

The Evolution of Carrier Ethernet Services — Requirements and Deployment Case Studies

*Luyuan Fang, Cisco Systems Inc., Nabil Bitar, Verizon,
Raymond Zhang, British Telecom
Michael Taylor, AT&T*

ABSTRACT

Carrier Ethernet is one of the fastest growing areas in the data communication field in recent years. In particular, the combination of MPLS core transport technologies in conjunction with IEEE 802.1ah in aggregation is one of the most interesting recent developments in the Carrier Ethernet evolution. This article provides a brief overview of the Ethernet technologies and requirements from the perspective of the service provider. We use three deployment case studies to illustrate the use of different technologies and discuss network design considerations. We identify the important factors influencing network design decisions and the future work in building Carrier Ethernet networks and services.

INTRODUCTION

Ethernet — the technology originally designed for simple data sharing over a local area network (LAN) in campuses or enterprises — is now becoming a major player in the next-generation carrier networks market, together with Internet Protocol (IP)/multiprotocol label switching (MPLS) technologies. Carrier Ethernet is being deployed for point-to-point links interconnecting IP/MPLS routers, for native Ethernet services to end users for Layer 2 (L2) virtual private networks (VPNs), and as an access technology to IP/MPLS services.

One of the main drivers for the development of Carrier Ethernet is the growing demand of high-bandwidth applications at increasingly lower costs. Therefore, Carrier Ethernet networks must support both existing and emerging services, including business mission-critical services, residential services, mobility services, and wholesale services. The applications supported by these services include video, voice, and high-speed data services for Internet access, private IP, or Ethernet VPNs.

The Ethernet services offered today in service provider networks are based on IEEE 802.1ad (802.1ad) provider bridges or Q-in-Q, IEEE 802.1Q (802.1Q) virtual LANs, MPLS, and the combinations of these technologies. Network evolution will involve the integration of IEEE 802.1ah (802.1ah) provider backbone bridges (PBB) [1] and possibly the future IEEE 802.1Qay provider backbone bridge traffic engineering in these networks. The choice of the specific technology depends on the service; geographical scope (e.g., metro, national, or international); the service evolution as a function of time and available technologies; service mix (e.g., BGP/MPLS IP VPN, any L2 over MPLS); and finally, operational constraints.

This article focuses on using MPLS and 802.1ah as a complementary design approach to provide more scalable and reliable network solutions for Carrier Ethernet service deployment. The MPLS and 802.1ah interworking architecture has been proposed in the Internet Engineering Task Force (IETF) [2].

The structure of this article is as follows. We first provide a brief overview of the Carrier Ethernet aggregation and access technologies available today, which include IEEE 802.1Q, 802.1ad, 802.1ah, virtual private line service (VPLS), and hierarchical-VPLS (H-VPLS). Second, we discuss the Carrier Ethernet requirements in the areas of security, scalability, reliability, quality of service (QoS), and operations, administration, and maintenance (OAM). Third, this article provides three deployment case studies and design proposals to illustrate the benefits and challenges of using different technologies to provide Carrier-class Ethernet services. Finally, we conclude our discussion by identifying the important factors influencing network design decisions and the future work toward building next generation Carrier Ethernet networks.

Based on the Metro Ethernet Forum's definitions, there are two broad categories of Carrier Ethernet services: point-to-point, referred to as E-Line services; and multipoint, referred to as E-LAN services.

CARRIER ETHERNET TECHNOLOGIES OVERVIEW

This section provides a brief overview of the building blocks of Carrier Ethernet aggregation or access technologies.

IEEE 802.1Q — VIRTUAL LANs (VLANs)

802.1Q provides for tagging Ethernet frames with VLAN IDs. It provides the mechanism that enables multiple-bridged networks to transparently share the same physical network while maintaining the isolation between networks. Ethernet switches deliver packets within the same VLAN and send the traffic between different VLANs to internal or external routers to perform the routing function. 802.1Q only supports up to 4094 VLANs, which is a scaling constraint for service providers.

802.1Q-IN-802.1Q (Q-IN-Q)

Q-in-Q enables VLAN stacking that supports the appending of multiple VLAN tags to the same Ethernet frame creating a hierarchy. Q-in-Q started as a proprietary implementation to overcome the 4094 VLAN limit in an 802.1Q network; however, it became the de facto standard for preserving customer VLAN settings and providing transparency across a provider network. It is widely used in Ethernet deployments today.

IEEE 802.1AD — PROVIDER BRIDGES (PB):

802.1ad standardizes the architecture and bridged protocols to allow Ethernet frames with multiple VLAN tags. It also defines the labels for customer VLANs (C-VLANs) and service VLANs (S-VLANs) and introduces a customer medium access control (MAC) address space and a provider MAC address space for L2 control protocols. 802.1ad provides separate instances of the MAC bridging services to multiple independent clients of a bridged local area network by adding/removing S-VLANs. This does not require coordination among the users and only requires minimum coordination between the users and the provider of the bridging service.

IEEE 802.1AH — PROVIDER BACKBONE BRIDGES (PBB) [1]

802.1ah defines bridge protocols and an architecture for the interconnection of provider bridged networks (PBNs). It is compatible and interoperable with PBN protocols and equipment. It allows a provider to support up to 2^{24} (~16 million) service instances, as opposed to 2^{12} (~4000) service instances in a PBN. At the edge of a PBB network (PBBN), an Ethernet frame is associated with a service instance based on the S-VLAN in the frame header; a PBB MAC header encapsulates the customer frame, including the C-MAC header. The PBB header is composed of a source and destination backbone MAC address (B-MAC), a backbone VLAN ID (B-VID) to segregate the backbone into broadcast domains, and a 24-bit service instance identifier (I-SID) in a service instance

tag (I-Tag). Figure 1 illustrates the network interconnections and the encapsulation protocol stacks.

