Developing Standards for Metro Ethernet Networks

Stephen Haddock
shaddock@extremenetworks.com
Chief Technology Officer

Agenda

- Metro Ethernet Networks
- Metro Ethernet Forum
  - Services Model and Definitions
  - Traffic Management
- IEEE 802.1ad Provider Bridging Standard
- Scaling Metro Ethernet Networks
  - MPLS L2VPNs
- New Standards Initiatives
  - IEEE 802.1ag Connectivity and Fault Management
  - IEEE 802.1ah Provider Backbone Bridging

Appeal of Ethernet for Metro Services

- Packet Optimized / High Bandwidth
  - Data traffic exceeded voice traffic in 2000, and continues to grow at a much faster rate.
  - TDM is inefficient for transport of packetized data.
- L2/L3 switches enables a rich and flexible service offering.
- The “end-points” are Ethernet
  - An all-Ethernet network architecture avoids additional layers that add complexity
- Fast and flexible provisioning
  - Fine grain bandwidth increments can be provisioned remotely
  - Widely available, well understood technology

Metro Network Overview

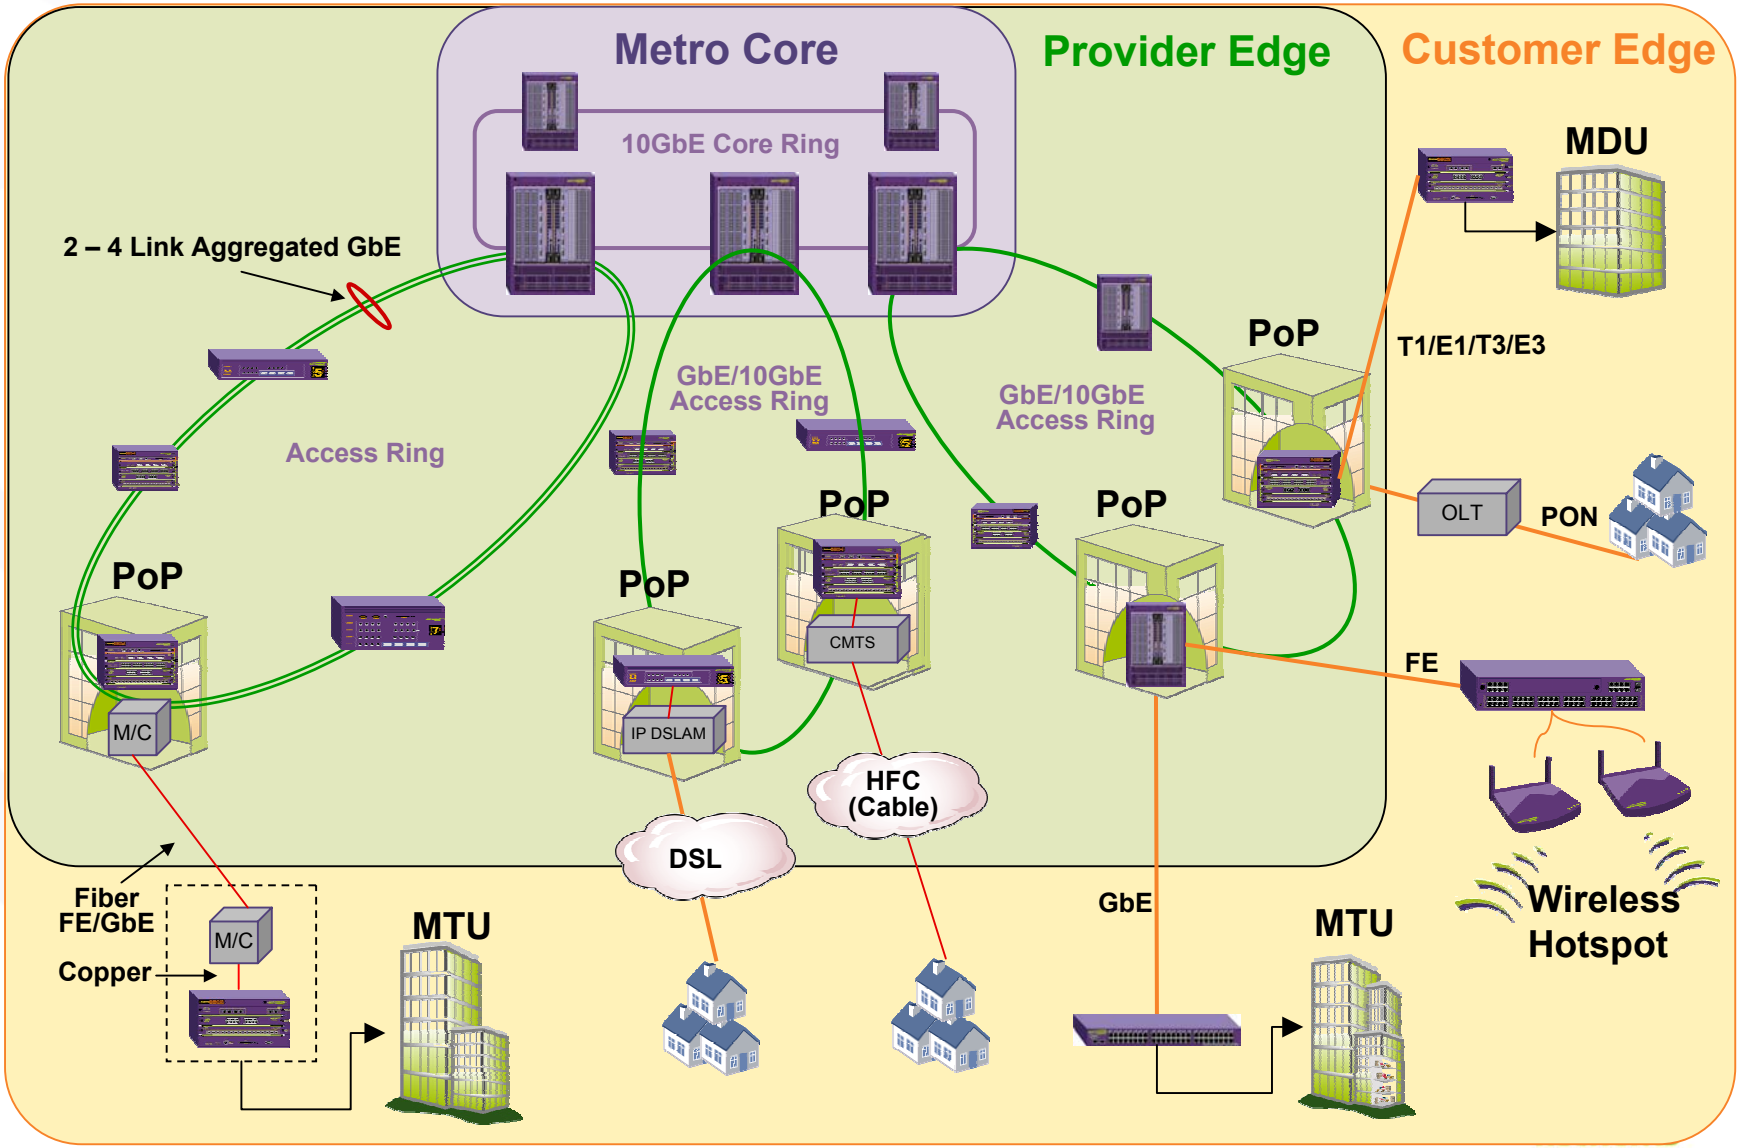

Metro Ethernet Forum

MEF Positioning Statement

- Mission

  "Accelerate worldwide adoption of carrier class Ethernet networks and Services"

- Objectives

  1. Build consensus and unit service proviedres, equipment vendors and end-customers on Ethernet service definition, technical specifications and interoperability.
  2. Facilitate implementation of existing and new standards, Ethernet service definition, test procedures and technical specifications of the MEF to allow delivery of Ethernet services and make Ethernet-based metro networks carrier-class.

