Service Provider Choices for Ethernet Services

Andrew G. Malis
Verizon Communications
andrew.g.malis@verizon.com

Future-Net Expo
14 April 2008

Introduction

- Public Ethernet services are exploding in popularity
- External Ethernet interface to the customer may not necessarily mean "Ethernet inside"
  - Providers have a choice of mechanisms to us
  - Ethernet switching not always the best choice for public Ethernet services
    - Scaling limitations that can limit the scope of an Ethernet service
    - Functional limitations that can restrict Service Level Agreements
  - This talk discusses technology available to providers to support scalable Ethernet services with SLAs

Why Are Ethernet Services Popular?

- Ubiquity and low cost of Ethernet interfaces in customer equipment, universal experience with Ethernet in LANs, and perceived simplicity
- Successful marketing of the “Ethernet” brand by vendors, IEEE, MEF, and others
  - Little resemblance with original DIX Ethernet specifications, from physical layer on up (e.g., today’s Ethernet is mostly point-to-point or ring-based rather than CSMA-CD at the physical layer)
  - Most everything has changed except for the basic frame format – and jumbograms change even that
- Favorable pricing by service providers

“Enterprise-Class” Ethernet Limitations

- **“Enterprise-class” Ethernet switching has shortcomings as a basic for public Ethernet services**
  - **Few features for high availability in protocols or LAN-based equipment**
  - **Scaling limits on MAC addresses, VLAN IDs, and spanning tree topology limit the size of native Ethernet networks**
  - **Spanning tree routing may take seconds to (occasionally) minutes to re-converge**
- **Early bleeding-edge Ethernet providers found the hard way that enterprise-class Ethernet cannot naively be deployed for reliable carrier services**

Emergence of “Carrier Ethernet”

- Limitations in enterprise-class Ethernet have led to the development of “Carrier Ethernet”
- Meant to address unique requirements for carrier Ethernet services
  - Scaling to support a large number of customers
  - Scaling to support large numbers of switches and customer interfaces
  - Support both point-to-point (E-Line) and multipoint (E-LAN and E-Tree) services
  - Support for both port-based and VLAN-based services
  - Support for QoS other than best-effort to support QoS-based SLAs
  - Sub-second outage restoration and routing convergence to support availability SLAs
  - Policing and shaping to support sub-rate services (e.g., 200 Mbps service on a physical GigE interface)

MEF Carrier Ethernet Service Definitions

Connectivity Model	Port-Based (All to One Bundling)	VLAN-Based (EVC identified by VLAN ID)
E-Line (point-to-point EVC)	**Ethernet Private Line (EPL)**	**Ethernet Virtual Private Line (EVPL)**
E-LAN (multipoint-to-multipoint EVC)	**Ethernet Private LAN (EP-LAN)**	**Ethernet Virtual Private LAN (EVP-LAN)**
E-Tree (rooted multipoint EVC)	**Ethernet Private Tree (EP-Tree)**	**Ethernet Virtual Private Tree (EVP-Tree)**

- **Three service types based on the three Ethernet Virtual Connection (EVC) types**
- **Two "UNI Types" determine whether services are 'private' or 'virtual'**
  - **Port-based (All to One Bundling) → single EVC (transparency, but uses an entire port per service)**
  - **VLAN-based → 'N' EVCs per UNI (not as transparent, but multiple services per port)**
- **Services are defined by combination of connectivity model and 'UNI Type'**
- **Also Ethernet-based access services to Layer 3 VPNs or dedicated Internet access**