VIRTUAL PRIVATE LINE SERVICES (VPLS) (RFC4762 [3], RFC4761 [4])

VPLS is a multipoint L2 VPN technology that enables multiple sites to be connected over a emulated Ethernet broadcast domain across an IP/MPLS network. VPLS evolved as a logical extension of virtual private wire services (VPWS) based on RFC4447 [5]. Ethernet VPWS provides point-to-point (P2P) Ethernet-based L2 VPN services. A VPLS can be defined as a group of virtual switch instances (VSIs) that are interconnected to form a single logical bridge domain. A VSI is similar to the bridging function defined in IEEE 802.1Q; a frame is switched, based on the destination MAC and membership in a L2 VPN. It floods unknown, broadcast, or multicast frames to all ports associated with the VSI (Table 1).

HIERARCHICAL VPLS (H-VPLS)

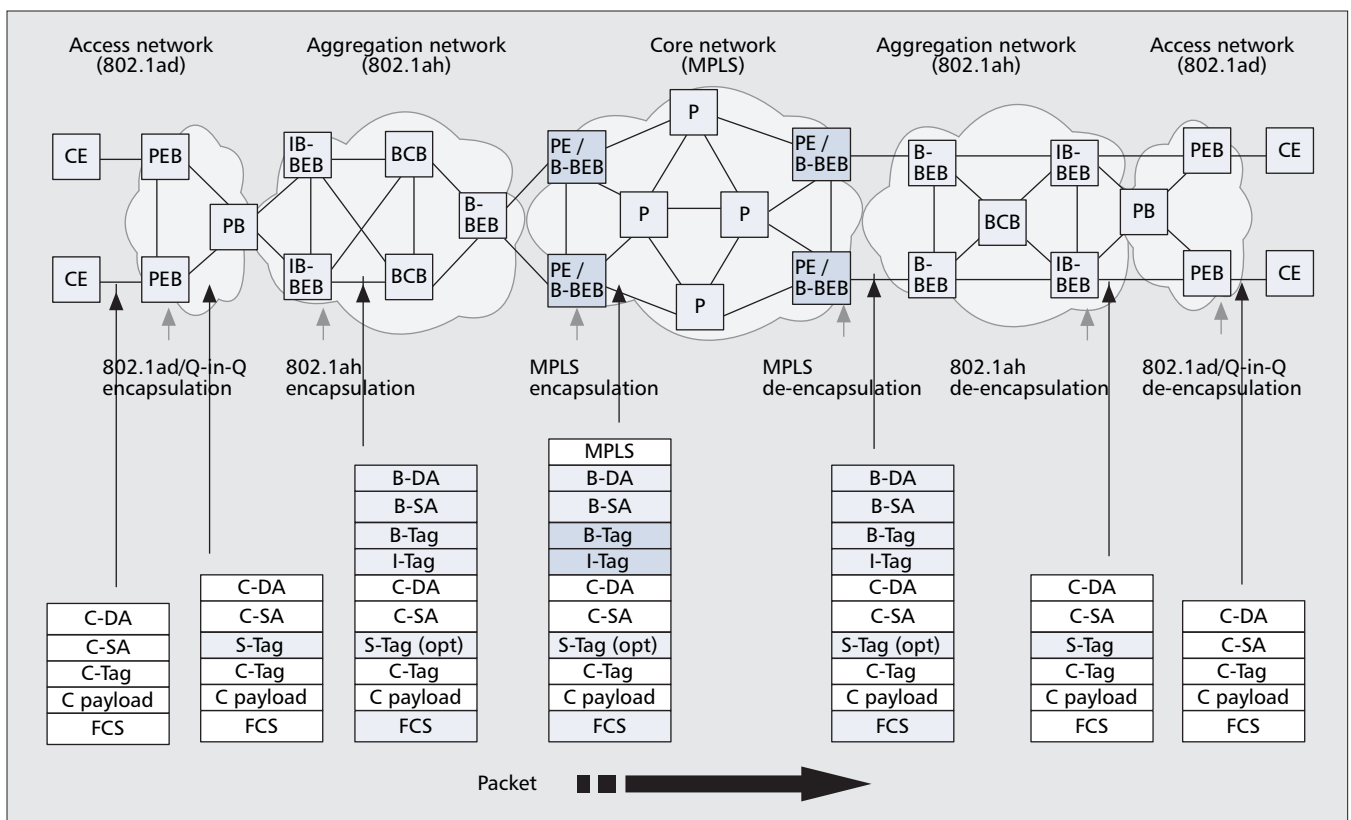
H-VPLS is introduced to improve the scalability of VPLS. H-VPLS partitions the network into several edge domains that are interconnected using an MPLS core. Full mesh VPLS connectivity is limited to only the core network among network provider edges (nPEs). The user PEs (uPEs) need learn of only their local nPE devices. A spoke pseudowire (PW) is used as an attachment circuit (AC) from a uPE to a VPLS instance at an nPE. This mechanism may subsequently be used to partition a core MPLS network into VPLS islands interconnected by PWs to reduce the requirements on PW meshness among nPEs or to optimize multicast by reducing replication at an nPE to a smaller set of nPEs. The edge domain also can be built using 802.1Q, 802.1 Q-in-Q, or 802.1ad Ethernet networks, described earlier in this section.

