes act as satellites of the PAN-SE. This in turn takes the role of the nV host and becomes the central configuration and management point for all the remote devices.

Note The current release of EPN system only validates the ring topology.

Figure 13: Wireline L3VPN Service with Ethernet Access

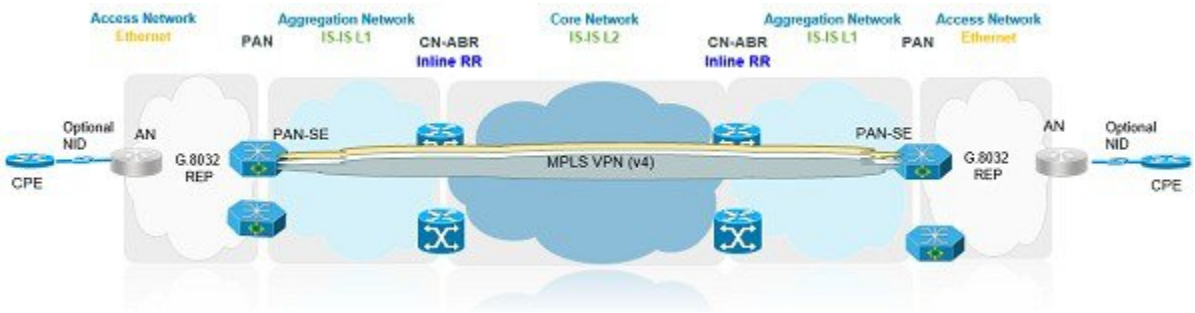

To ensure the end-to-end connectivity through the core:

- Nodes with fixed access UNIs map the UNI to the proper 802.1Q or Q-in-Q Ethernet NNI, for transport to the PAN-SE.
- PAN-SE maps the S-VLAN or C-VLAN from the UNI or Ethernet NNI, to the MPLS VRF.

The SE node implements the service SLA enforcement through per-subscriber QoS policies and required ACLs. The access node provides the aggregate class enforcement through QoS.

L3VPN Service with MPLS Access—On-Demand Next-Hop

This model uses PCE or XTC to provide end-to-end data path provisioning between service edge nodes that are part of the same L3VPN service. For information on XTC and data plane programming, see *Cisco EPN Transport Design Guide, Release 5.0*.

The SE nodes learn about the data path for the end-to-end service, by sending a ODN request to the XTC. The request is initiated upon receipt of service specific routes and BGP network layer reachability information (NLRI) from other service endpoints. The request is then relayed to the local SE by the network service route reflector with next-hop unchanged.

To resolve the network layer reachability requests from the SE nodes, the XTC learns about the full network topology by exchanging BGP-LS advertisements with the various inter-domain ABRs. The BGP-LS advertisements convey the topology information for each domain and other additional information required for XTC , based on one of the following constraints:

- Shortest path (IGP metrics)
- Traffic-Engineering metrics
- Disjoint path (link/node/SRLG)

To optimize the number of endpoints controlled by the XTC device, the service edge are enabled at the aggregation ABRs nodes for providing the L3VPN service.

The access network connected to the ABRs are deployed either using Ethernet or unified MPLS access. The L3VPN traffic is then carried from the access node to the SE through:

- Q-in-Q or 802.1ad transport for Ethernet-based deployment.
- Access PW for Unified MPLS access based deployment.

The figures given below provide the details about the path established through ODN with MPLS access. The BGP on SE node, receives the VPN prefixes from service route reflector. The SE node requests and receives the path information from XTC. The path information is in the form of a segment routing prefix-SID list. The SR Traffic Engineering on SE node installs prefix-SID list and allocates Binding Label or Binding-SID. The

binding label is installed to FIB and BGP table, to resolve remote SE reachability and VPN prefix reachability respectively. The BGP sets the best-path to FIB, by creating a recursive route.