E-Line Services

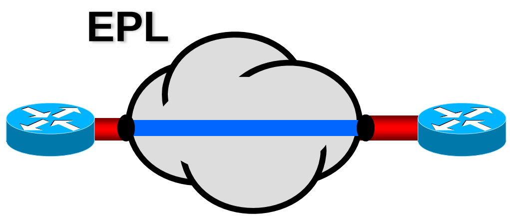

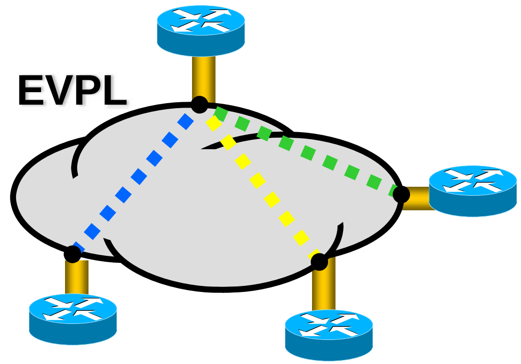

Key Characteristics

- **EPL**
  - p2p service, transparent, single service, uses a port on CE for each service
  - Ideal for customers wanting a ‘private line’ like service model
- **EVPL**
  - p2p service, not as transparent, multiple services on a UNI
  - Ideal for customers wanting a ‘frame relay’ like service model

All to One Bundled UNIs

Service Multiplexed UNIs

E-LAN Services

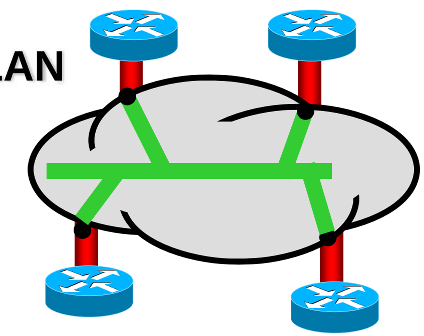

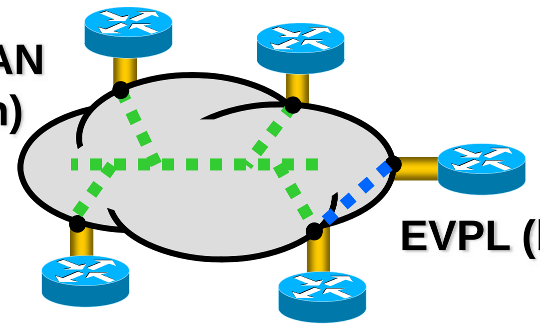

Key Characteristics

- EP-LAN
  - mp2mp service, transparent, single service, uses a port on CE for each service
  - Ideal for customers wanting a ‘Transparent LAN’ like service model
- EVP-LAN
  - mp2mp service, not as transparent, multiple services on a UNI
  - Ideal for customers wanting a multipoint service for LAN interconnect and one or additional services on one or more UNIs

All to One Bundled UNIs
Service Multiplexed UNIs

E-Tree Services

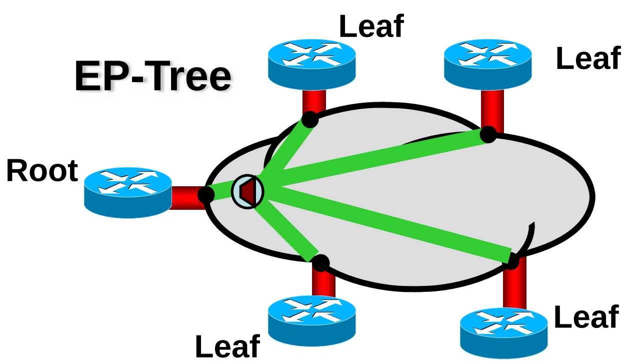

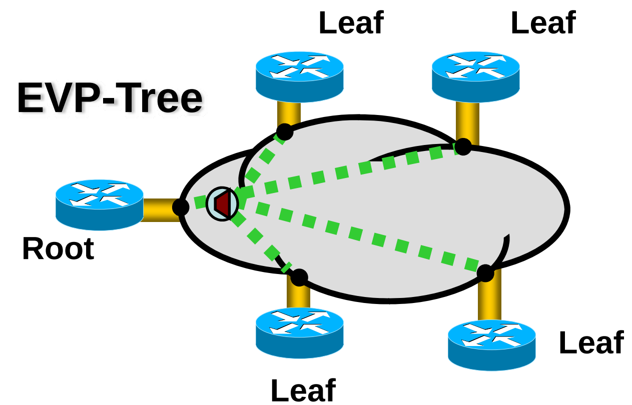

Key Characteristics

- EP-Tree
  - Rooted mp service, transparent, single service, uses a port on CE for each service
  - Ideal for customers wanting a ‘broadcast’ like service model from one or more roots to many leaves; allows for some upstream b/w
- EVP-Tree
  - Rooted mp service, not as transparent, multiple services on a UNI
  - Ideal for customers wanting one or more rooted multipoint services and other services on a UNI (e.g., market data feed, ISP, mobile backhaul, distance learning)