CARRIER ETHERNET SERVICE REQUIREMENTS

Based on the Metro Ethernet Forum's (MEF) definitions [6], there are two broad categories of Carrier Ethernet services: point-to-point, referred to as E-Line services; and multipoint, referred to as E-LAN services. Both E-line and E-LAN services are often provided with multiple classes of service (CoS); where a single Ethernet virtual connection (EVC) can carry traffic with one or more CoS. Service providers desire to build networks that offer all services simultaneously on a single converged infrastructure.

In addition to the previous MEF services (i.e., E-Line and E-LAN), there is the advent of E-Tree, which is a rooted-multipoint EVC. E-Tree is targeted toward hub-and-spoke applications and multicast streaming from a source to multiple receivers. E-Tree is a part of the ongoing MEF Phase 2 service work. Because this work is in its infancy stage at the MEF, the service is not discussed in detail in this article.

Table 2 highlights the carrier-class networks and service requirements.



■ **Figure 1.** Example of using PBB (802.1ah) for aggregation and MPLS for core transport.

BEB	Backbone edge bridge — encapsulates customer frames for transmission across a PBBN.
BCB	Backbone core bridge — bridges frames based on B-TAG and B-MAC information, similar to an 802.1ad bridge, in the PBBN core.
B-BEB	B type BEB — contains a B-component. It supports bridging in the provider backbone based on B-MAC and B-TAG information.
I-BEB	I type BEB — contains an I-component for bridging in the customer space based on customer MAC and service VLAN ID.
B-TAG	Backbone VLAN tag — has a similar format to an 802.1ad S-TAG.
I-TAG	Service Instance tag — encapsulates customer addresses and contains the Service Instance identifier (I-SID).
I-SID	Service Instance identifier — A field of the Service Instance tag which identifies the service instance of the frame.
S-TAG	A field defined in the 802.1ad Q-in-Q encapsulation which identifies the Service VLAN (S-VLAN).

■ **Table 1.** IEEE 802.1ah terminologies.

DEPLOYMENT CASE STUDIES

SCENARIO 1: USING 802.1AH FOR AGGREGATION AND VPWS/VPLS IN MPLS CORE FOR EXTENDING ETHERNET SERVICES

The scenario described in this section illustrates how multiple technologies could be used in a hierarchical fashion in a service provider network to address the service provider's requirements for scalable carrier-class global Ethernet services networks.

Figure 2 depicts a potential global service provider network offering Ethernet services. There are three main layers in the network: an

access layer based on 802.1ad or Q-in-Q, referred to as a PB domain; an IEEE 802.1ah PBB layer; and an MPLS layer. The MPLS layer is built on an IP-MPLS network that can be decomposed into multiple layers from an IP routing viewpoint. An E-Line or E-LAN service can span a single PBN island, two or more PBN islands interconnected by a PBBN, or two or more PBBN islands interconnected by an MPLS core network.

In a PB network island, an EVC is identified by a service-provider VLAN identifier (S-VLAN ID). An S-VLAN ID is 12 bits, which imposes the limitation of 4094 service instances in a PB network. Routing in a PB network is based on spanning tree protocols (STPs), such as STP

Services	Support E-Line, E-LAN services currently defined by MEF, and E-Tree in the future. Service multiplexing on the same UNI of multiple EVPL and ELAN EVCs. A service in this case must be identified by port+VLAN ID and mapped to a service provider tag that encapsulates the customer tagged Ethernet frame.
Security	Prevent non-authorized access to the network. Drop traffic tagged with unassigned VLAN on a UNI or external NNI (E-NNI). Provide layer2 control separation between customer and provider using 802.1ad or MAC translation in 802.1Q networks. Resource usage control per customer, including MAC entries and bandwidth per CoS. Maintain service privacy by keeping traffic for an EVC constrained to the EVC. Apply security techniques, including protocol authentication, encryption, access control, and infrastructure isolation at MPLS I-NNI and E-NNI.
Scalability	Support tens of thousands of EVCs in a large metro to hundreds of thousands across a WAN for enterprise services. Residential service transport over Ethernet often requires support for hundreds of thousands to millions of EVCs in a large metro. Support tens to hundreds of thousands of MAC addresses for E-LAN services in a large metro and hundreds of thousands or millions across a WAN for global services. Support a large number of edge devices. The number of these devices varies per service.
Reliability	Support resilient network elements, including stateful control plane switchover, to minimize the likelihood of service interruption during planned software upgrades or unplanned control plane failures. Support nonstop forwarding for established services during control plane failure of a device; support fast reroute for node and link protection, fast network convergence after a network event, and fast failure detection using OAM mechanisms to invoke fast switchover of traffic to an available alternative path in the network.
QoS	Support a hybrid of priority scheduling, and Weighted Fair Queuing with Weighted Random Early Discard or tail drop for buffer management. Support traffic profile enforcement on an UNI or an E-NNI per EVC per CoS. This enables the carrier to control resource sharing in the core among customers, and provide service level agreements (SLAs) based on a clear traffic contract defined by the traffic profile. Preserve customer markings in the P bits of the customer VLAN tag by encapsulating the customer frame with a service tag and provider assigned p-bits in a native Ethernet network or an MPLS header and provider assigned exp value in an MPLS network without modifying customer markings.
Manageability	Minimize network touchpoints in provisioning: support IEEE 802.1ak — Multiple Registration Protocol and Ethernet Local Management Interface (E-LMI) for native Ethernet networks; BGP auto-discovery [7], and RADIUS Dynamic Authorization, or LDP/BGP for signaling Ethernet service over MPLS networks. Support standards-based OAM mechanisms for failure detection, notification, and performance measurements, for example, IEEE 802.1ag (802.1ag) — Connectivity Fault Management, IEEE 802.3ah — Ethernet in the First Mile; ITU Y1731 — OAM functions and mechanisms for Ethernet-based networks; and IETF Pseudowire Virtual Circuit Connectivity Verification (VCCV) [8] and MPLS Ping [9] for MPLS networks carrying Ethernet services. Support native Ethernet and MPLS OAM interworking, such as VCCV and LDP failure notification interworking with 802.1ag and Y.1731.