Figure 14: Large Network—End-to-End Programmable Segment Routing

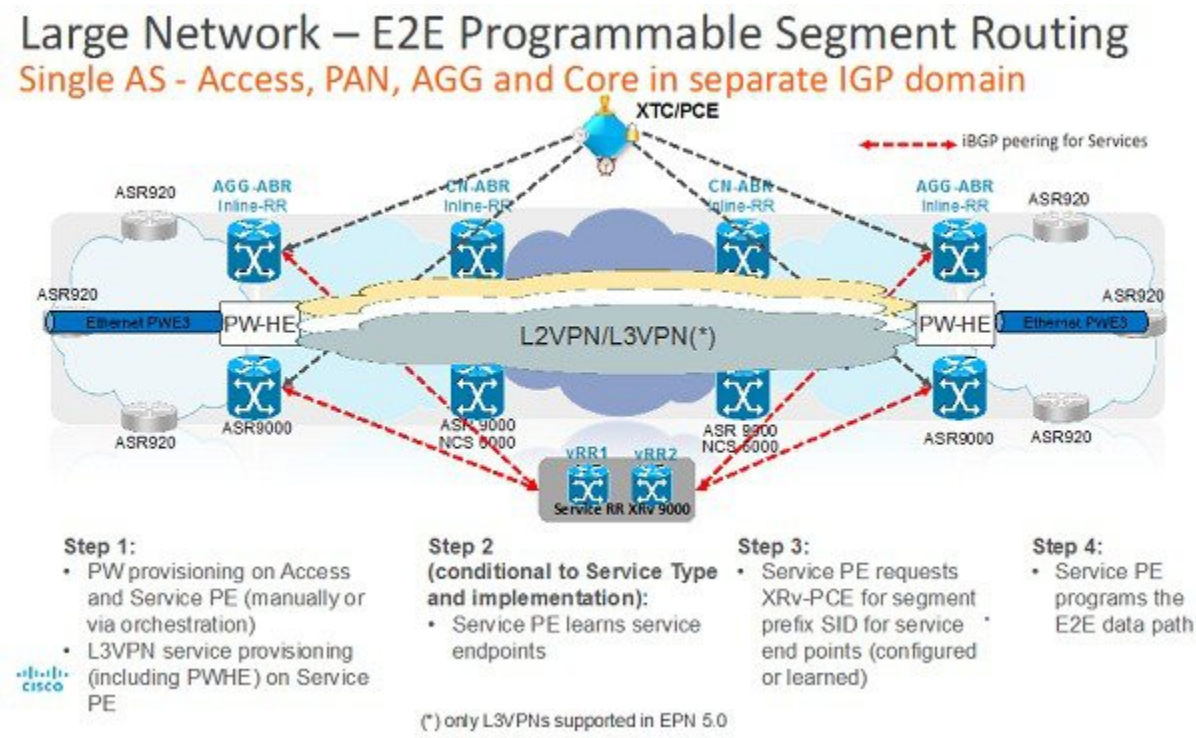

Figure 15: On-Demand Next-Hop

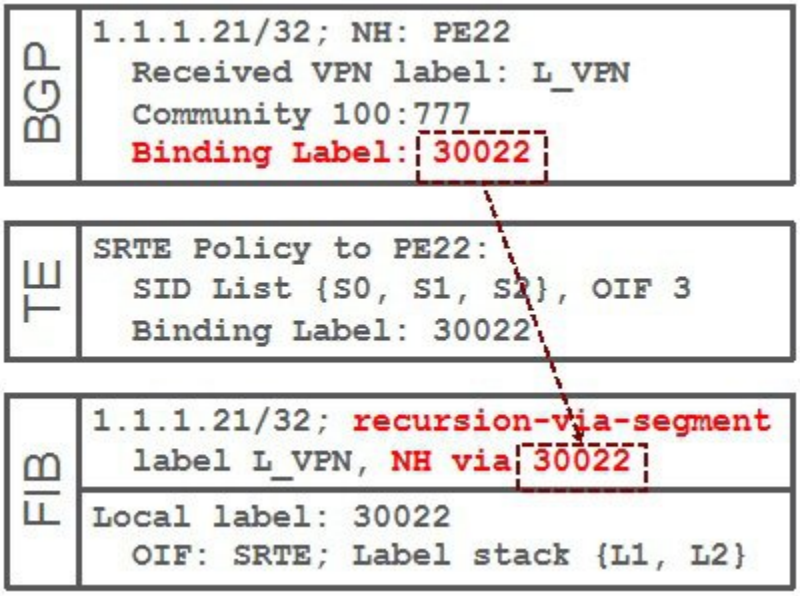

CHAPTER 4

Functional Components Design

This chapter contains the following sections:

Quality Services

The Cisco EPN system uses a Differentiated Services (DiffServ) QoS model across all the network layers of the transport network to guarantee proper treatment of all the services being transported. This QoS model guarantees the SLA requirements of all MEF and business services across the transport network. The QoS policy enforcement is accomplished with flat QoS policies using DiffServ queuing on all the NNIs and Flat QoS policies with parent shaping and child queuing on the UNIs and SE node interfaces.

