Evolved Programmable Network (EPN) Foundation—Expanding MPLS to the Access

How to Deliver Scalable and Resilient Networks with Simplicity

Last updated: January 2013

Introduction

The Cisco Evolved Programmable Network (EPN) solution builds upon years of experience in delivering successful multi-service network deployments for service providers. As network technology evolved, Multiprotocol Label Switching (MPLS) became the de-facto transport protocol for core networks by providing superior performance and resiliency. Today, the Cisco solution expands MPLS into the edge with the benefit of carrier-grade level features that support the design of multi-service, massive scale networks with exceptional reliability and simplicity.

Service Provider Challenges

Service providers face the challenge to provide intelligent services that can quickly adapt to constantly changing market needs. New services such as cloud transport require unprecedented flexibility and elasticity from the network. Increasing bandwidth demands and decreasing ARPU put pressure on reducing network cost. At the same time, services need to be deployed faster and more cost effectively to stay competitive.

Expanding MPLS protocol to the access is intended to address the following challenges:

- Implementing multiple services over a single infrastructure with a unified operational model and achieve CAPEX and OPEX optimizations. Avoiding the maintenance of costly parallel SONET/SDH infrastructures for legacy services.
- Scaling multi-service deployments and maximizing the use of the capacity offered by the network.
- Increasing network reliability to protect potentially large failure domains with multiple services.
- Deploying network elements and services faster to improve time to market for new services, and shorten service delivery times.

While benefits of MPLS have been understood and implemented in the core networks, the expansion to the edge has been slower because of perceived increased complexity. Today, the Cisco MPLS solution brings additional value with the simplified operational aspects of MPLS, by introducing unique features for node insertion with addressing auto-negotiation, and plug and play prefix independent 50ms resiliency with remote LFA.

Evolution of the Network Transport Solution

The Carrier Edge Transport Network has gone through a significant evolution over recent years. Moving from TDM to Ethernet, service providers were able to increase bandwidth efficiency and reduce costs, but also faced the operational challenges of managing a large flood domains and limited scalability.

Over the years, service providers adopted MPLS in the core and aggregation layers in order to improve network scale, manageability, and resiliency. However, access networks relied predominantly on Ethernet with multiple protocols providing loop management and resiliency.

With the LTE mobile transport evolution to native IP, a strong desire exists to replace multiple networks with a unified network infrastructure managed by a single control plane. In the EPN design, by extending MPLS from the core/aggregation all the way to the access, service providers are able to realize a single resilient network that is capable of carrying multiple services and offers unprecedented scalability and operational simplicity.

MPLS in the Access

The Unified MPLS solution expands MPLS transport to the access platforms. It brings the benefit of an end-to-end converged network, unifying Ethernet, ATM, and TDM support. A single delivery model for all services reduces operational complexity, while the MPLS network offers superior end-to-end service resiliency.

Protocols

What differentiates aggregation and core networks from the access networks are the number of network elements that need to be managed, and frequent operational interventions to add new services and insert new nodes. Therefore, the key requirement for the solution is operational simplicity.

In order to deploy the IP/MPLS solution, access nodes require the IGP routing protocol, ISIS or OSPF, along with label distribution protocol. For bandwidth management, the solution will implement infrastructure QoS with aggregate traffic classes that are assigned to separate queues, delivering the required performance characteristics.

Scalability

In order to scale the infrastructure further, IGP domains can be partitioned into different areas, preventing unnecessary IP addresses advertisement between access and pre-aggregation rings. iBGP with RFC 3107 implemented at the access nodes provides controlled prefix and label distribution to build end-to-end hierarchical LSPs. Services, such as L2 or L3 VPNs, are transported over end-to-end hierarchical LSPs. This mechanism optimizes label and prefix learning at access nodes only for those necessary to build services for those access nodes. Transit nodes transport services transparently to the end destination, or ASBR, as illustrated in the figure below.