■ **Table 2.** Carrier Ethernet service requirements.

defined in IEEE 802.1ad, Rapid STP (RSTP) — IEEE 802.1w, or Multiple STP (MSTP) — IEEE 802.1s. Signaling is based on Generic VLAN Registration Protocol (GVRP) — 802.1Q, or Multiple VLAN Registration Protocol (MVRP) — IEEE 802.1ak. GVRP provides 802.1Q compliant VLAN pruning and dynamic VLAN creation on 802.1Q trunk ports. MVRP is part of Multiple Registration Protocol (MRP). Furthermore, MVRP will provide for the rapid healing of network failures without interrupting services to unaffected VLANs. Alternatively, manual configuration can be used. Multicast, broadcast, and flooding in a PB are confined to an S-VLAN and use native Ethernet capability.

The PBB layer provides scalability in two areas:

- Support large numbers of service instances
- Customer MAC address (C-MAC) hiding

The ingress backbone edge bridge (BEB) maps a 12-bit VLAN ID from the PBN to a 24-bit I-SID in the I-Tag. A B-VID is used to build point-to-point or multipoint tunnels between BEBs. The core of the PBBN uses provider MAC bridging. Customer MAC addresses are tunneled through the PBBN. As a result, PB islands, including connected BEBs with shared EVCs, are required to learn only the customer MAC addresses at the ends of these EVCs. Similar to PBNs, routing

in the PBBN is based on STP, and signaling for B-VID registration is based on GVRP or MVRP. Spanning tree instances, GVRP, and MVRP in a PB island (including a connected I-component of a BEB) are confined to that island and are separate from those used in the PBBN. Multicast, broadcast, and flooding in a PBBN are confined to a service instance. A PBBN assigns a multicast MAC address per I-SID and uses Multicast Multi Registration Protocol (MMRP) to register the MAC address by BEBs.

When extending Ethernet services over MPLS, E-Line is based on VPWS (RFC4447), whereas E-LAN is based on VPLS (RFC4762 and RFC4761). Each E-Line EVC is mapped to a PW and each E-LAN EVC is mapped to a VSI at an MPLS PE. RFC4762 uses the Label Distribution Protocol (LDP) and PW control in RFC4447 to build a mesh of PWs that interconnect PEs for a VSI. A network may use Multi-Border Gateway Protocol (MP-BGP) for VSI auto-discovery [7] or rely on manual or operation support system (OSS)-driven configuration to tell a PE which other PEs have the same instance in a L2 VPN service. RFC4761 uses MP-BGP for VPLS autodiscovery and signaling. At a PE, an S-VLAN is used as an AC to identify the service instance (PW or VSI). If the PE is PBB-capable and connected to a PBBN, the ser-