The QoS design aims to satisfy the SLA requirements of business traffic, including real-time, such as telepresence, business-critical, and best-effort applications.

OAM

For EPN services, the Cisco EPN system utilizes a combination of protocols to provide the required service and transport OAM and PM functionality between the customer CPEs.

At a high level, for Service OAM, the EPN system employs the following tools:

- Ethernet OAM tools:
  - Ethernet Y.1731 Performance Management (PM) and Service Activation (Y.1564)—help to monitor the network performance such as, throughput, delay, jitter and loss that will help the customer to meet the SLA requirement.
- MPLS OAM tools for Layer 3 and Layer 2 transport:
  - MPLS VCCV PW OAM—monitors the state of the access PW for Business services that leverage an MPLS Access.

- MPLS VPN OAM—monitors Business VPN service from end-to-end.

For transport OAM, the EPN system employs MPLS LSP OAM to monitor the health of the unified MPLS transport. The Performance monitoring is based on Cisco IP SLA tools running between service end points or between any two points in the unified MPLS domain to find performance bottlenecks.

Static OAM

In case of Static Pseudowire (PW) scenario, the PW are statically provisioned and there is no control plane protocol, such as Transport LDP (T-LDP) to transmit the status of the PW or attachment circuit (AC) between the two Provider Edge (PE) devices at each end of the point to point service. The Static OAM (RFC 6478) is used to compensate the lack of signaling through a control plane protocol, by carrying the PW status code in-band with the PW data. By doing so, critical High Availability related behaviors, such as Remote Port Shutdown, can still be supported.

High Availability

The Cisco EPN system architecture implements high availability at the transport network level and the service level.

The Cisco EPN design is capable of meeting the stringent network recovery times that customers expect in their SLAs by utilizing various technologies throughout the network.

The *Cisco EPN Transport Design Guide, Release 5.0* addresses the implementation of high availability technologies at the transport layer that are common to all services.

For access service redundancy against local SE node and link failures, the following mechanisms are supported for different service and access types:

- For E-LAN services with IP and Unified MPLS access and VPLS or PBB-EVPN core, the AN runs actively and standby access PWs to the local SE nodes.
- For E-LAN services over G.8032 and REP enabled access, the access ring terminates on a redundant pair of SE nodes. During steady state conditions, the traffic load is shared across these nodes based on VLAN assignment, with each SE node acting as the RPL Owner role in the corresponding G.8032 instance. The ring is closed by the VFI/EVI core instances ensuring that services terminating at different nodes on the same ring are fully protected against ring failures.

For business services based on MPLS VPNs, the BGP Edge protection and BGP FRR Edge protection mechanisms are supported. The following mechanisms are supported for different access types:

- For IP/Unified MPLS access, the customer CE node is either:
  - Single homed to a AN node, with active/standby PWs toward a pair of SE nodes, with BFD-protected PE-CE static routing.
- For G.8032 access, the access ring terminates on a redundant pair of SE nodes that can either operate in active/active mode for the following:
  - A pair of VLANs that terminate at Layer 3 on different SE nodes. The SE node that does not terminate the VLAN bridges it to the other SE over an inter-SE PW to close the ring. This ensures

that a ring failure does not cause the CE to change primary SE node. The eBGP is used as PE-CE protocol and the primary SE is selected based on BGP local preference.

For core service redundancy against remote SE node and link failures:

- For E-line services, one-way PW redundancy is supported with active and standby PWs terminating at pairs of redundant SE nodes for the same customer site.
- For E-Line services with Unified MPLS access, two-way PW redundancy is supported.
- For E-LAN services, redundancy is ensured by the multipoint nature of the services, with pairs of SE nodes serving the same customer site.

CHAPTER 5

Orchestration Layer Design and Device Management Tools

This chapter contains the following sections:

Components of Service Orchestration Layer

The Cisco EPN solution is unique and not only de-couples the provisioning of overlay services from the underlying transport, but it also introduces a comprehensive orchestration/controller layer that enables faster service deployment and network programmability.