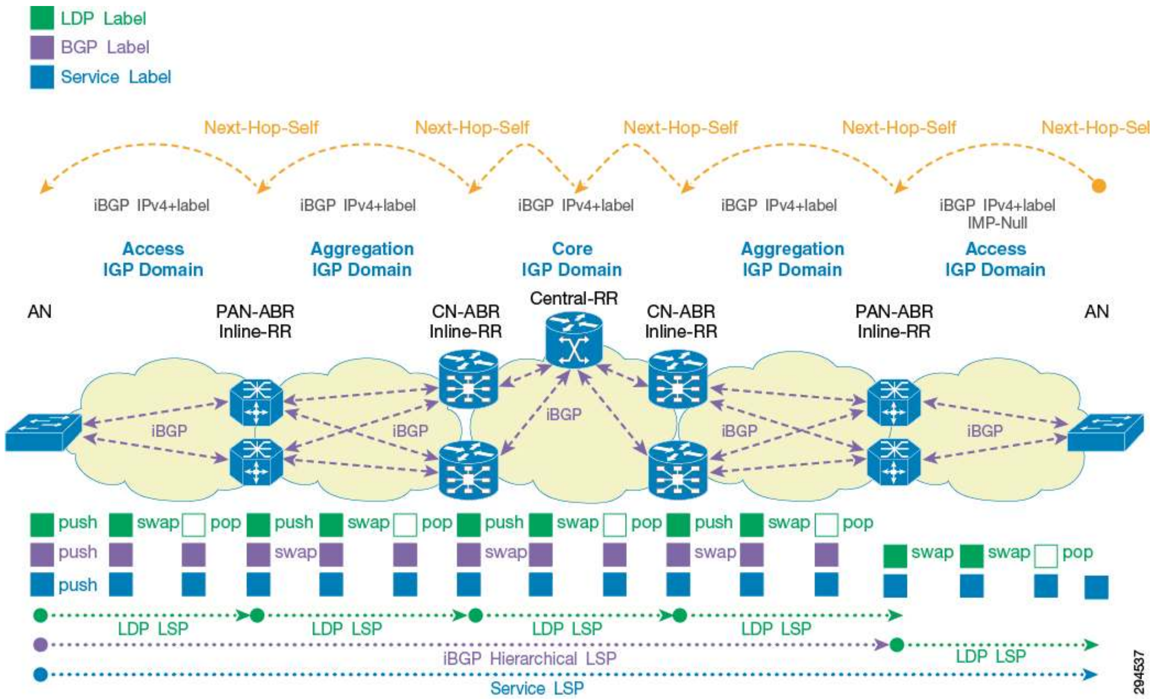

Resiliency

The mechanism widely deployed for 50ms resiliency in MPLS networks is Traffic Engineering Fast Re-Route (TE-FRR). However, in the access network, implementing MPLS-TE would represent a very high number of tunnels and introduce operational overhead.

The alternative resiliency method is Loop Free Alternate (LFA) technology, a Cisco-driven innovation providing 50ms convergence for IP/MPLS networks. LFA is based on the local IGP protocol knowledge about alternate routes. The alternate route is pre-installed in hardware and ready to take over immediately if the primary route becomes unavailable. LFA technology is the foundation of simple resiliency in MPLS access.

Remote LFA extends the LFA capability to non-ECMP networks, addressing, among others, ring topology, which is pre-dominant in Carrier Ethernet pre-aggregation. With Remote LFA (rLFA), IP/MPLS traffic is protected with no operational overhead. The operator needs only to activate rLFA for the prefixes that are to be protected. Moreover, LFA provides protection to all services implemented over the MPLS access infrastructure.

Services Implementation

Multi-service networks must support revenue-generating services:

- Mobile backhaul services continue driving networks growth in support of the increasing bandwidth requirements of mobile services. MEF CE 2.0 E-Line or E-LAN services, as well as RFC4364 L3 VPN services, can deliver scalable mobile backhaul. Synchronous Ethernet and IEEE 1588v2 timing protocols are also required for the mobile backhaul infrastructures.
- Business VPN and wholesale services rely on MEF CE 2.0 E-Line, E-LAN, E-Access, or L3 VPN options. These transport technologies, paired with strong EOAM implementation, deliver SLA-managed business services.

- Multicast transport—E-Tree or L3 multicast can be deployed over the MPLS infrastructure. With strong replication implementation, the network can carry very high number of flows under efficient and stringent service levels.
- Cloud services—L2 VPN or L3 VPN, paired with SDN capabilities, provides a set of tools to provision flexible private cloud services with the capability to adjust to application demands.

The following diagram illustrates service delivery models for Unified MPLS.