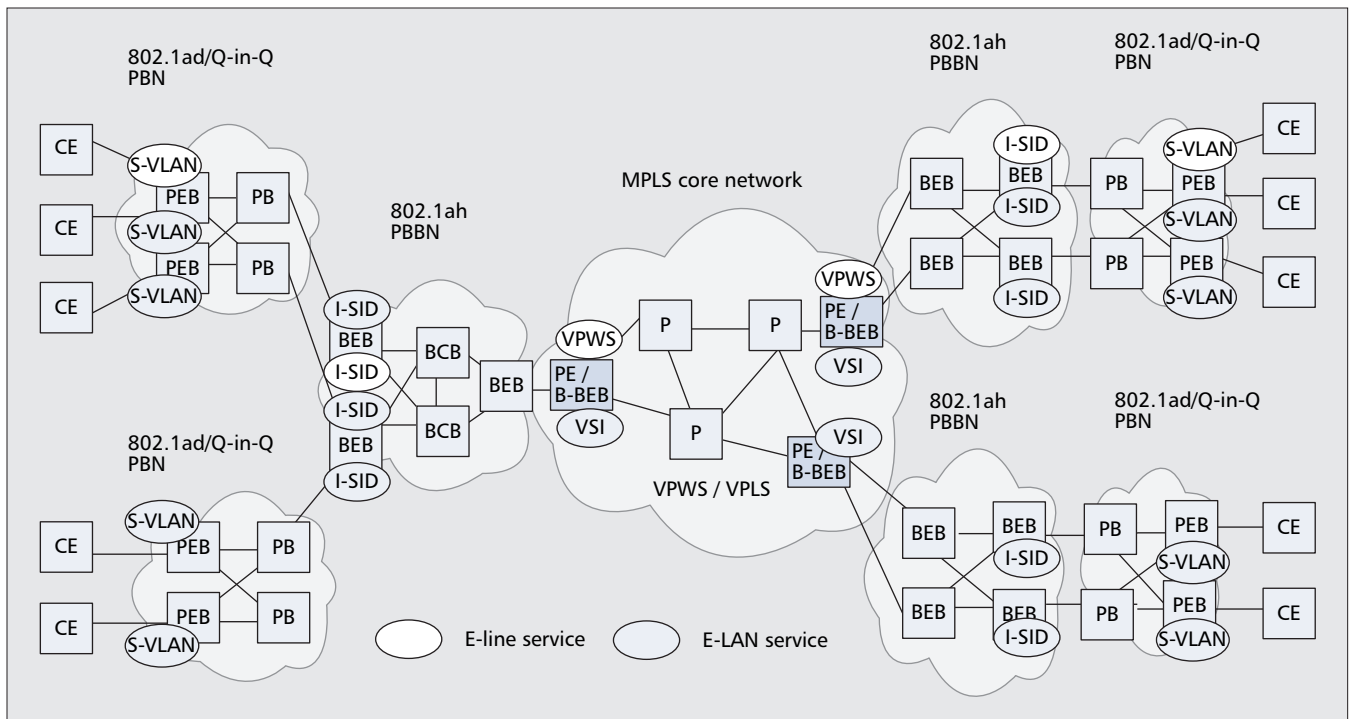


Figure 2. 802.1ad islands in a metro network are interconnected by 802.1ah networks. Ethernet services are extended over an MPLS core network using VPWS and VPLS.

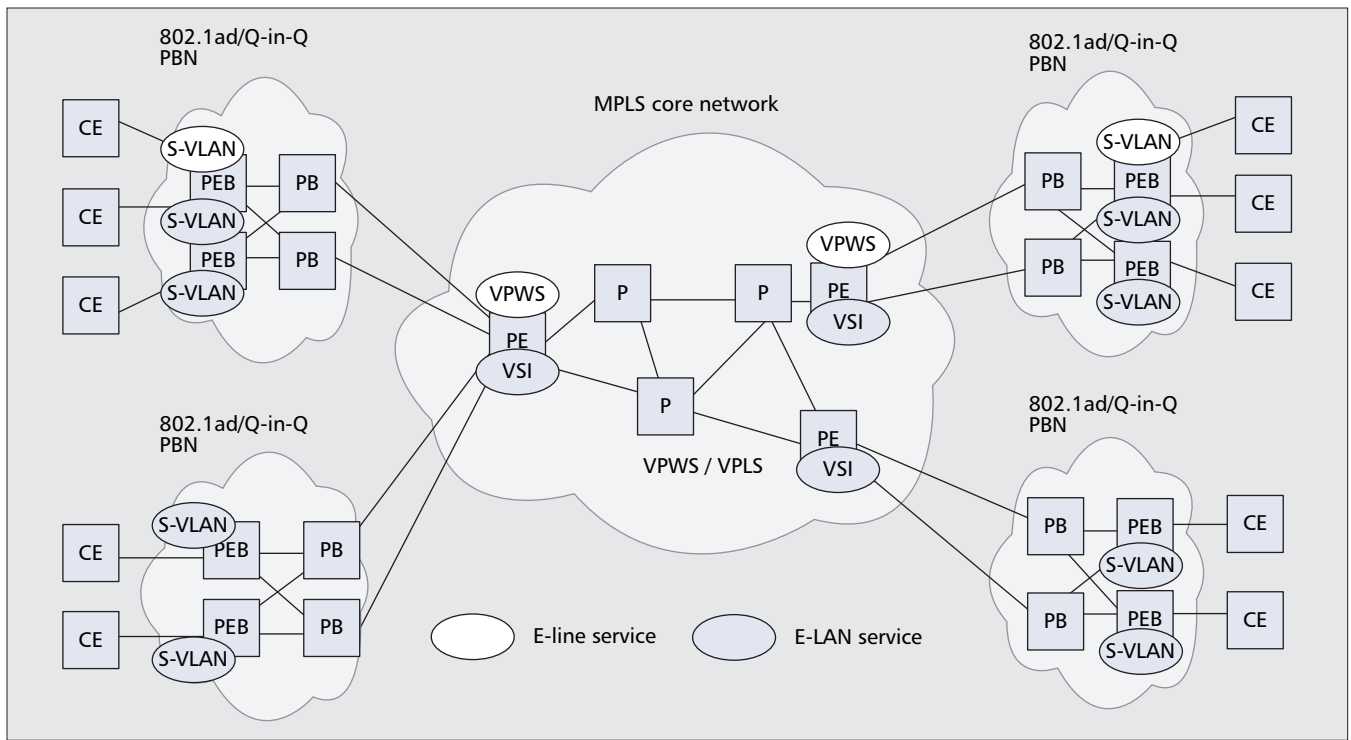
vice instance can be identified by a B-VID for interconnecting PBB islands or an I-SID for interconnecting PBB islands with an instance for a given service [2]. In either case, PBBN access to the PE provides for customer MAC hiding. Multicast, broadcast, and MAC unknown flooding over a VPLS often relies on replication at the PE connected to the source (edge replication). Further optimization for this traffic may rely on multicast capabilities in the MPLS core to build multicast trees among PEs that can be shared by VSIs or dedicated per VSI. Optimization for multicast can rely on Internet Group Management Protocol (IGMP) and protocol independent multicast (PIM) snooping at the PE. For the network depicted in Fig. 2, it is important that resiliency, CoS, and OAM requirements are met using native Ethernet and MPLS mechanisms and their interworking.