There are mainly four components involved in Cisco EPN services deployment and device management that play different roles in the network:

1. Network Services Orchestrator (NSO)
2. Evolved Programmable Network Manager (EPN-M)
3. XR Transport Controller (XTC)
4. Cisco Software Manager (CSM)

The above components play different roles in the network and therefore will be used for different purposes and applications. These include the following:

- Reliable and Optimized Transport to the Destination.
- Fast and Reliable Service Provisioning.
- Existing Network Service Discovery and Service Management.

- Service Assurance or Service Level Agreements (SLAs).

Figure 16: Services Workflow

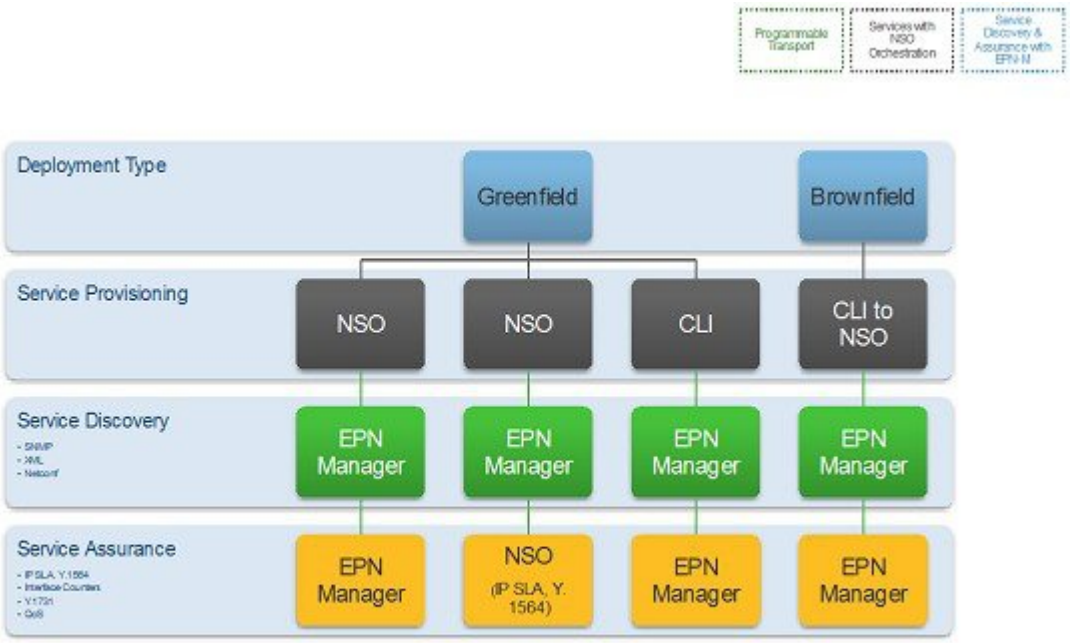

Network Service Orchestrator

The NSO is a management and orchestration (MANO) solution for network services and Network Functions Virtualization (NFV). The NSO includes capabilities for describing, deploying, configuring, and managing network services and VNFs, as well as configuring the multi-vendor physical underlay network elements with the help of standard open APIs such as NETCONF/YANG or a vendor-specific CLI using Network Element Drivers (NED).

In Cisco EPN, the NSO is used for Services Management, Service Provisioning, and Service Orchestration.

Note Cisco EPN 5.0 is not currently using NSO for NFV and that will be included in later EPN releases.

The NSO provides several options for service designing as given in the following:

- Service model with service template.
- Service model with mapping logic.
- Service model with mapping logic and service templates.

The selection of an option from the above list depends on the customer requirements.

The service model is a way of defining a service in a generic way. Once service is defined, the service model is also used to take an user input to enable the service. For example, the E-Line service requires two endpoints and a unique virtual circuit ID to enable the service. The end devices, attachment circuit UNI interfaces, and a circuit ID are parameters that are required and provided by the user to bring up the E-Line service. The service model uses the YANG modeling language (RFC 6020) inside NSO to define any service. For example, L3VPN.