MEF Work Items

MEF Technical Committee Work Dashboard				
Service Area	Protocol and Transport Area	Architecture Area	Management Area	Test and Measurement Area
MEF 1 - Ethernet Services Model (TS)	MEF 2 - Protection Framework and Requirements (TS)	MEF 4 – Metro Ethernet Network Architecture Framework Part 1 (TS)	MEF 7 - EMS-NMS Information Model (TS)	MEF 9 - Abstract Test Suite for Ethernet Services at the UNI (TP)
MEF 6 - Ethernet Services Definitions (TS)	Transport Multiplexing Function (formerly EMF) (TS)	UNI Framework & Requirements (TS)	Service OAM Requirements and Framework (TS)	
MEF 5 - Ethernet Traffic Management (TS)	QoS Framework (TS)	UNI Type 1 (IA)	Ethernet Service OAM (TS)	
MEF 3 - Circuit Emulation Service Requirements (TS)	Ethernet over SONET (IA)	External NNI (E-NNI) (TS)	EMS Requirements (TS)	
Services Phase 1 Unified Document	MPLS Protection (IA)	Eth-Layer Architecture Framework (TS)	Ethernet Performance Monitoring (TS)	
PDH Emulation Services (IA)	Frame Relay Service Interworking Function (TS)	UNI Type II (Ethernet LMI) (TS)		
MEF 8 - CES over MENs (IA)		MPLS-based NI-NNI		

TS – Technical Specification
IA – Implementation Agreement
TP – Test Procedure

Legend
- Approved Draft
- Sent to Straw Ballot
- Sent to Letter Ballot
- Approved

METRO Ethernet Forum

extreme networks

MEF Phase I Service Documents

- Phase I consists of 3 technical specifications
  - **Ethernet Services Model** (MEF 1 standard)
    - Defines Ethernet service building blocks (service attributes)
    - Defines a framework describing how to build an Ethernet service
      - does not define Ethernet services
  - **Ethernet Services Definitions** (MEF 6 standard)
    - Defines how to apply the ESM building blocks to create services
    - Defines Ethernet Line (E-Line) and Ethernet LAN (E-LAN) service types and instances of them:
      - Private Line, Virtual Private Line, Internet Access, TLS
  - **Ethernet Traffic Management** (MEF 5 standard)
    - Defines traffic management and service performance requirements to create CoS-based SLAs

Ethernet Service – Basic Model

- Customer Equipment (CE) attaches to UNI
- CE can be
  - router
  - IEEE 802.1Q bridge (switch)
- UNI (User Network Interface)
  - Standard IEEE 802.3 Ethernet PHY and MAC
    - 10Mbps, 100Mbps, 1Gbps or 10Gbps
    - Full Duplex
  - Untagged or VLAN-tagged Ethernet Frames
- Metro Ethernet Network (MEN)
  - May use different transport and service delivery technologies
    - IEEE Provider Bridging (Q-in-Q), MPLS L2VPN, Provider Backbone Bridging (MAC-in-MAC), Ethernet over SONET/SDH, WDM

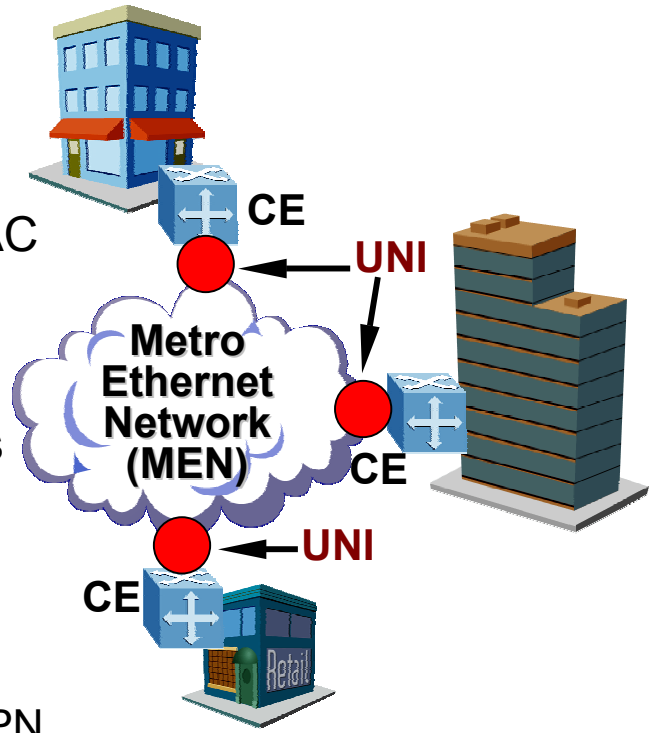

E-Line and E-LAN Service Types

- E-Line Service used to create
  - Private Line Services
  - Ethernet Internet Access
  - Point-to-Point VPNs

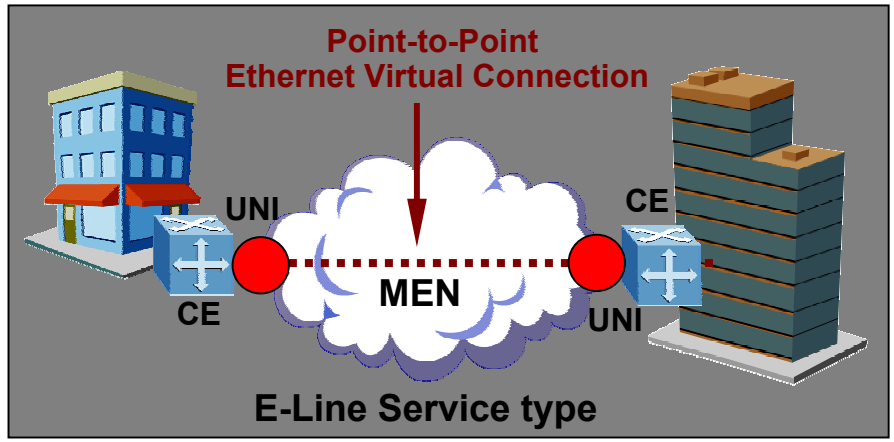

- E-LAN Service used to create
  - Multipoint VPNs
  - Transparent LAN Service

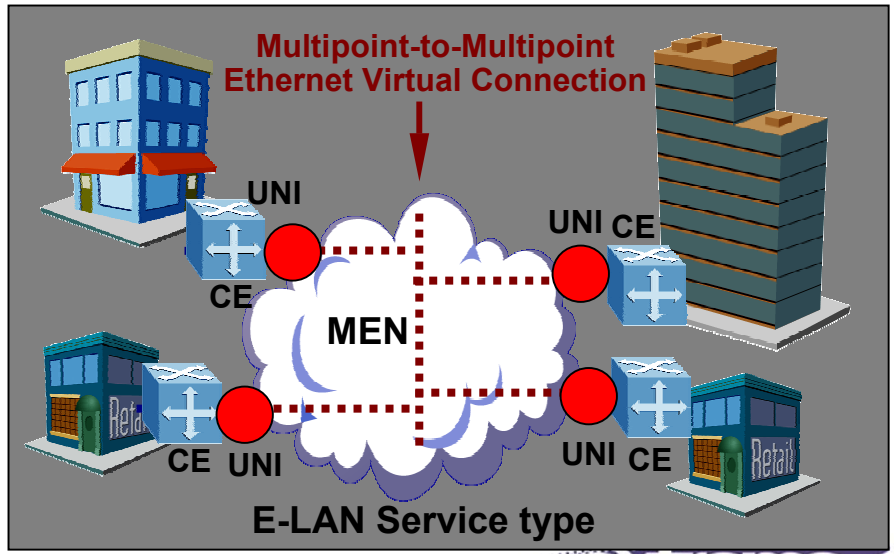

Private Line / Virtual Private Line

- Ethernet Private Line
  - Dedicated UNIs
  - Analogous to TDM Circuits
  - Customer VLAN transparency
- Ethernet Virtual Private Line
  - Service Multiplexed UNI
  - Analogous to Frame Relay
  - Service selected by C-VLAN