Resiliency entails failure detection and switchover actions. In a PBN or PBBN, reroute around node failure relies on STP/RSTP/MSTP to converge and VLAN re-registration to take place. In an MPLS network, reroute around a core router failure relies on Interior Gateway Protocol (IGP) convergence and PE-PE tunnel re-establishment. If Resource Reservation Protocol with Traffic Engineering Extensions (RSVP-TE) is enabled in the MPLS core, reroute may take place faster using fast reroute (FRR) mechanisms. In all cases, link protection may rely on link aggregation based on IEEE 802.3ad. The challenge arises when stitching the network layers together for an end-to-end service across all layers. For instance, if an STP topology changes in a PB island that changes routing for a given I-SID in a PBBN, this must be rapidly reflected in all remote BEBs with an instance for that I-SID. Fault detection and

propagation across layers is important for fast failure recovery.

CoS is identified in a PB, PBB, and on an Ethernet AC to a VSI or VPWS based on the priority code point (PCP) in the outer VLAN tag. Drop priority is identified by the PCP or by the discard eligibility identifier (DEI) in the outer tag. Ethernet frames are mapped to a queue based on the PCP value. A queue can be per port or per VLAN + port, and it is assigned a scheduling behavior. At the MPLS PE, PCP and DEI are mapped to the EXP value in the PW MPLS header. In the case of a uniform tunneling model [10], the PW EXP value is copied to the outer MPLS tunnel header used to transport the PW across the MPLS network. In the E-LSP model [10], EXP is used in the MPLS network to define the CoS that a packet is assigned to and consequently, the per-hop behavior (PHB) at each MPLS node.

Ethernet OAM is based on IEEE 802.1ag and ITU Y.1731 to provide for continuity checks and loopback tests per VLAN (S-VLAN and B-VID), as well as alarm indication signal (AIS) and performance measurements. Ethernet OAM is used to check the liveness of an EVC or a PBBN tunnel and to trigger events (alarms or switchover) when a failure occurs. In the MPLS domain, OAM is based on MPLS mechanisms including virtual circuit connection verification (VCCV) for PWs and MPLS bidirectional forwarding detection (BFD), and IP BFD. MPLS ping [9] and traceroute are also used to provide for LSP ping and trace. Failure notification for PWs also can be based on LDP failure status notification and BGP network layer reachability information withdrawal, as they apply. When relying on failure notification rather than on end-to-end liveness, check to detect an EVC



■ **Figure 3.** 802.1ad islands interconnected by an MPLS network.

failure; fault correlation across the layers is important. For instance, failure of a PW must trigger the generation of AIS on the corresponding AC. Similarly, failure of an AC must trigger the generation of status-down failure notification in LDP for the corresponding PW. At a BEB, the failure of a PBB tunnel carrying multiple EVCs identified by I-SIDs must generate AIS for each affected EVC. Similarly, the failure of an EVC must generate an AIS and remote defect indicator (RDI) for the corresponding I-SID across the PBB network.

SCENARIO 2: VPWS/VPLS INTERCONNECTING 802.1AD ISLANDS

While 802.1ah can be used to aggregate 802.1ad islands to provide for service scalability across the islands, an MPLS network can accomplish the same task using VPWS and VPLS.

Figure 3 depicts a network where 802.1ad (or Q-in-Q) islands (referred to as PBN islands) are interconnected via a MPLS network. E-LAN services are extended using VPLS across the MPLS network, and E-Line services are extended using VPWS. At an MPLS PE, each E-LAN EVC identified by an S-VLAN ID (S-VID) is mapped to a VSI, and each E-Line EVC is mapped to a PW. PEs are interconnected by a mesh of PWs per VSI instance for VPLS services. Resiliency, OAM, and QoS in the PBN islands and MPLS network are similar to what is discussed in Scenario 1 and are not repeated here.

There are several differences between PBB and VPLS/VPWS approaches when used for interconnecting PBN islands:

- PBBN provides for customer MAC hiding in the PBB core. Thus, when a VPLS interconnects PBB islands as in Scenario 1, the customer

MAC addresses are not seen at the MPLS PE. However, it should be noted that a BEB and an MPLS PE are visible to the same customer MAC addresses when they connect to the same PBN islands.