Example of a service definition is given below:

```
container 12vpn—xc {
      tailf: cli-add—mode;
      when "(../../service-type = 'vpws') or (../../service—type = 'pbb-evpn-vpws') or
(../../service—type = 'l3vpn-pwhe')";
        leaf context {
          type string;
          config false;
        }

        leaf rem-pe-lpbk {
          type string;
        }

        leaf vc—id {
          type string;
          description "Virtual Circuit id";
        }

        container member {
          leaf xc—nember {
            type string;}
          }
        }
      }
```

Once the service characteristics are defined based on the requirements, the next step is to build the mapping logic in NSO to extract the user inputs. The mapping logic can be implemented using Python or Java. The purpose of mapping logic is to transform the service models to device models. It includes mechanisms of how service related operations are reflected on the devices. This involves mapping a service operation to available operations on devices. For example, when you add an access link to a VPN, how is this reflected on Customer Edge (CE) and Provider Edge (PE) routers.

Note The system uses Python only for mapping logic.

For more information, refer to the *Cisco EPN Service Orchestrator User Guide, Release 5.0*.

Finally, service template needs to be created in XML for each device types (IOS XE, XR, and so on) as shown in figure. In NSO, the service templates are required to translate the service logic into final device configuration through CLI NED. The NSO can also directly use the device YANG models using NETCONF for device configuration. These service templates enable NSO to operate in multi-vendor environment.

Note The Cisco EPN release only validates the service provisioning using CLI NED.

```
config—template xmlns:"http://tail—f.com/ns/config/1.Ø">
  <devices xmlns="http://tail—f.com/ns/ncs">
  <device>
    <name>{$PE}</name>
      <config tags="merge">
      <l2vpn-xconnect xmlns="urn:ios" tags="merge">
      <l2vpn>
      <xconnect>
```

```
    <context>
<name>{$L2_SERV_NAME}</name>
       <interworking>ethernet</interworking>
       <member>
         <member-list>
     <name>{$L2_SERV_IF}</name>
     <service-instance>{$L2_SERV_INST_ID}</service-instance>
  <member-list>
  <address-list>
           <name>{$REM_PE_LPBK_IP}</name>
    <vcid>{$L2_SERV_VC_ID}</vcid>
    <encapsulation>mpls</encapsulation>
  <address-list>
<member>
   <context>
   <xconnect>
   <l2vpn>
   <l2vpn-xconnect>
```

There are different options available in NSO for service designing. To understand all the different component involved in service designing we need to look at each option from design perspective.

- Option 1—service model with service template is normally suitable for those customers who offers limited amount of services and have fixed sertice parameters. Option 1 is easy to implement but service models and templates are specific for certain use case with this approach. It does not give flexibility to the user if customer wants to take some action based on some events.
- Option 2—provides lot of flexibility to the developer but is quite complex from operations point of view. Option 2 is suitable for those customers who wants to do a lot of customization as per their network needs and would like to take series of actions for certain events.
- Option 3—takes an optimal approach compared to option 1 and 2 and is suitable for most of the customers. It allows the flexibility to developer and also keep the operation simple.

The Cisco EPN uses option 3; service model with mapping logic and service template. This design approach is very modular, fits most customer requirements, and can be modified easily to adapt to specific customer needs

From operation perspective, NSO is used in cluster and HA mode to avoid point of failure in case of any disruption in the network environment.

Note The Cisco EPN uses NSO in non-cluster mode.

For details on MEF and business services implementation details, see *Cisco EPN Service Orchestrator User Guide, Release 5.0* Below is the table showing the service coverage of NSO from EPN 5.0 perspective.

Table 3: EPN 5.0 Services via NSO and CLI

EPN 5.0 Services	**Service Provisioning**	
	CLI	NSO
LDP VPWS	Y	Y
EVPN VPWS	Y	Y
Static VPWS	Y	Y

EPN 5.0 Services	Service Provisioning	
LDP VPLS	Y	Y
BGP-AD VPLS	Y	Y
L3 VPN	Y	Y
LDP H-VPLS	Y	Y
BGP AD H-VPLS	Y	Y
PBB EVPN H-VPLS	Y	Y
L3VPN with PW-HE	Y	Y

Y indicates that it is supported and tested in EPN 5.0

Evolved Programmable Network Manager

The Cisco Evolved Programmable Network Manager (EPN-M) is an element and network management system solution. It provides simplified, converged, end-to-end lifecycle management for carrier-grade networks of all sizes. The business can improve agility and operational efficiencies through automated device operations, fast provisioning, and proactive assurance.