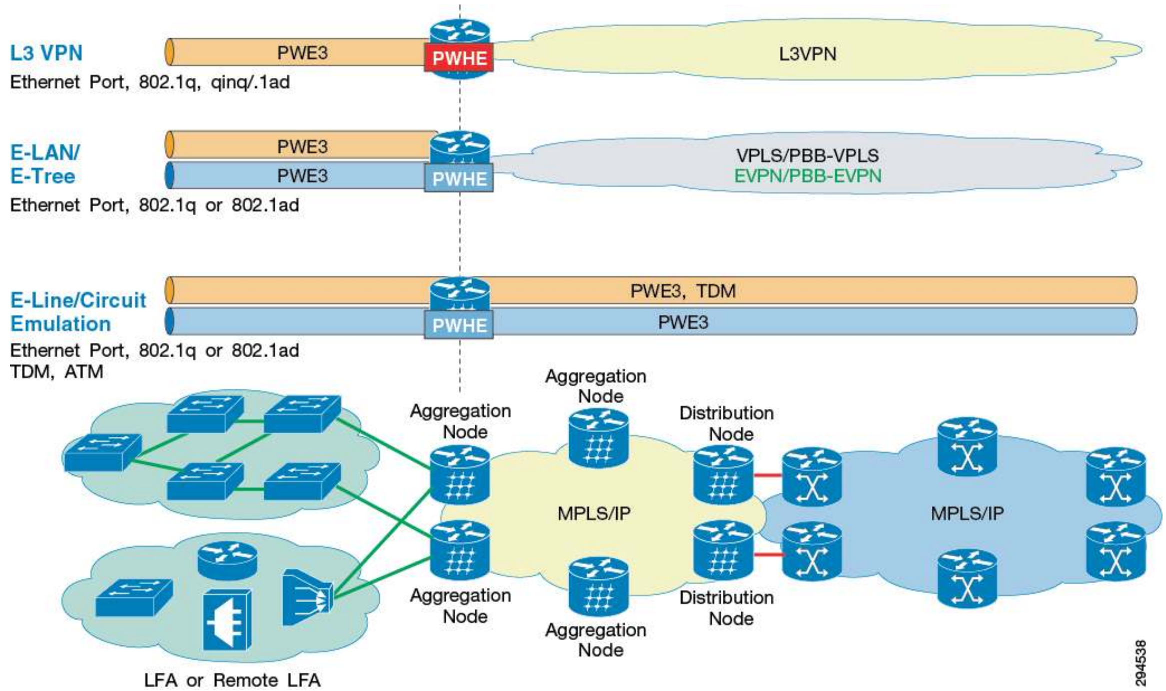

L2 VPN Services

MEF CE 2.0 services are implemented over MPLS via the use of pseudowire or VPLS, both well-known standard protocols. E-LAN and E-Tree services can now be delivered via PBB-EVPN, bringing significant advantages to L2 VPN implementation.

EVPN consists of routing MAC addresses via BGP. MAC addresses are learned through data flows and advertised via BGP control plane. The technology eliminates the requirement for T-LDP and pseudowires, simplifying the overall solution and improving scalability.

PBB encapsulation is critical when it comes to high scale and multi-service networks. Tier-1 L2 VPN or data centers, in particular, can be exposed to high MAC scale with significant MAC learning and aging activity. For these use cases, PBB brings the advantage of advertising only provider bridge addresses into BGP. This eliminates potential increased BGP activity through the constant customer MAC routes advertisements.

E-Access delivers a standardized way of implementing wholesale services between service providers. E-Access implements a point-to-point hand-off between providers with standardized QoS and VLAN transparency handling. It can be delivered with 802.1ad, as per standard specification for supported platforms, or with 802.1 QinQ encapsulation.

L3VPN Services with Pseudowire Head-End and Service Multiplexing

Layer 3 services are delivered via RFC4364 MPLS VPN. Cisco has simplified the solution further by implementing L3 VPN service termination, logical sub-interface, directly on a pseudowire. The technology, Pseudowire Headend (PWHE), provides integration of transport and service termination into one system.

PWHE allows further service multiplexing to transport multiple VLANs, which can provide separate services over a single transport pseudowire. Logical subinterfaces represented via an inner/outer VLAN combination can be configured on a pseudowire and associated with different services, VRF, VPLS, or multicast. The implementation provides full support for QoS and service attributes on a pseudowire, such as security ACLs, NetFlow, or storm control.

Service definition is centralized in a single network element, which reduces both design and implementation time, and results in a significant operational improvement.

In addition, avoiding the L2 hand-off between pre-aggregation and the Ethernet service edge simplifies significantly the resiliency design. In the traditional L2 hand-off, resiliency has to be designed between the transport and service edge via L2 protection mechanisms (e.g., spanning tree, MC-LAG). In the MPLS approach, MPLS links between the transport and service edge are, by default, protected via IGP resiliency, or LFA/rLFA as per recommended design, and do not require any additional design.