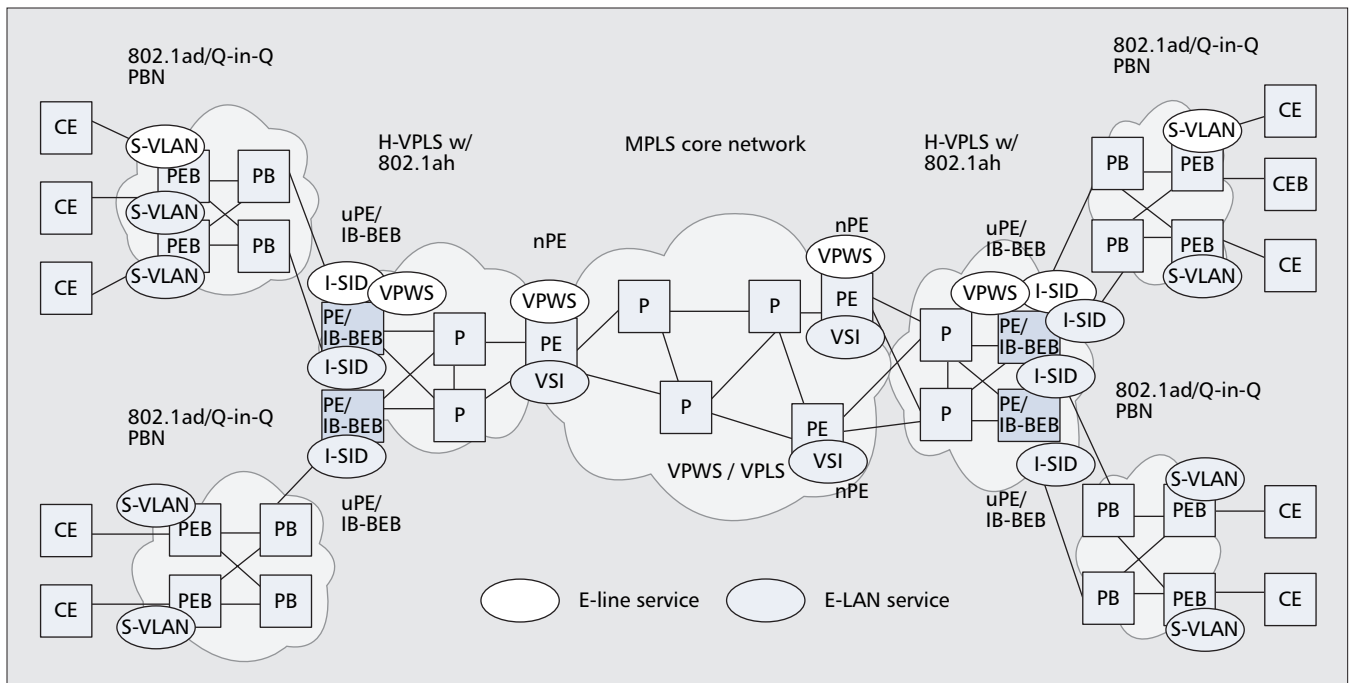
- A PBBN uses the native Ethernet mechanisms for flooding and sending broadcast traffic across the PBBN. Further optimization can be obtained as discussed in Scenario 1. On the other hand, an MPLS PE replicates flooded traffic and broadcast/multicast traffic to every other PE with an instance of the same E-LAN service to which the traffic belongs. Flooding/broadcast/multicast forwarding for VPLS can be optimized by using IP/MPLS multicast mechanisms in the MPLS core as discussed in Scenario 1.

- A PBBN maintains Ethernet OAM mechanisms (Y.1731, 802.1ag) and Ethernet Layer2 control protocols across the PBN and PBBN. PBBN may require the enablement of new L2 control protocols (e.g., MMRP). VPLS/VPWS requires MPLS OAM and interworking between MPLS OAM mechanisms and native Ethernet OAM mechanisms. In addition, an MPLS network relies on the IP control plane for routing and MPLS and PW signaling.

SCENARIO 3: H-VPLS AGGREGATION WITH 802.1AH EXTENSION

This design scenario aims to improve the scalability of VPLSs as aggregation networks with the 802.1ah extension [2].

Figure 4 illustrates this design option of the service provider network supporting Ethernet services. There are also three layers. The access layer is based on 802.1ad or Q-in-Q; the aggregation layer is using H-VPLS technologies with an 802.1ah extension; and the core is MPLS. For



■ **Figure 4.** H-VPLS with 802.1ah extension aggregation networks over an MPLS core.

E-LAN service, it is extended using enhanced VPLS, with the combination of VPLS and 802.1ah at uPE, across the aggregation network and over the MPLS core. The uPEs in the aggregation network are upgraded to PE/IB-BEB to support 802.1ah and MPLS/VPLS. Each E-LAN EVC identified by an S-VID first is mapped to an I-SID; then a single I-SID or a group of I-SIDs is mapped to a spoke PW. For the E-Line service, one could bypass the 802.1ah encapsulation to map the AC directly to a VPWS, using MPLS through the aggregation and core networks.

This design approach has certain similarities and differences when compared with Scenario 1 — native 802.1ah aggregation over MPLS core, and scenario 2 — a pure MPLS/VPLS approach.

Scenario 1 and scenario 3 both use 802.1ah technologies in the aggregation networks to interconnect the 802.1ad or Q-in-Q islands. In scenario 1, the aggregation networks are native Ethernet 802.1ah PBBN; therefore, they use shortest path tree (SPT), Ethernet OAM, and Ethernet failure protection mechanisms. In Scenario 3, the aggregation networks are H-VPLS networks; they use MPLS control plane, MPLS OAM, and MPLS failure recovery mechanisms.

Scenario 2 and scenario 3 both use VPLS in the aggregation networks. The difference is that scenario 3 uses 802.1ah scaling enhancement in the VPLS aggregation networks and across the core, whereas scenario 2 uses regular VPLS/VPWS. The uPE in scenario 3 functions as an 802.1ah IB-BEB when facing access networks, and it functions as an MPLS/VPLS PE when facing the core of the aggregation network.

This scenario takes the advantages of the technology maturity of MPLS and the scalability of 802.1ah. The scalability improvements are

- Service instances: from 4 K to 16 M by using I-SID

- MAC addresses reduction — by adding 802.1ah B-MAC encapsulation in uPE, the nPE is required to learn only the MAC addresses of the uPEs, not all customer MAC addresses
- PW reduction in the core because of H-VPLS

When using n:1 for I-SID to B-VLAN mapping and per B-VLAN VPLS instance mapping, the number of PWs and VSIs can be significantly reduced.

CONCLUSIONS AND FUTURE WORK

In this article, we reviewed the evolution of Ethernet technologies and the recent developments in Carrier Ethernet services. We centered our discussion on service requirements and on three deployment scenarios for building scalable and reliable Carrier Ethernet networks and services.