The EPN-M is a single end to end management solution for IP, optical and cable access networks and capable of service provisioning, service assurance, device management and network management.

The Cisco EPN release leverages EPN-M capabilities and functions for the service discovery and assurance over the packet transport layer. The system also leverages EPN-M for overall network management functions including network topology discovery, device inventory, software management and fault/performance management.

For best practices, it is recommended to use EPN-M in High Availability mode to avoid single point of failure.

The EPN-M provides both APIs (Northbound and Southbound) and Graphical User Interface (GUI) for user interaction. The Large SPs can use NSO and EPN-M together. The NSO can be used for Multi-Vendor Service Provisioning and EPN-M can be used for service assurance and network management.

Below table shows EPN Use Cases tested and verified with EPN-M. The details of each use case can be found in respective implementation guides.

Table 4: EPN-M Service Usecase Coverage

EPN 5.0 Services	EPN-M Discovery	EPN-M Service Assurance	
		Y.1564/ IP SLA	Y.1731
LDP VPWS	Y	Y	Y
EVPN VPWS	NA	NA	NA

EPN 5.0 Services	EPN-M Discovery	EPN-M Service Assurance	
Static VPWS	Y	Y	Y
LDP VPLS	Y	Y	Y
BGP-AD VPLS	Y	Y	Y
L3 VPN	Y	Y	Y
LDP H-VPLS	Y	N	N
BGP AD H-VPLS	Y	NA	NA
PBB EVPN H-VPLS	NA	NA	NA
L3VPN with PW-HE	Y	NA	NA

- Y = Yes. It indicates that the feature is supported and tested in EPN 5.0.
- NA = Not Available. It indicates that the feature is not available and not tested in the current release of EPN 5.0.
- N = No. It indicates that the feature is not tested.

XR Transport Controller

The XR Transport Controller (XTC) is an IOS XR-based stateful PCE. As shown in below Figure, XTC is capable to collect the network topology using standards-based protocols such as BGP-LS, and also able to compute and deploy SR paths in the network through stateful protocol PCEP.

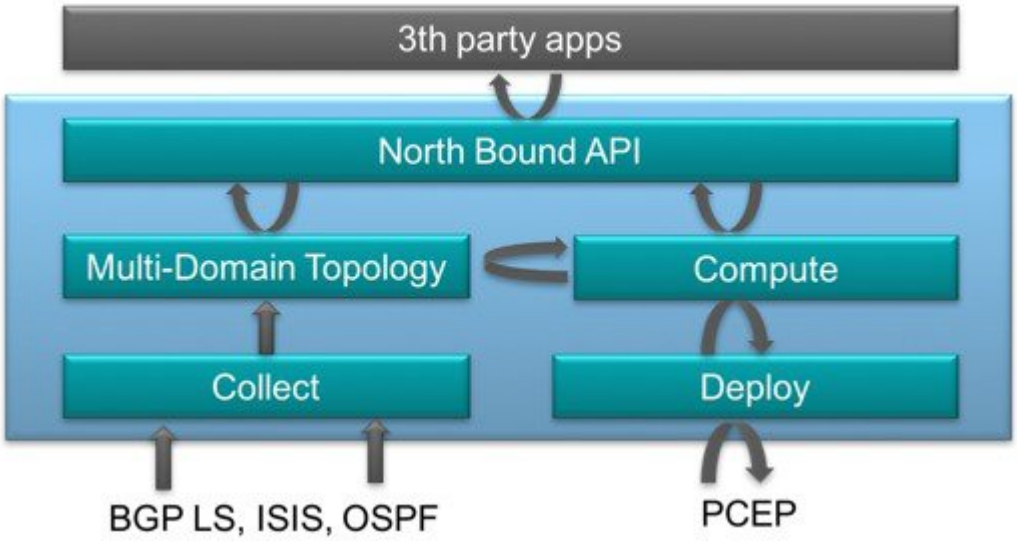

The XTC can be deployed as a centralized PCE Controller or it can be deployed as distributed PCE Router. The algorithms used by XTC for path computation is capable of handling multiple user constraints as well:

- Shortest path (IGP metrics).
- Traffic-Engineering metrics.
- Disjoint path.