All to One Bundled UNIs
Service Multiplexed UNIs

Ethernet-Based Access

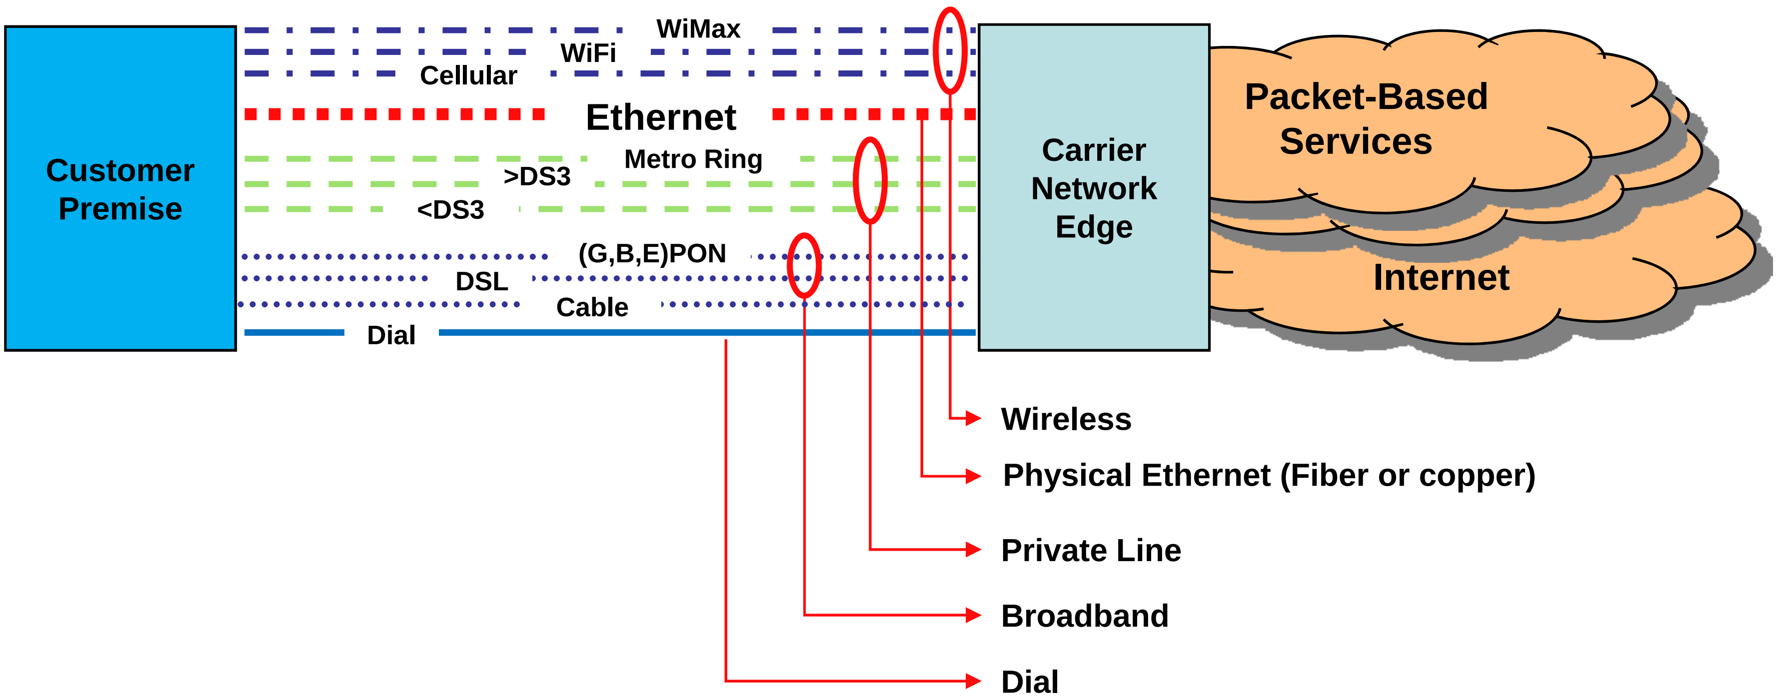

- **All of the access solutions can present an Ethernet handoff**
- **Does not mean Ethernet service – source of confusion in the market**