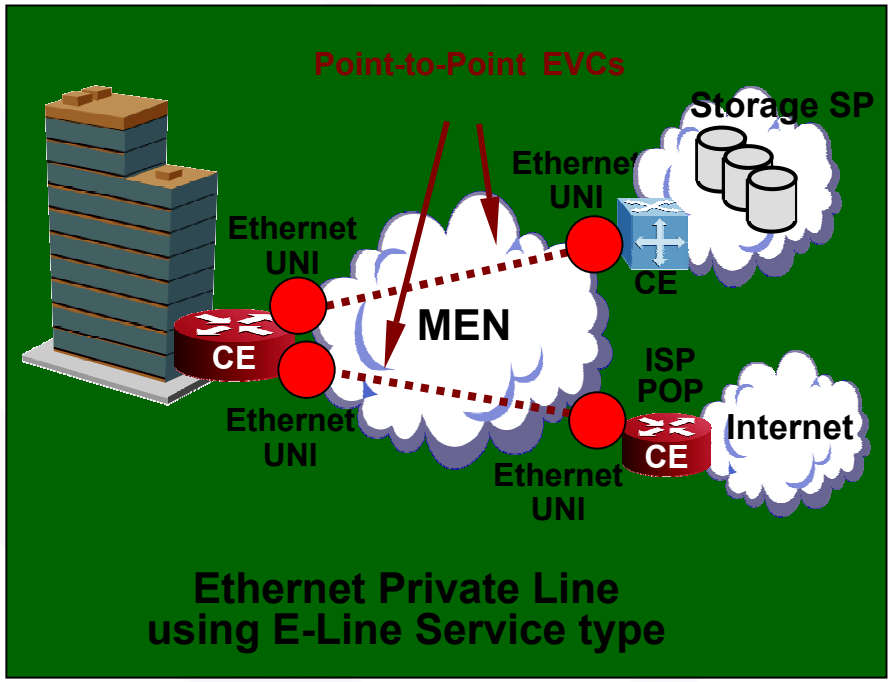

Ethernet Private Line using E-Line Service type

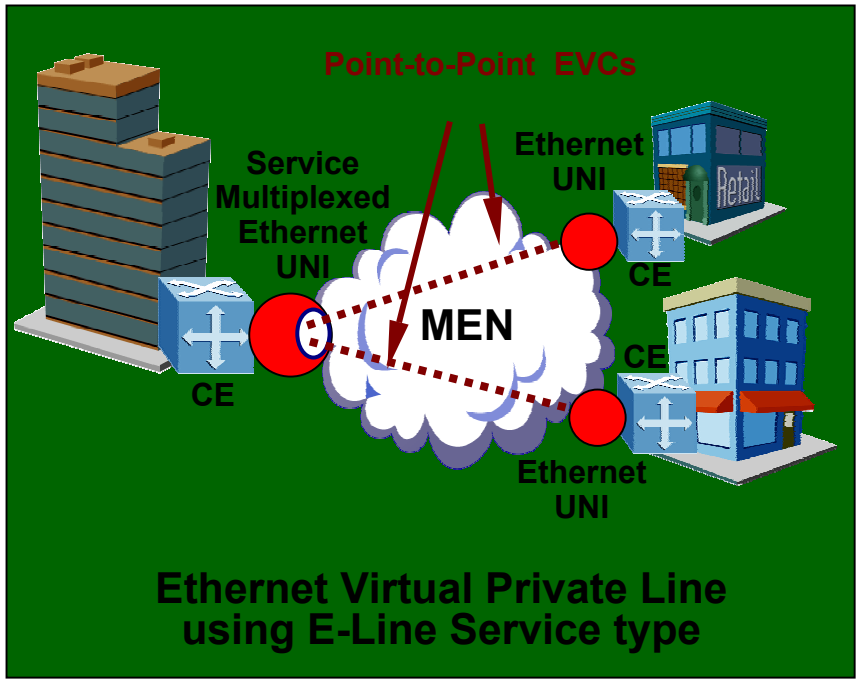

Ethernet Virtual Private Line using E-Line Service type

Transparent LAN Service

- Transparent LAN Service (TLS)
  - Uses E-LAN Service Type
    - (multipoint EVC)
  - Dedicated UNIs or Service-Multiplexed UNIs
  - Full transparency of L2 control protocols
- TLS makes the MEN look like a private LAN

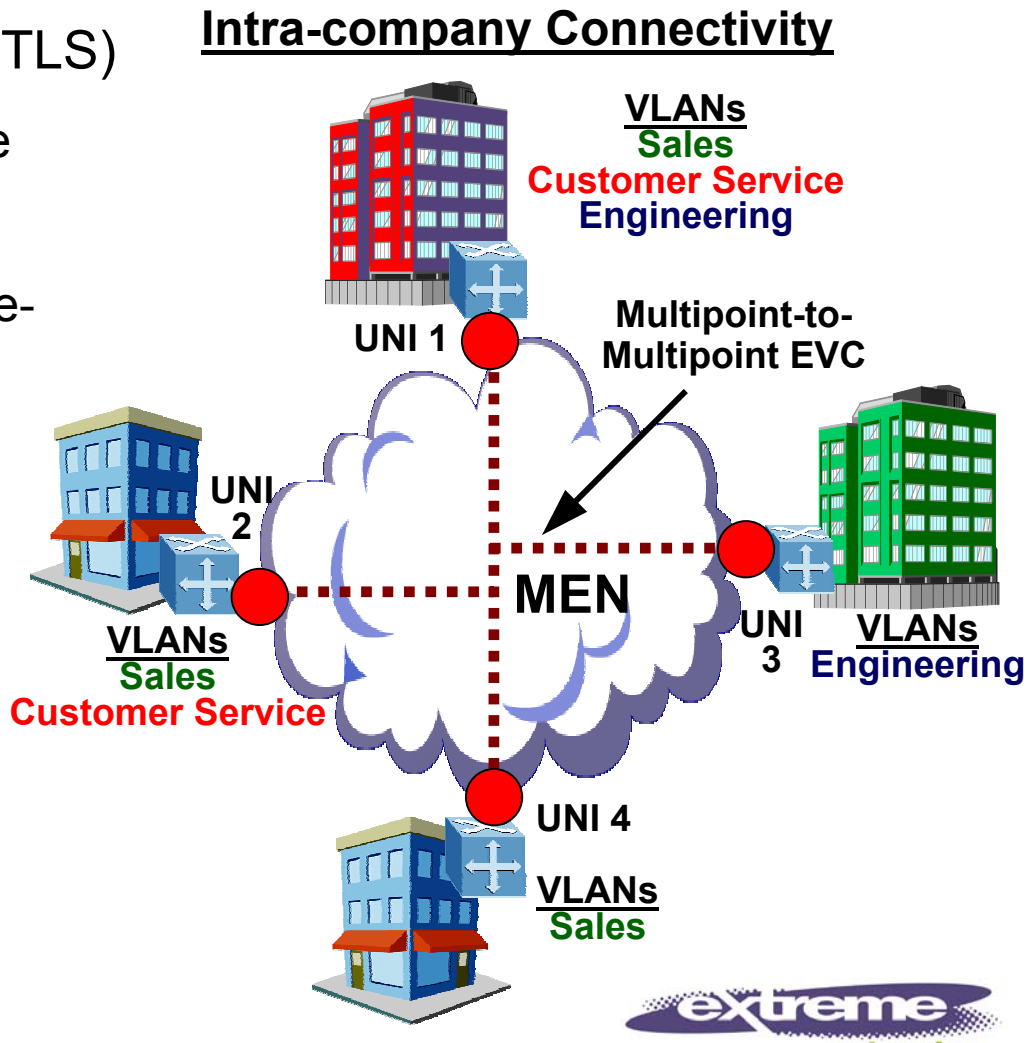

Bandwidth Profiles and Parameters

- Choice of Bandwidth Profiles
  - Ingress BW Profile per User-Network Interface (UNI)
  - Ingress BW Profile per Ethernet Virtual Connection (EVC)
  - Ingress BW Profile per Class of Service (CoS)
- Bandwidth Profile Parameters for Dual Rate Control
  - Committed Rate (CIR) and Burst Size (CBS)
    - assures frame delivery meets service level performance objectives
  - Excess Rate (EIR) and Burst Size (EBS)
    - controls amount of excess frame delivery allowed
- Service Performance Parameters
  - Frame Delay (Latency)
  - Frame Jitter (Latency variation)
  - Frame Loss