In EPN, XTC is used as a centralized PCE and plays a critical role in deploying a dynamic SR paths for L3VPN service. Therefore, it is one of the critical component in EPN systems which brings programmability aspect inside the solution.

Note This PCE functionality can be hosted directly on a physical device or a virtual router. In the EPN 5.0, XRv9000 (Virtual Router) is used for PCE functionality.

Cisco Software Manager

The Cisco Software Manager (CSM) is a tool developed for the end-to-end automation of software deployment. The tool is capable of doing zero-touch upgrades of software images on network nodes, including SMUs. It can be customized to download such software directly from Cisco Connection Online (CCO) or from any local TFTP/HTTP servers as per customer needs.

The CSM supports all network devices included in Cisco EPN. It is responsible for bringing up all its nodes across all operative systems, and for the overall software and device management.

The Northbound REST APIs is created inside CSM to integrate with an application.

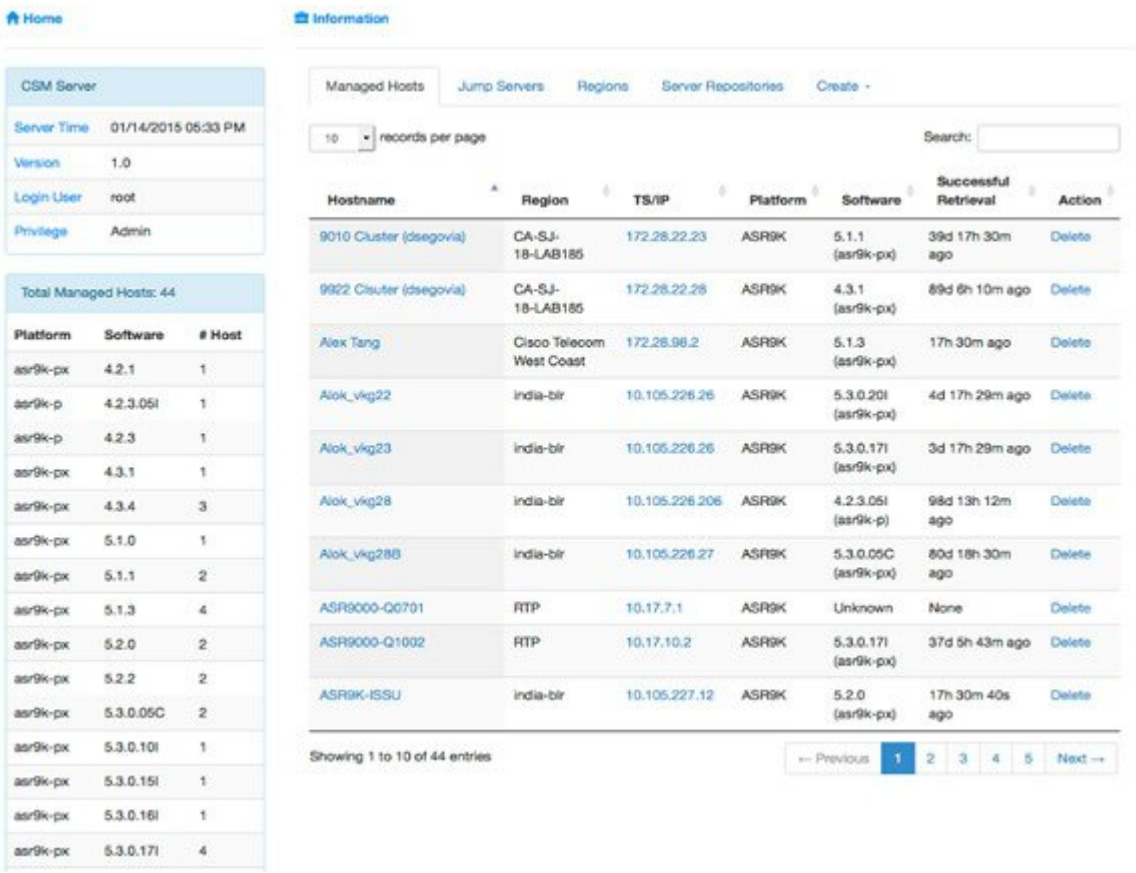

Note In the current release of EPN system, the CSM is only used for IOS XE and IOS XR software management.