Network Operations

Network operations simplicity is critical for access and aggregation infrastructures because of the large amount of equipment involved and the frequency of changes required. New equipment insertions need to have zero-to-minimal impact on existing services and need to be able to carry new services as soon as possible. At the same time, truck rolls need to be minimized.

Cisco is introducing multiple innovations that are aimed at simplifying the operations of the MPLS infrastructure. Some of those innovations are described below.

Infrastructure Activation

While activating the network elements, service providers perform multiple activities: a device needs to be sent to the site, pre-loaded with a configuration or provisioned on-site, connected to the network, provisioned with IP addresses, etc. Only after the protocols are up and reachability has been established is the device ready for provisioning from the NOC.

Cisco has introduced Auto-IP technology across all platforms to ease one of the aspects of node insertion, IP address management, and provisioning. As illustrated below, the technology avoids the necessity of neighboring nodes reconfiguration. It also minimizes the need to manage local link addresses.

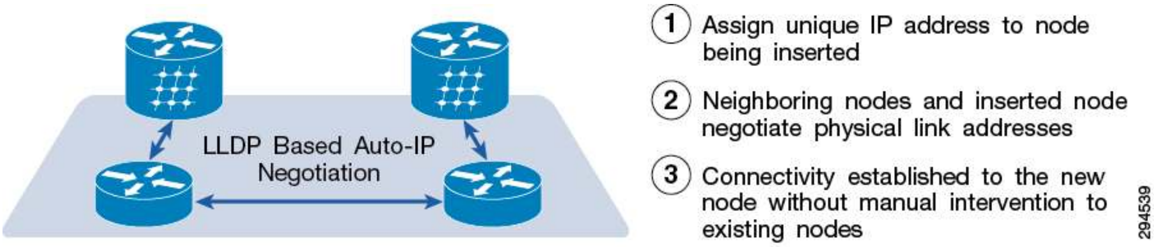

Further development is made to automate the process of bringing a device up. The technology umbrella, Autonomic Networking, allows the node to connect to a secure management infrastructure and boot up without operator intervention on-site.

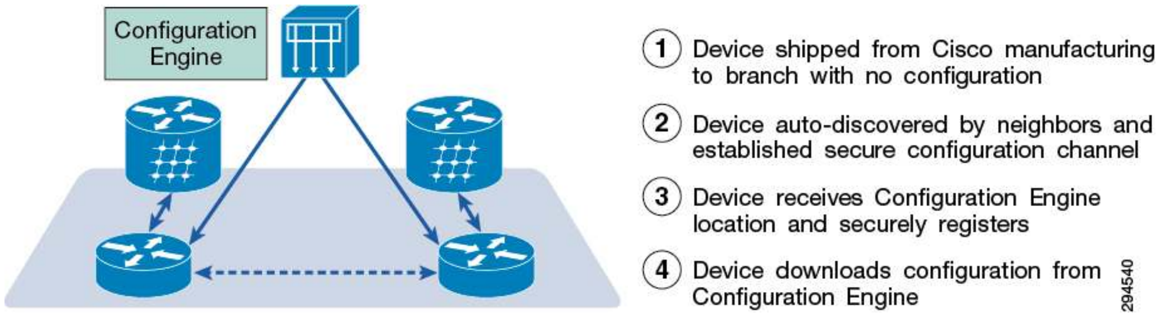

Service Activation

Once the infrastructure is ready for service activation, service provisioning is performed. The Cisco Service Activation framework reduces truck rolls and streamlines service turn-up procedure. With the Y.1564 traffic generator integrated in the access platforms, traffic is injected into newly provisioned service. On the other side of the circuit, intelligent loopback executes MAC swap and loopback for traffic under test per specified parameters, leaving the remaining traffics flow unaltered.

The following diagram represents the components of service activation framework.