US Ethernet Forecast

Ethernet Forecast Components ($M)

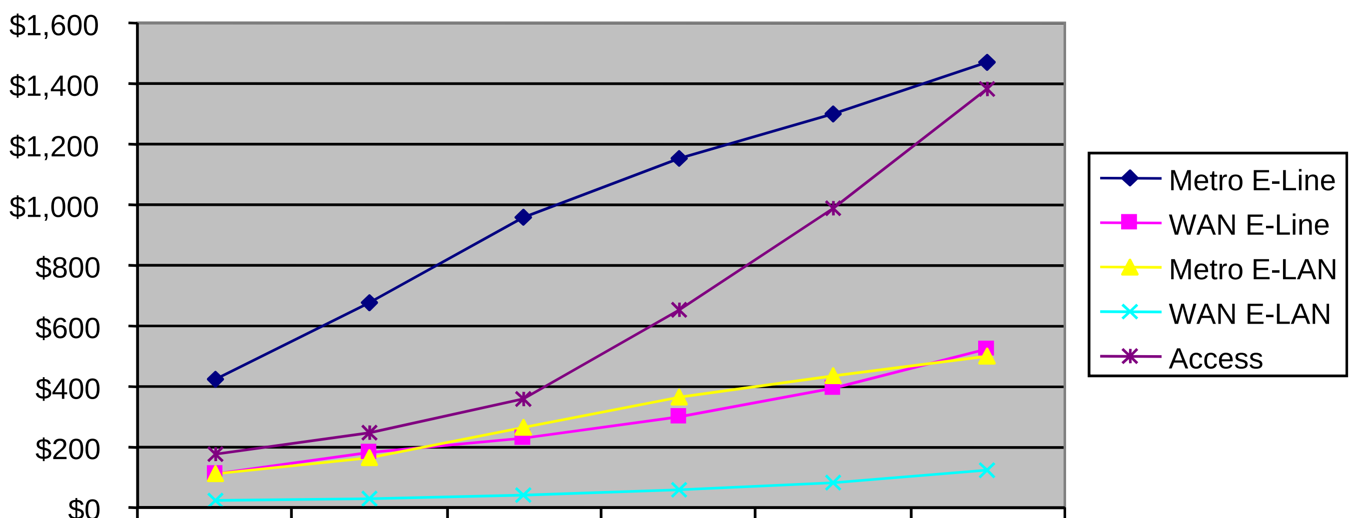

	2006	2007	2008	2009	2010	2011
Metro E-Line	$423	$675	$957	$1,156	$1,300	$1,473
WAN E-Line	$113	$179	$231	$300	$394	$522
Metro E-LAN	$109	$162	$263	$366	$438	$497
WAN E-LAN	$24	$32	$42	$57	$83	$122
Access	$175	$249	$360	$651	$990	$1,380
Total	$845	$1,297	$1,854	$2,529	$3,204	$3,993

Source: IDC

How Carrier Ethernet Requirements Are Being Addressed

- IEEE 802 has completed or is working on protocol extensions such as provider bridging, MSTP, RSTP, PBB, PBB-TE, and OAM
- MEF is creating Carrier Ethernet interface and service specifications and addressing issues such as management, resiliency, and QoS
- IETF and IP/MPLS Forum have added Ethernet transport and interworking capabilities to IP/MPLS routers, such as pseudowires ,VPLS, and multi-service interworking
- Ethernet equipment vendors are adding high availability features to Ethernet switches, such as redundant power supplies, fans, and crossbar switches, and high availability software features such as nonstop software upgrades
- Optical and transport equipment vendors are incorporating Ethernet interfaces and mappings to enable packet-based Ethernet transport