Three types of Bandwidth Profiles

- per UNI

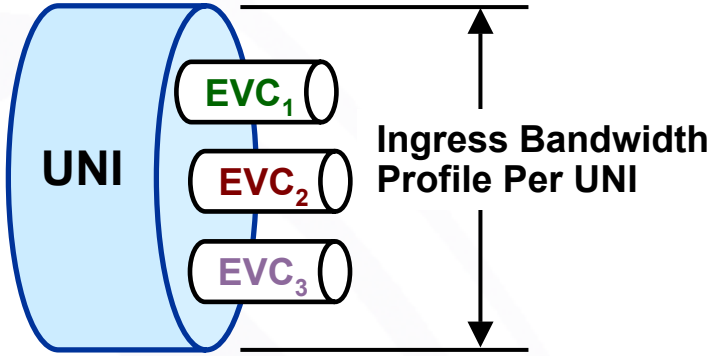

- per Ethernet Virtual Connection

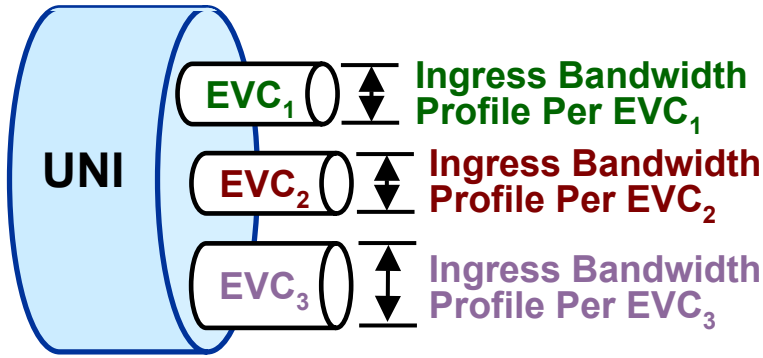

- per Class of Service

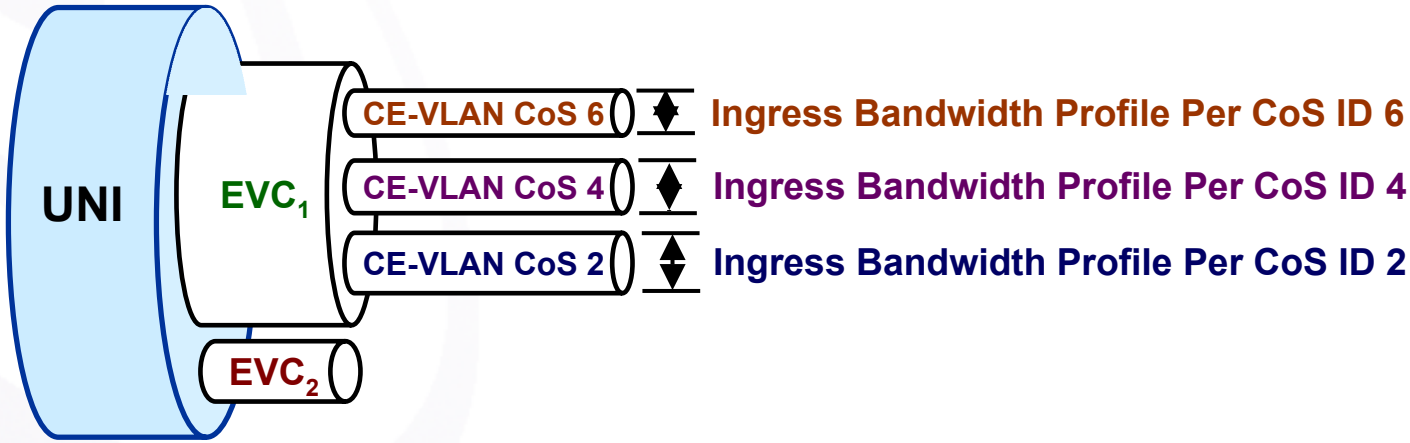

Two Rate Three Color Marker

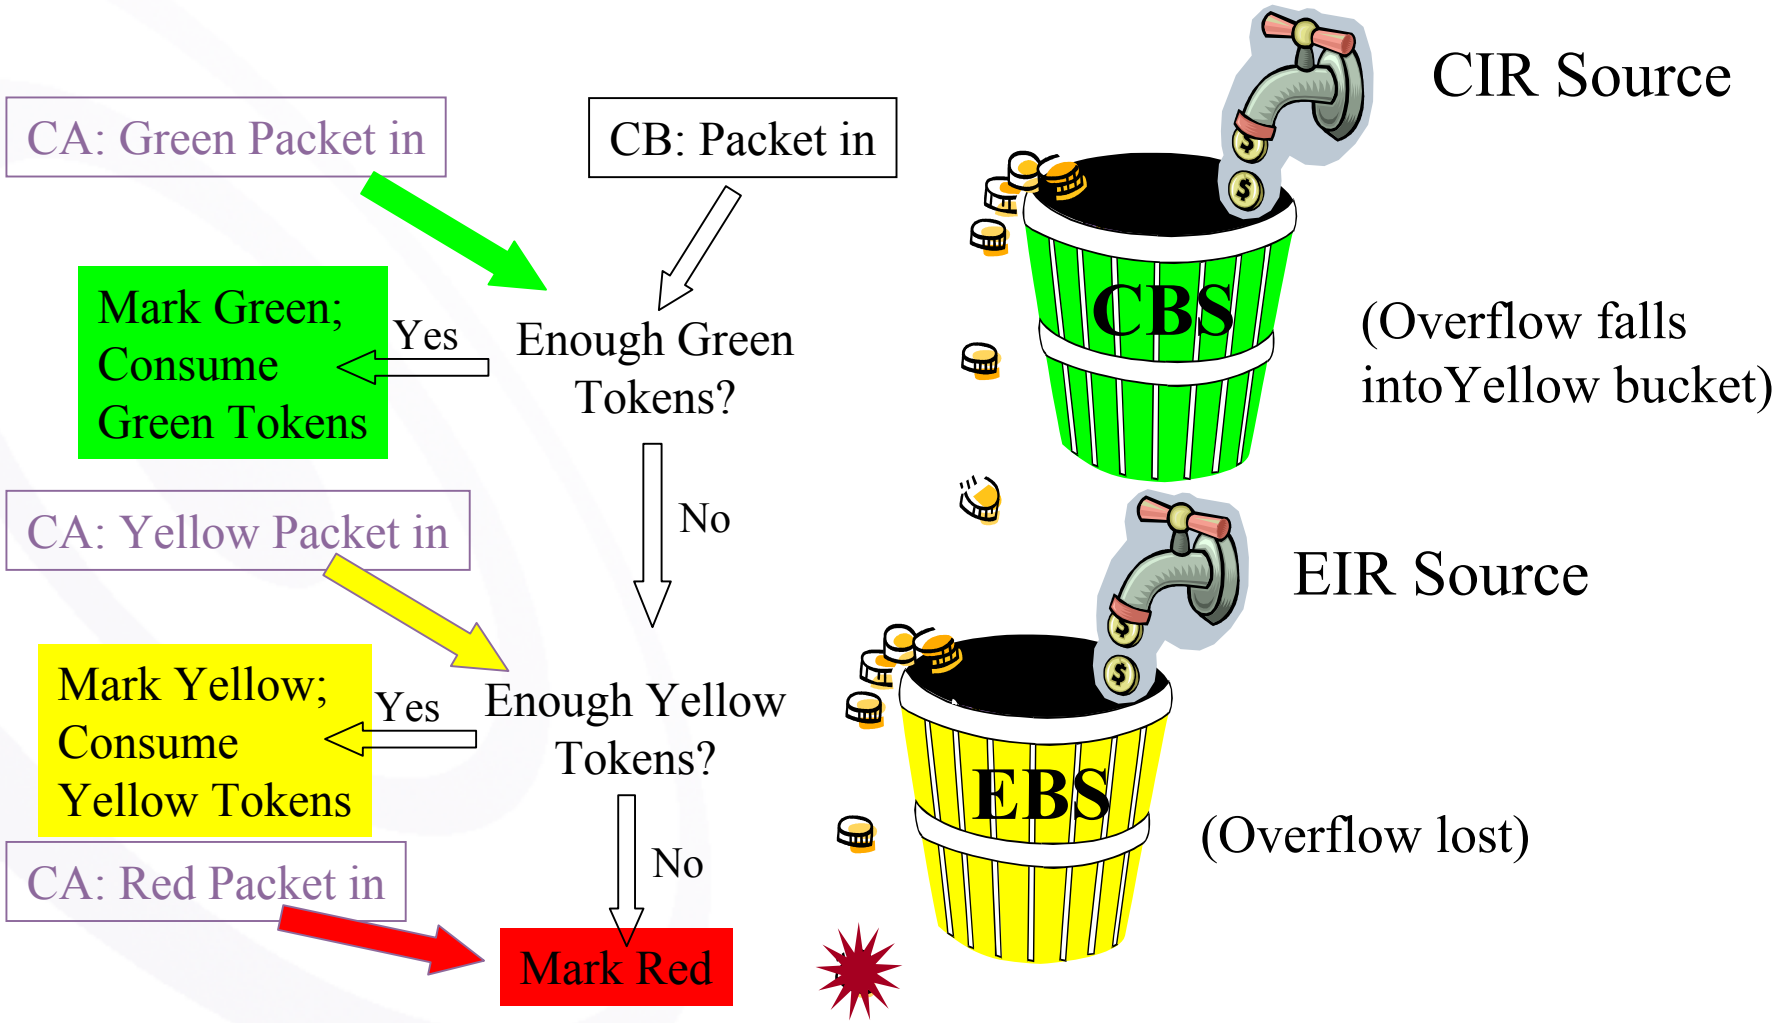

Example Metro Ethernet SLA

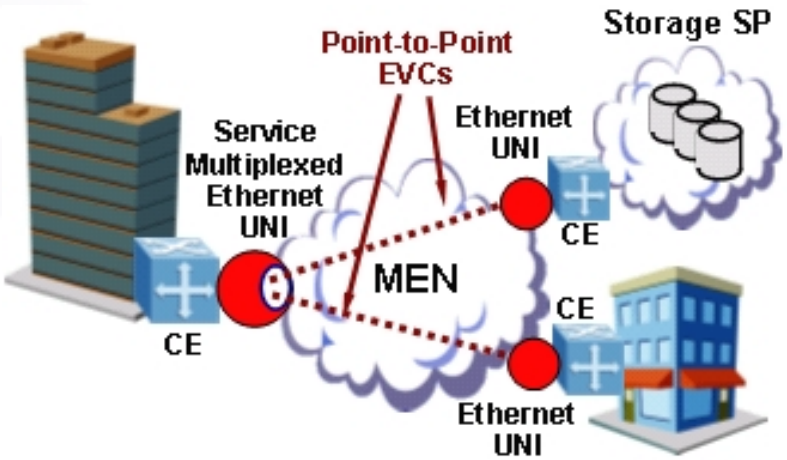

- E-Line Service
- 4 Classes of Service
- CoS determined via 802.1p CoS ID
- Common type of SLA used with CoS-based IP VPNs

Service Class	Service Characteristics	CoS ID	Bandwidth Profile per EVC per CoS ID	Service Performance
Premium	Real-time IP telephony or IP video applications	6, 7	CIR > 0 EIR = 0	Delay < 5ms Jitter < 1ms Loss < 0.001%
Silver	Bursty mission critical data applications requiring low loss and delay (e.g., Storage)	4, 5	CIR > 0 EIR ≤ UNI Speed	Delay < 5ms Jitter = N/S Loss < 0.01%
Bronze	Bursty data applications requiring bandwidth assurances	2, 3	CIR > 0 EIR ≤ UNI Speed	Delay < 15ms Jitter = N/S Loss < 0.1%
Standard	Best effort service	0, 1	CIR=0 EIR=UNI speed	Delay < 30ms Jitter = N/S Loss < 0.5%

Metro Ethernet Forum Summary

- Specifies the User Network Interface
- Specifies Service Types
  - E-Line for point-to-point services
  - E-LAN for multipoint services
- Defines Service Attributes
  - Service Multiplexing and Service Selection
  - Class of Service Selection
  - Bandwidth Profiles
  - Performance Parameters

IEEE 802.1ad
Provider Bridging

p802.1ad Provider Bridging PAR

- Purpose

  This standard will enable a Service Provider to offer the equivalent of separate LAN segments, Bridged or Virtual Bridged LANs, to a number of users, over the providers bridged network. This standard will enable the use of the architecture and protocols of IEEE Std 802.1Q, and provide for interoperability and consistent management.