There are several technologies available today to provide Ethernet services, with different trade-offs among them. When deciding which technology to deploy, in which part of the network, and for what Ethernet services, one must consider many factors, both technical and operational. Some of the main factors are:

- Capital expenses (CAPEX) and operational expenses (OPEX)
- Whether the deployment is a green field network or is expanding existing networks
- Existing technologies already deployed in the network for Ethernet and other services
- Operational staff experience
- The portfolio of desired services (E-Line, E-LAN, multicast) and technology-based optimality and ease of operation
- The availability and maturity of the desired technologies
- The preferred network and service manage-

When deciding which technology to deploy, in which part of the network, and for what Ethernet services, one must consider many factors, both technical and operational.

- ability paradigm, which may vary from one operator to another
- The operation support systems (OSS) to be used
- Future work in Carrier Ethernet services may address the following critical areas:
- **Scalability.** This includes scaling improvement for MPLS technologies, new developments in 802.1ah, interworking of VPWS/VPLS with 802.1ah, and systems hardware and software scaling improvements.
 - **New standards and technology evolution to make Ethernet qualify for carrier class.** These include: QoS, OAM, high availability (HA), traffic engineering (TE), and FRR, some of which are not supported by traditional Ethernet technologies.
 - **Network management systems or control plane to meet different customers' requirements.** There are currently several efforts in IETF and IEEE to provide the control plane for replacing STP in Ethernet. A unified solution, if possible to achieve, will contribute to cost savings and interoperability.
 - **CAPEX and OPEX analysis.** The attraction of Ethernet is its simplicity and economics. Though challenging, it is desirable to maintain that appeal when adding the additional features.
 - **Additional Carrier Ethernet services.** In addition to E-Line and E-LAN, E-Tree services are being developed. Additional services that offer fast provisioning, flexibility, and reliability will emerge to meet end users' requirements.

REFERENCES

- [1] IEEE Draft 3.7, "IEEE 802.1ah — Provider Backbone Bridges," Aug. 2007, work in progress.
- [2] A. Sajassi *et al.*, "VPLS Interoperability with Provider Backbone Bridges," Mar. 2007, Nov. 2007; work in progress.
- [3] M. Lasserre and V. Kompella, "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling," RFC 4762, Jan. 2007.
- [4] K. Kompella and Y. Rekhter, "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling," RFC 4761, Jan. 2007.
- [5] L. Martini *et al.*, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)," RFC 4447, Apr. 2006.
- [6] Metro Ethernet Forum, "MEF 6 — Ethernet Services Definitions — Phase 1," June 2004.
- [7] E. Rosen *et al.*, "Provisioning, Autodiscovery, and Signaling in L2VPNs," RFC ED queue, May 2006.

- [8] T. Nadeau and C. Pignataro, "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires," RFC ED queue, July 2007.
- [9] K. Kompella and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures," RFC 4379, Feb. 2006.
- [10] F. Le Faucheur, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services," RFC 3270, May 2002.

BIOGRAPHIES

LUYUAN FANG (lufang@cisco.com) received her Ph.D. in computer science from Flinders University of South Australia in 1991. She joined Cisco in 2006 as a product manager with responsibilities in the service provider wireline segment, including Ethernet technologies and solutions. Prior to joining Cisco, she worked as a lead network architect for AT&T in IP/MPLS VPN design and deployment. Prior to joining AT&T in 1998, she worked for Racal Datacom, Nortel, and Telstra. She is an active contributor in IETF. She has co-authored seven RFCs and has several active Internet drafts in the MPLS, L2/L3 VPN, and Traffic Engineering Working Groups. She is a co-editor of the *IEEE Communications Magazine* Feature Topic on VPN technologies. She has served as a technical committee member and a frequent speaker for several prominent MPLS and carrier Ethernet conferences worldwide. She has over 70 technical publications.

NABIL BITAR holds B.S., M.S., and Ph.D. degrees in electrical engineering from Boston University, Massachusetts. He is currently a principal member of technical staff at Verizon in the Packet Network Architecture department. He leads the network architecture for metro packet services, including Ethernet services, metro packet transport, video packet transport, IPVPN, VoIP routing, and IP services over FTTP. Prior to joining Verizon in 2004, he worked at Ascend/Lucent as a system architect for ATM and IP-MPLS services (forwarding, traffic management, signaling, and routing). Prior to Ascend/Lucent he worked at GTE Laboratories for five years on wireless AIN, ADSL architecture, IP IntServ and DiffServ, MPLS, VoIP, and traffic management. He is a regular contributor to the IETF and IP/MPLS Forum.

RAYMOD ZHANG is director of Advanced Engineering, BT Design, and a principal architect of BT's 21CN Networks. His main interests are in areas of large-scale backbone routing, traffic engineering, performance and traffic statistical analysis, information architecture and service modeling, software driven computing, MPLS, and Ethernet related technologies. He participates in several IETF drafts relating to MPLS, BGP-based MPLS VPNs, inter-AS TE, and more recently PCE-based work. He received his M.S. degree in electrical engineering from City University of New York.

MICHAEL TAYLOR is a network planner/designer in the metro Ethernet space at AT&T Laboratories, San Ramon, California. He has over 20 years of networking experience in telecommunications. Part of this was spent working in Silicon Valley with startups and Cisco Systems. Additionally, he has held positions in management, engineering support, engineering quality management, and services design. He received his M.B.A. from St Mary's College, Moraga, California; he also holds a B.B.A. from the University of Texas, Arlington.