IEEE 802 Key Recent/Current Carrier Ethernet-Related Projects

Specification	Key Focus
Provider Bridging (802.1ad-2005)	•Standardize Q-in-Q (Ethertype: 88A8); Priority Code Point + Drop Eligibility Indicator bit
Connectivity Fault Management (802.1ag)	•Fault Management for EVCs and links •Continuity Check, Loopback and Link Trace •Extended by ITU SG13 (Y.1731) to include AIS, PM for point-to-point EVCs
Multiple Registration Protocol (802.1ak)	•Automated VLAN (Multicast Multiple Registration Protocol) and MAC (Multiple MAC Registration Protocol) management
Provider Backbone Bridges (802.1ah)	•MAC-in-MAC tunneling, solves VLAN and MAC scaling issues; multipoint and point-to-point based tunnels
PBB-TE (802.1Qay)	•Provisioned TE-paths for p2p tunnels; 1:1 path protection (50 ms)
Frame format expansion (802.3as)	•2000 byte MTU (allows for sandwich of protocols – IPSec, MPLS, PBB; retains 1500 byte packet payload)
Shortest Path Bridging (PLSB, 802.1aq)	•Control plane for Ethernet PBBNs, based on IS-IS (Provider Link State Bridging)

MEF Key Current Carrier Ethernet-Related Projects

Specification	Key Focus
UNI Type 2 Implementation Agreement (IA)	•Ethernet Local Management Interface, Service OAM •Link OAM and protection (link aggregation) •Two types: 2.1 (scaled down), and 2.2 (full feature set)
Ethernet Services Definitions, Phase 2	•3 service types (E-Line, E-LAN and E-Tree); 6 services •UNI, EVC service attribute requirements; use cases
Service OAM IA	•Fault Management (standard and tunnel access services; UNI, E-NNI) •Performance Management (significant work focusing on implementation)
E-NNI Phase 1	•External Network-to-Network Interface •S-tag; Link protection via Link Aggregation •Standard and tunnel access services (no E-Tree in Ph 1) •Virtual UNI •Management (service, link, tunnel)
CoS	•Basic 2, 3 and 4-class relativistic models; map typical apps to CoS; CoS by service type?; stretch goal → quantify performance)
Abstract Test Suites	•UNI Type 1, UNI Type 2 (per protocol), E-NNI

IETF Ethernet Services Support

- Point-to-point pseudowires (PWs) to carry layer two frames, including Ethernet, over IP/MPLS networks
- Extremely popular, implemented by most every router vendor and in wide use by service providers world-wide
- Extends the MPLS LDP protocol to signal pseudowire establishment
- IETF extended PWs to a multipoint Ethernet service, VPLS (Virtual Private LAN Service)
- IETF also standardized PWs over L2TPv3 for those few service providers not using MPLS
- IETF's CCAMP WG is just beginning work on Generalized MPLS-based signaling for two Ethernet-based applications
  - To automate traffic engineering path computation and provisioning for IEEE 802.1Qay (PBB-TE) (also known as GMPLS Ethernet Label Switching or GELS)
  - End-to-end service signaling for MEF-defined carrier Ethernet service interfaces (may be over non-Ethernet networks)

IP/MPLS Forum Ethernet Services Support

- **Extended IETF PWs to support non-similar endpoint interworking**
  - Supports point-to-point Ethernet-to-Frame Relay, Ethernet-to-ATM, and ATM-to-Frame Relay interworking over MPLS PWs
  - Very useful for multiservice convergence, and to support customers with a variety of access methods
  - Can support applications such as hub location with GigE access, and low-speed Frame Relay spokes
  - Supports interworking of IP packets via ARP Mediation, and bridged services by interworking native Ethernet with Ethernet frames encapsulated by FR or ATM
  - Can also support VPLS endpoints with FR or ATM-attached customer equipment

Constructing Carrier Ethernet Networks

- Given standardization work in various venues, service providers have a choice of technologies to use to instantiate Ethernet services
  - Ethernet switch-based networks
  - Router-based networks
  - Optical switch-based networks
  - Some combination of the above
- This choice is further complicated by the fact that standards work is still in progress
  - Pre-standard implementations in vendor equipment
  - Different vendors make different choices of what to implement, since vendors are resource-constrained
  - Technology choices may be constrained by vendor choices, or vice versa