- Scope

  To develop an architecture and bridge protocols, compatible and interoperable with existing Bridged Local Area Network protocols and equipment, to provide separate instances of the MAC service to multiple independent users of a Bridged Local Area Network in a manner that does not require cooperation among the users, and requires a minimum of cooperation between the users and the provider of the MAC service. To define basic management of users' MAC service.

Simplified Model

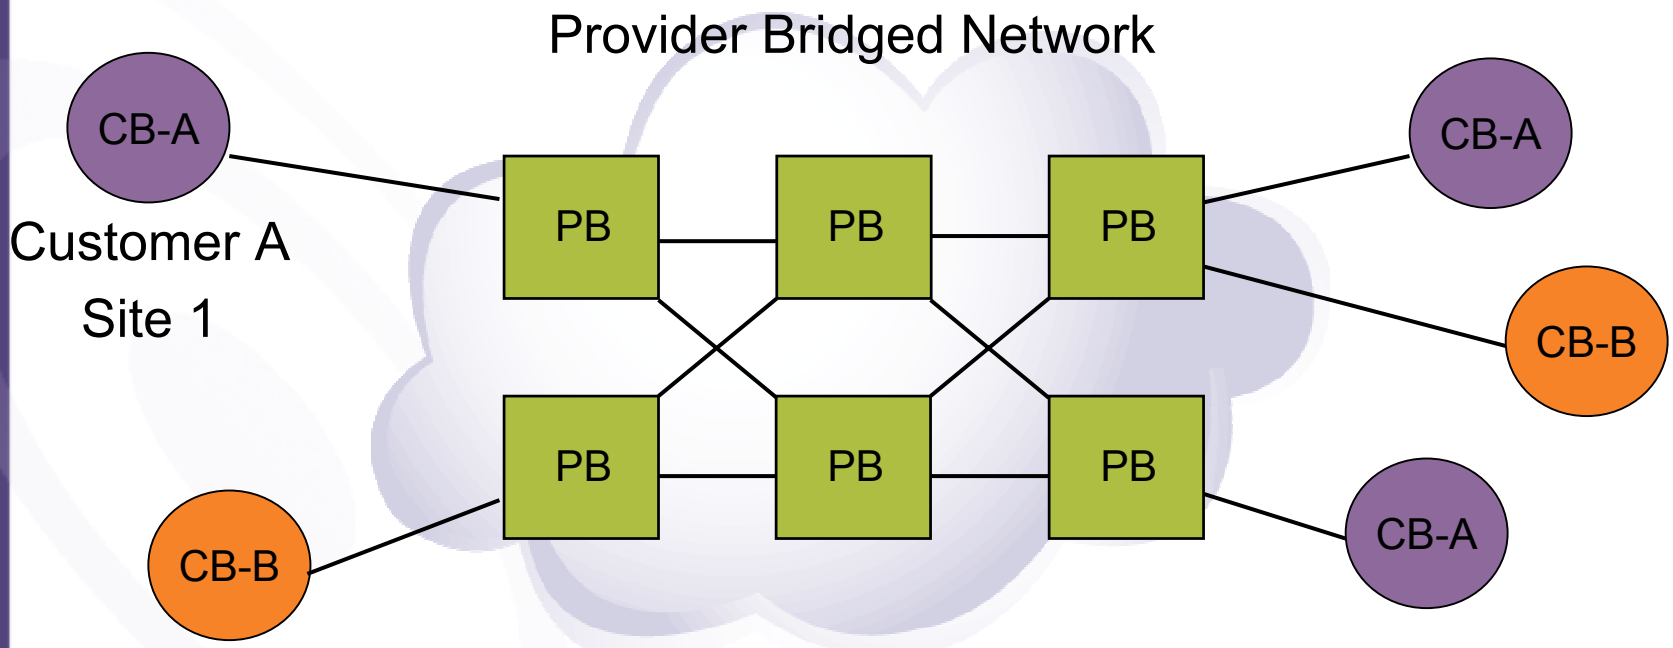

Goal: Transparently interconnect all of Customer A sites and all of Customer B sites while maintaining complete isolation between Customers A and B.

802.1Q Bridges almost meet the goal

- VLAN tag can be used as a Customer ID
  - VLANs constrain broadcast domain so one customer never sees another customer's packets.
  - Ingress/Egress VLAN filtering rules per port enable access control enforcement.
- But there are problems:
  1. Customer packets must be untagged.
     - Customer assigned VLAN tags cannot be transported.
     - No means of indicating packet priority.
     - Cannot access multiple services through a single port.
  2. No customer/customer or customer/provider separation in the control plane (for control protocol packets such as Spanning Tree BPDUs).

First Level Solution

1. Give the Provider network it's own VLAN tag
   - Create a "Service VLAN Tag" (S-TAG) that has analogous format and function as a VLAN tag, but is present only on the Provider network and is separate from the Customer VLAN Tag (C-TAG).
   - Proprietary implementations known as "Tag Stacking", "Q-in-Q", or "VMAN tag".