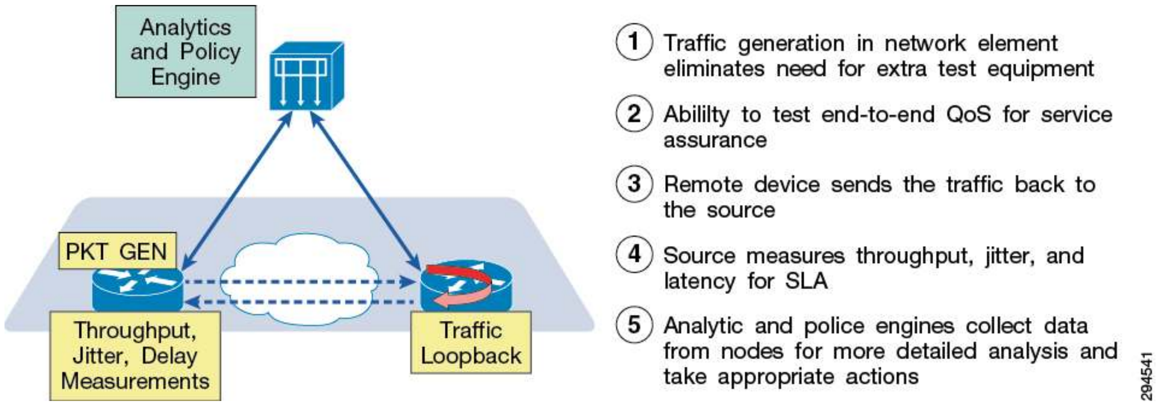

Service Monitoring

As part of the Carrier Ethernet portfolio, the Cisco solution provides a comprehensive suite of Ethernet OAM protocols to monitor infrastructure and services:

- Link OAM—A IEEE 802.3ah-compliant protocol that is used to monitor point-to-point PE-CE links
- E-LMI—This standard protocol monitors the status of Ethernet Services on PE-CE links
- CFM—IEEE 802.1ag-compliant implementation provides a set of tools for connectivity monitoring for point-to-point and multi-point services
- Y.1731 Fault Management—AIS and RDI mechanisms that allow propagating alarms within the network and between maintenance domains
- Y.1731 Performance Management—Delay measurements and synthetic loss measurement provide a mechanism for service providers to monitor services proactively and ensure SLA delivery

Cisco Edge e2e Portfolio

The Cisco Unified MPLS Architecture is powered by a very flexible hardware portfolio. The access and aggregation product portfolio ranges from compact one-rack unit form factor equipment, to multi-rack redundant and modular platforms for Ethernet and TDM transport. The Ethernet service edge provides ultra-high service scalability and feature richness. All platforms provide the support for Unified MPLS as well as Ethernet aggregation, providing a smooth evolution from technologies deployed today.

Cisco Access, Aggregation and Edge Portfolio with Prime Management.

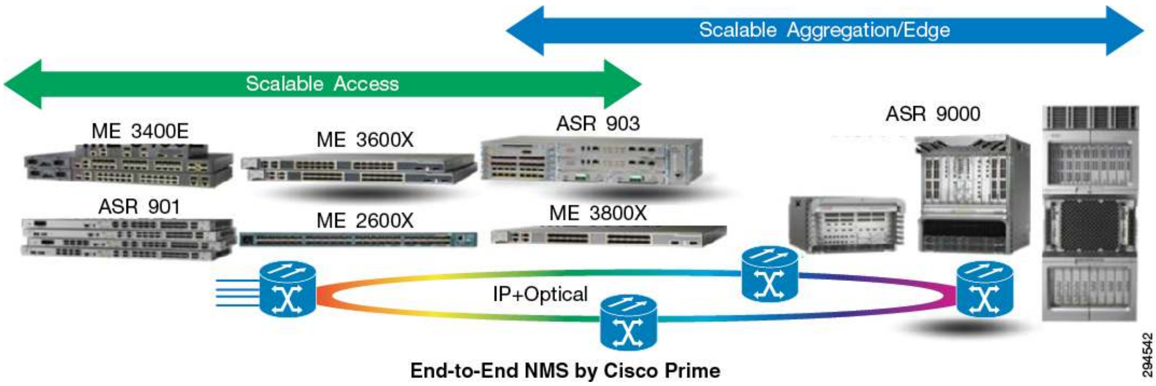

Cisco Validated Design

The Unified MPLS solution is a building block of the evolved programmable network from Cisco Validated Design (CVD). L2VPN and L3VPN services, along with multicast transport, are validated over a scalable and resilient infrastructure using Cisco industry-leading technologies such as rLFA, Pseudowire Head-End (PWHE), and PBB-EVPN. Cisco Validated Design provides a set of tested configurations that service providers can use as building blocks for scalable and reliable architectures.

Conclusion

Cisco EPN provides an industry leading end-to-end Carrier Ethernet transport solution that addresses major challenges faced by service providers, shortening delivery time for new services, decreasing CAPEX to allow for network growth, and lowering the OPEX of service activation.

With significant innovation implemented across the portfolio, Cisco provides the option to integrate with today's technologies and evolve to the next generation/state of the art within the same hardware portfolio.

With CVD quality, the solution provides fast and reliable deployment choice for service providers access, aggregation and edge networks.

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 C11-730477-00 01/14