Further Constraints in Constructing Carrier Ethernet Networks

- Full multivendor interoperability usually requires maturity in standards and in vendor implementations
  - Providers are usually reticent to deploy technologies that may lock them into a particular vendor
  - But at the same time, they don't want to be forced into using "old" technology that may not meet their ongoing requirements
- Providers may have additional constraints on technologies that they may deploy
  - Existing management systems and personnel
    - Personnel may need to be retrained
    - Management systems may need to be upgraded
  - Evolution or revolution from existing networks
    - Revolution may require "fork-lift" upgrade

Evolving and Scaling Ethernet Services

- A typical "early" public Ethernet service provider probably uses Ethernet switches and Q-in-Q for customer separation
- Typical end user services are
  - Ethernet Private LAN (EP-LAN)
  - Ethernet Virtual Private LAN (EVP-LAN)
  - Ethernet Private Line (EPL)
  - Ethernet Virtual Private Line (EVPL)
  - Each of these services requires the use of a provider VLAN tag
- As the service becomes successful, the provider will encounter the usual Ethernet scaling limitations
  - MAC address scaling
  - VLAN tag scaling (4K customer limit)
  - Switching capacity limits

Typical "Early" Ethernet Service Network

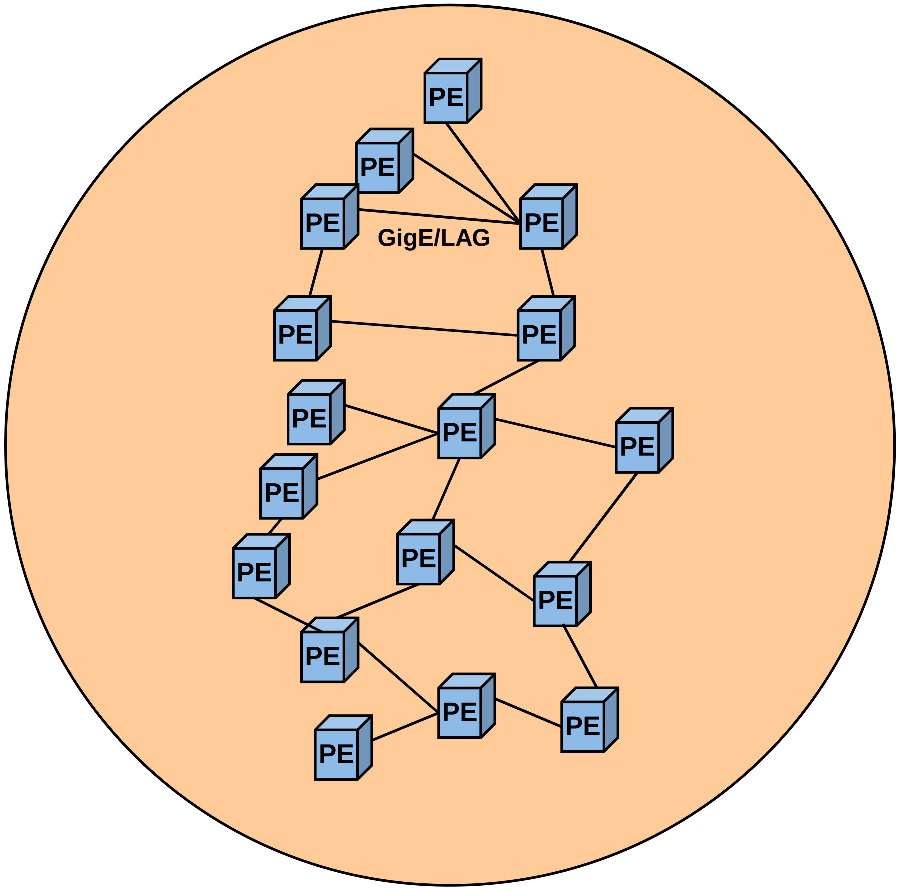

PE – Provider Edge Switch

- Characterized by organic growth driven by customer location
- All switches are "edge switches"
- May be some number of redundant links
- 802.3ad Link Aggregation may also be used for resiliency or for additional BW between switches
- Flat network with spanning tree routing
  - Network diameter is limited, often to metro scope