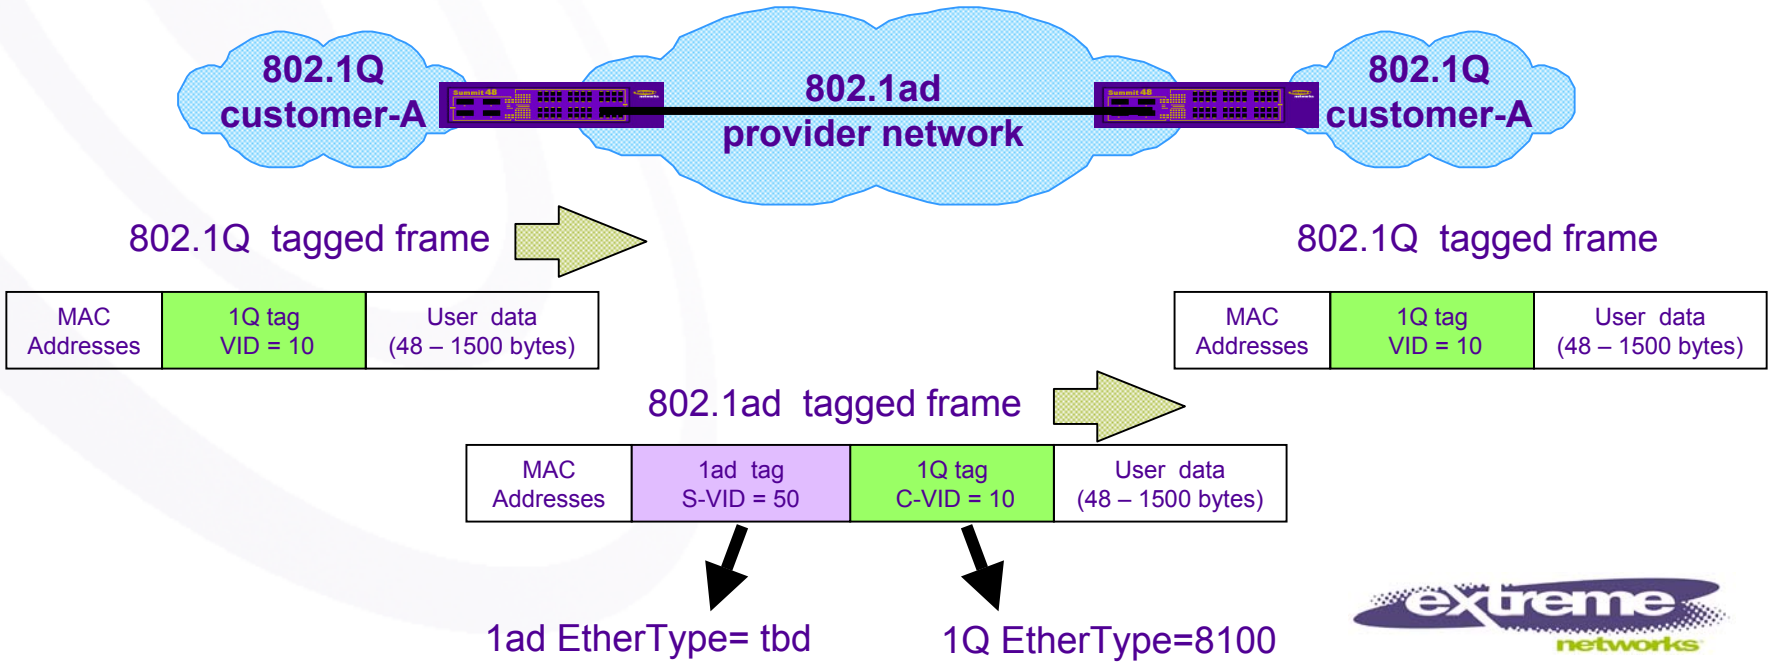

802.1D – 1998 Transparent Bridge

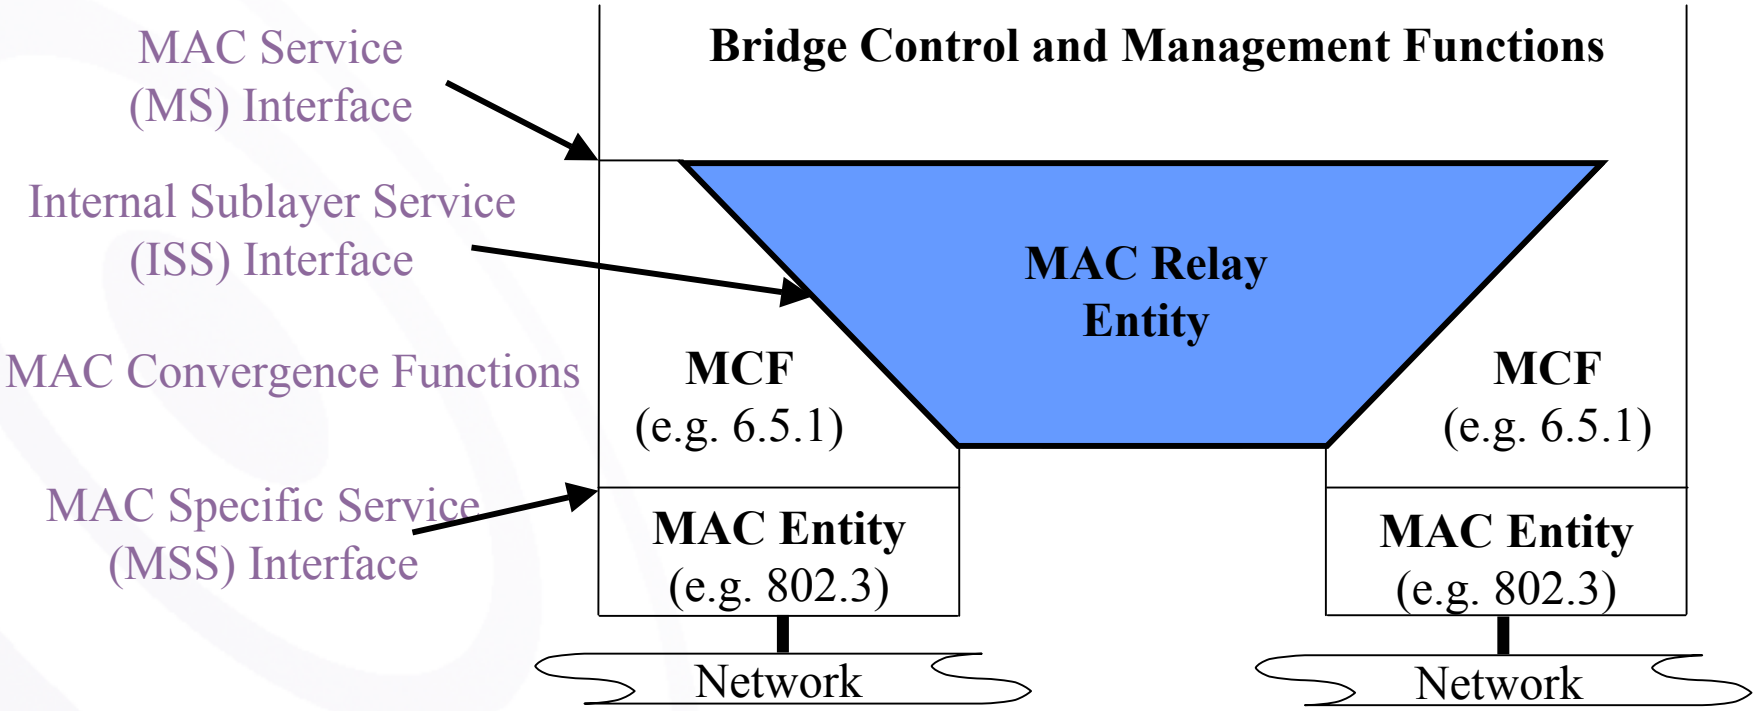

802.1Q – 1998 VLAN Bridge

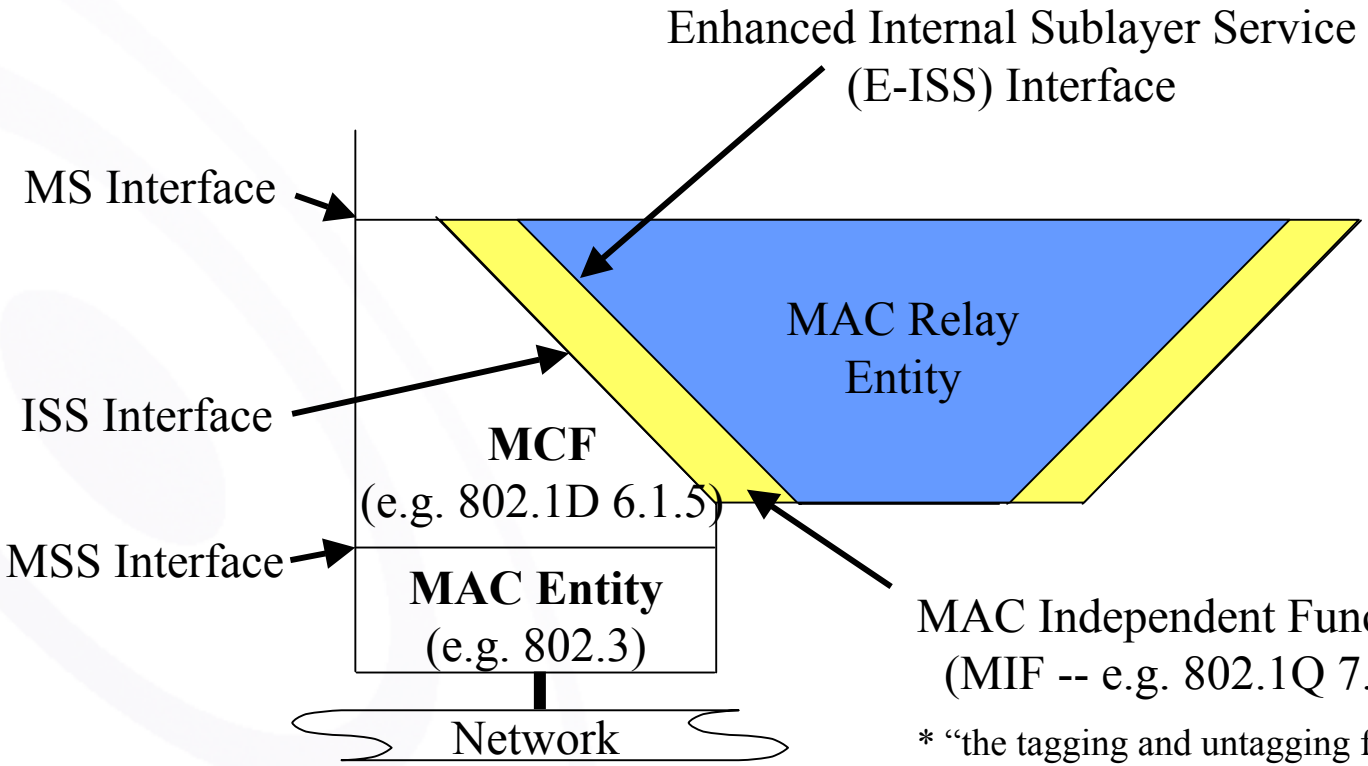

802.1ad – Provider Bridge

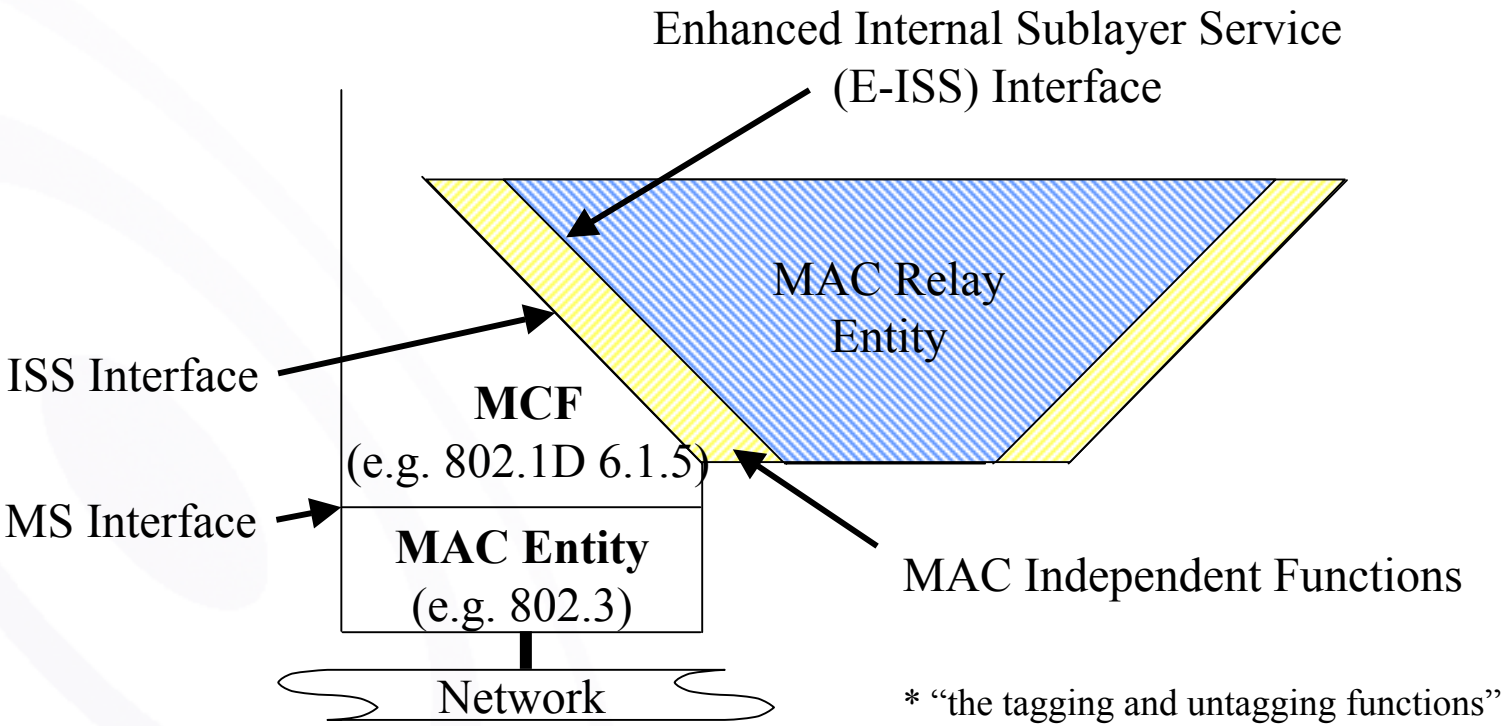

The mapping between the ISS and the E-ISS is the same as in 802.1Q 7.1.2 except that the operations are performed on a different tag – the Service VLAN Tag (S-TAG) rather than the Customer VLAN Tag (C-TAG).