Need To Evolve Service To Next Step

- Support more customers, beyond 4K VLAN tag limit and MAC address limits in flat networks
- Increase network reliability and feature support
  - Ethernet OAM (Y.1731, IEEE 802.1ag) for failure detection and notification
  - Dual-homing customers and switches to protect against single switch failures
  - In-service SW upgrades, non-stop control plane and forwarding
  - Support for multiple classes of services
    - Allows writing QoS SLAs, supporting multimedia and real-time applications

Introducing Hierarchy For Scaling

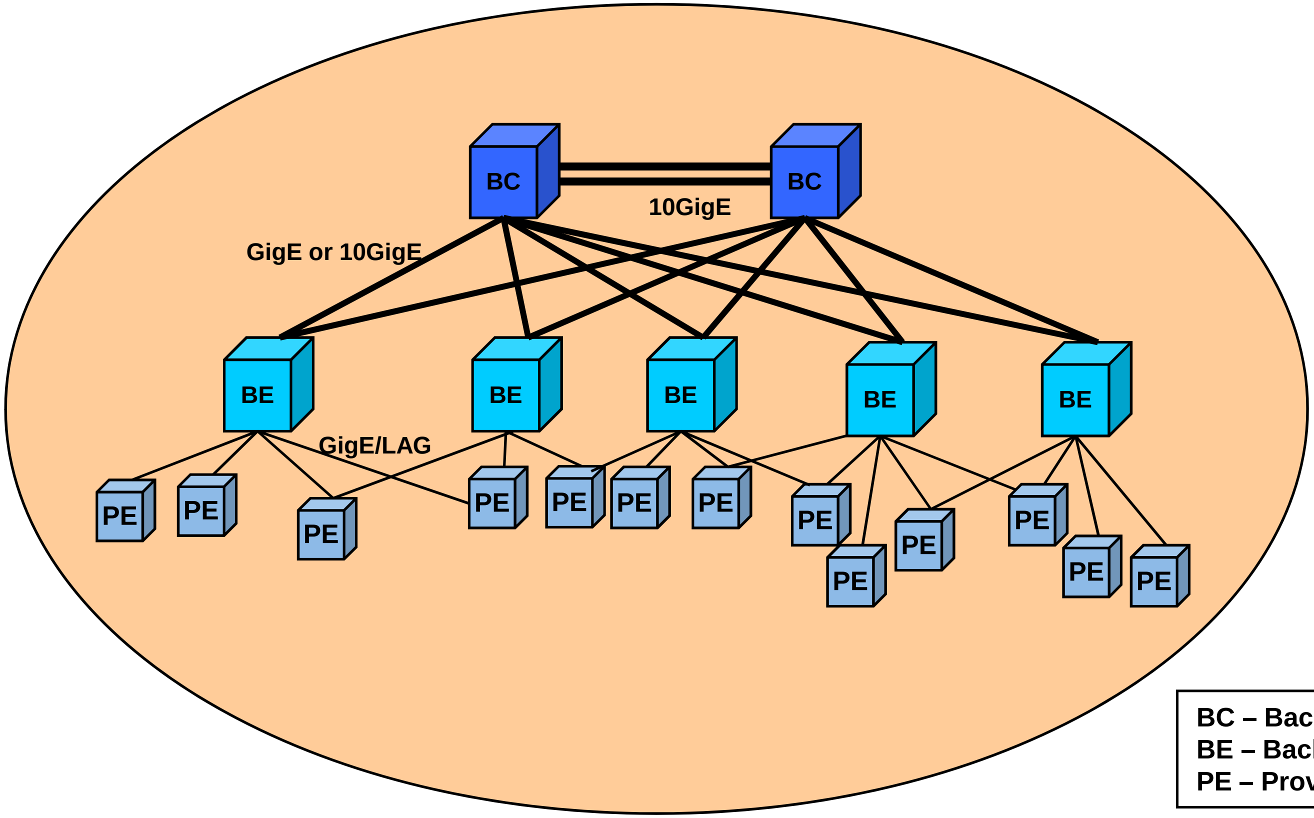

BC – Backbone Core Switch
BE – Backbone Edge Switch
PE – Provider Edge Switch

Physical Hierarchy Helps

- Physical scaling addresses
  - Switch capacity limits
  - Spanning tree diameter
  - Improves resiliency in core backbone network
- Allows optimizing equipment choices at each level of the hierarchy
  - Higher switching capacity in the core
  - GigE interface density for backbone edge
  - End customer interface cost and diversity at provider edge
- BUT … still have VLAN and MAC address scaling limitations

Also Need Tunneling For Scaling

- Tunneling enables scaling by:
  - Restricting customer MAC addresses to switches where they are used
  - Replaces provider VLAN tags with much larger provider service identifiers/labels
- Two primary tunneling alternatives
  - IEEE 802.1ah Provider Backbone Bridging (PBB)
  - MPLS-based VPWS for EPL and EVPL, and VPLS for EP-LAN and EVP-LAN services

Quick PBB Introduction

- Primary purpose is scaling provider Ethernet backbone networks
- MAC-in-MAC: provides tunneling for customer MAC header (inner MAC) in SP bridges' MAC header (outer MAC)
  - Hides customer MAC addresses from service provider switches
  - Reduction in MAC address table in the core – contains only bridge MAC addresses as opposed to customer MAC addresses
- B-VID: identifies tunnels between provider backbone bridges
- I-SID: represent a service instance in a B-VID carried in the 802.1ah header
  - A service instance would typically be a customer E-Line or E-LAN EVC
  - 24 bits – provides the capability of instantiating up to 16 million service instances in a metro network
  - Enables meeting market demand for large number of EVCs – addresses today's limit of 4094 maximum EVCs in a metro
- Standardization largely complete in IEEE, pre-standard implementations are being shipped

MPLS Strengths and Weaknesses for Ethernet Services

- **MPLS Strengths**
  - Allows converged infrastructure based on MPLS
  - MPLS and VPLS are widely available and interoperable
  - Traffic engineering allows optimization of backbone use
  - Allows the use of non-Ethernet trunking
- **MPLS Weaknesses**
  - Need to translate between Ethernet and MPLS OAM
  - May need to retrain operations staff of an existing Ethernet service
  - May require a fork-lift upgrade of an existing Ethernet service
  - Edge replication for multicast, broadcast and flooding – solution is being worked on in the IETF L2VPN working group (see draft-ietf-l2vpn-vpls-mcast-03.txt and related drafts)
  - Concerns about MAC address scaling in H-VPLS
  - Concerns about full-mesh tunnel scaling in large VPLS networks

PBB Strengths and Weaknesses for Ethernet Services

- **PBB strengths**
  - Closest to existing architecture for existing Ethernet networks, least amount of disruption during deployment
  - Minimal need to retrain operations staff
  - Most efficient for multicast support
- **PBB weaknesses**
  - Still pre-standard, need to wait for standards-based and interoperable implementations
  - Link-state routing coming in the future, but currently depends on spanning tree routing
  - May need PBB-TE in future to optimize backbone utilization

VPLS-PBB Interoperation

- **One possibility to take advantage of both MPLS and PBB strengths is a combined approach using a VPLS core with PBB access for edge scaling**
  - **See draft-sajassi-l2vpn-vpls-pbb-interop-02.txt for details**
  - **Requires new pseudowire type proposed in draft-martini-pwe3-802.1ah-pw-01.txt**
- **Work is very preliminary – still individual contributor drafts at this time**

The Answer?

- **At this point in time, there is no one "right" answer**
  - **Two major toolsets available, with possible future interworking**
  - **Different providers may reach different conclusions on which direction to take based upon their particular requirements and current architecture**
- **But the good news is that both toolsets have promise for scaling Ethernet services, and may interwork in the future**

Thank You!

Andrew G. Malis
Verizon Communications
andrew.g.malis@verizon.com