Simple Provider Service

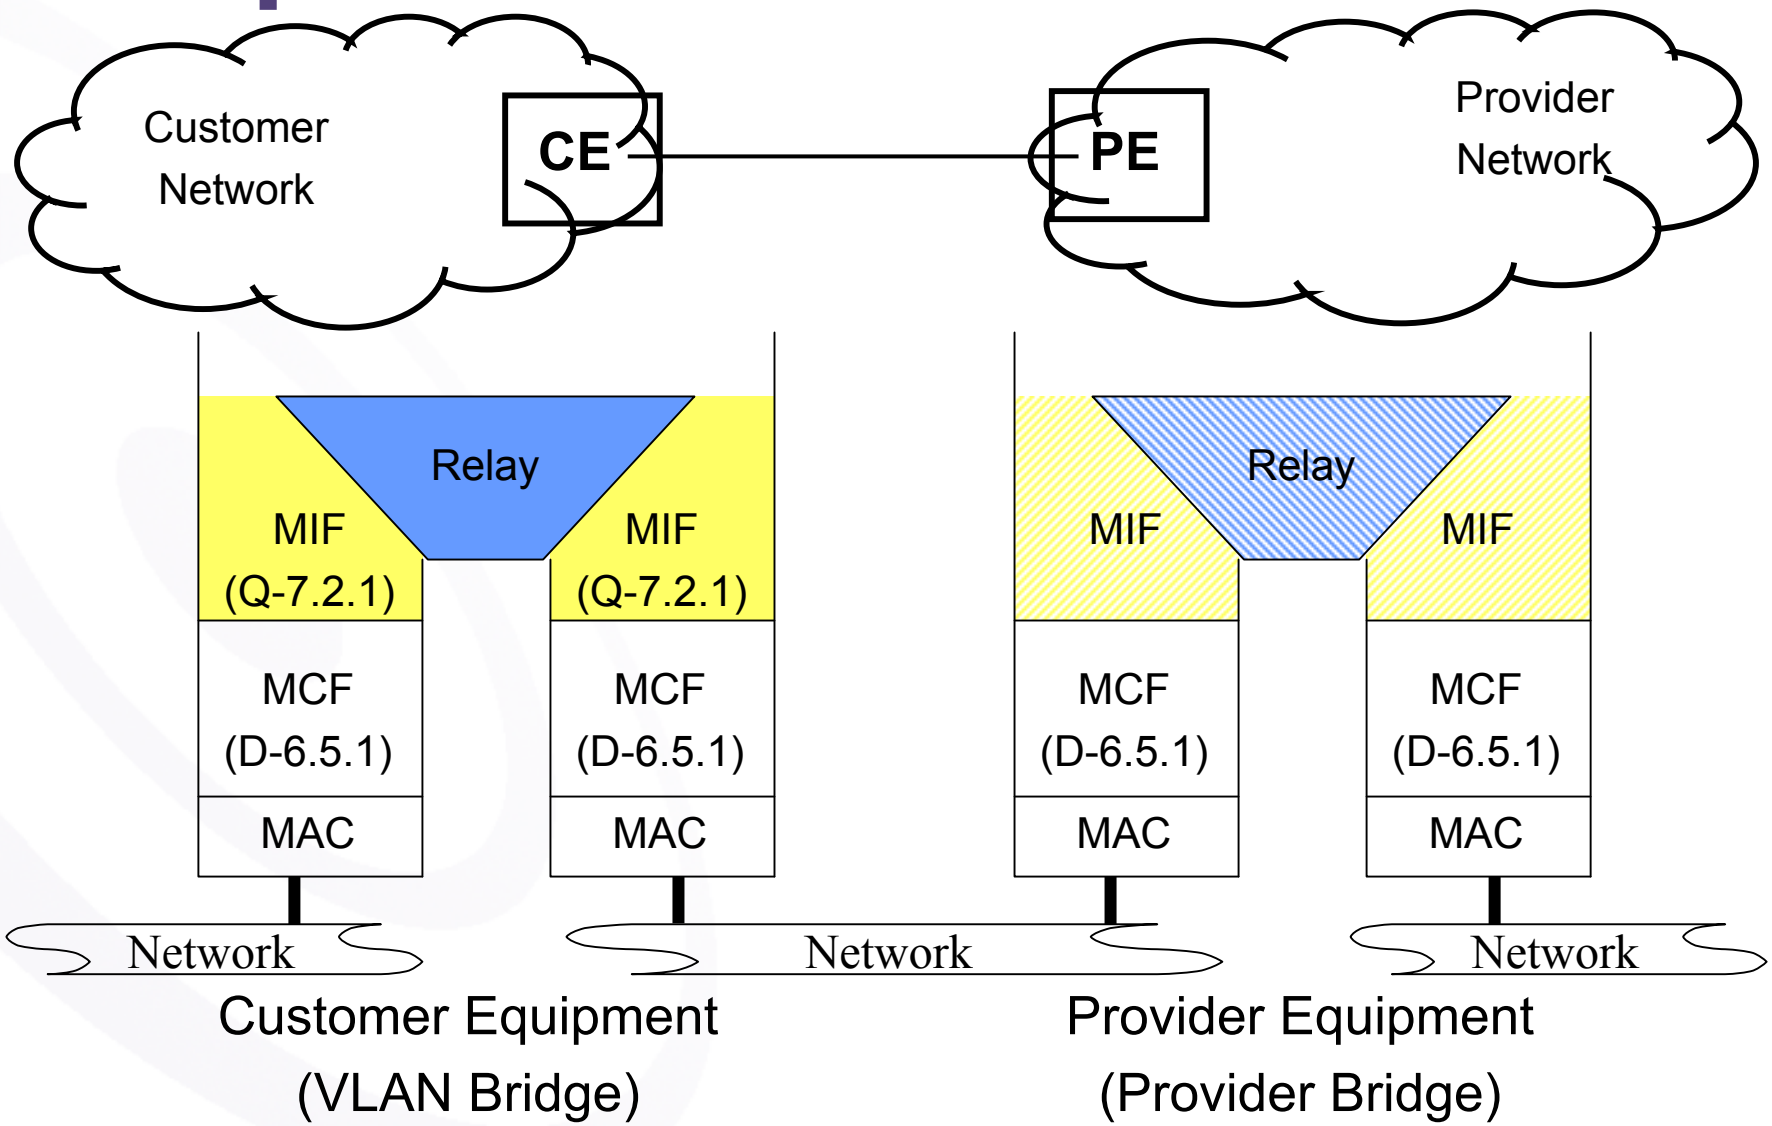

Simple Provider Service

- All the Provider Bridge does is insert a Service Tag in all frames received from the Customer Equipment.
- Minimal changes to make a 802.1Q bridge a provider bridge:
  - assign a new Service Tag Ethertype, and
  - assign a new Provider BPDU Address.
- This is sufficient provided that:
  - All customer traffic maps to a single provider service instance.
  - All customer traffic has the same priority in the provider network.
- Can this be extended to support service multiplexing (accessing multiple service instances through a single Customer-Provider connection) and prioritized services?

“Dual Bridge” Provider Edge Model

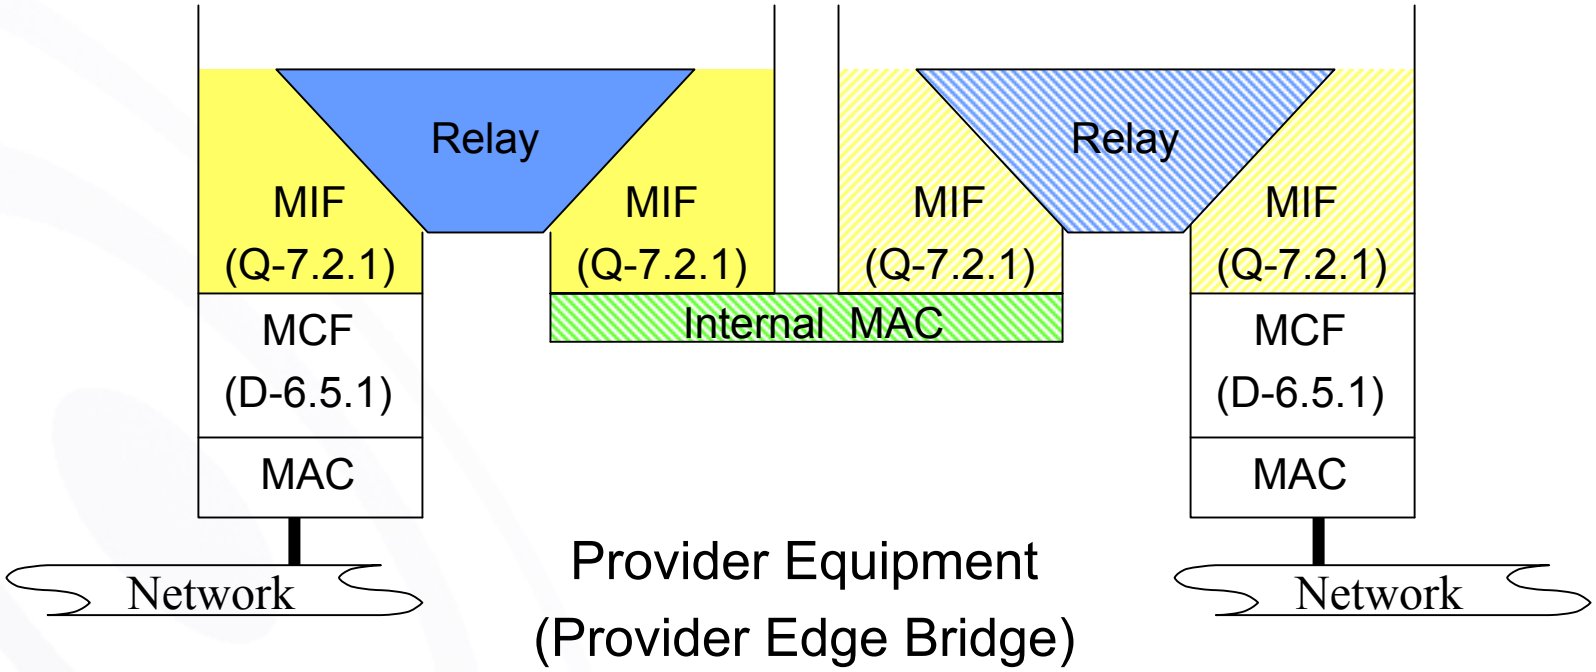

- Specify behavior of a Provider Edge as two bridges in one box.
  - Customer facing side operates on Customer VLAN Tags and BPDUs
  - Provider facing side operates on Provider Tags and BPDUs
  - Interconnect with an “internal port” per service instance

802.1ad Provider Edge Service

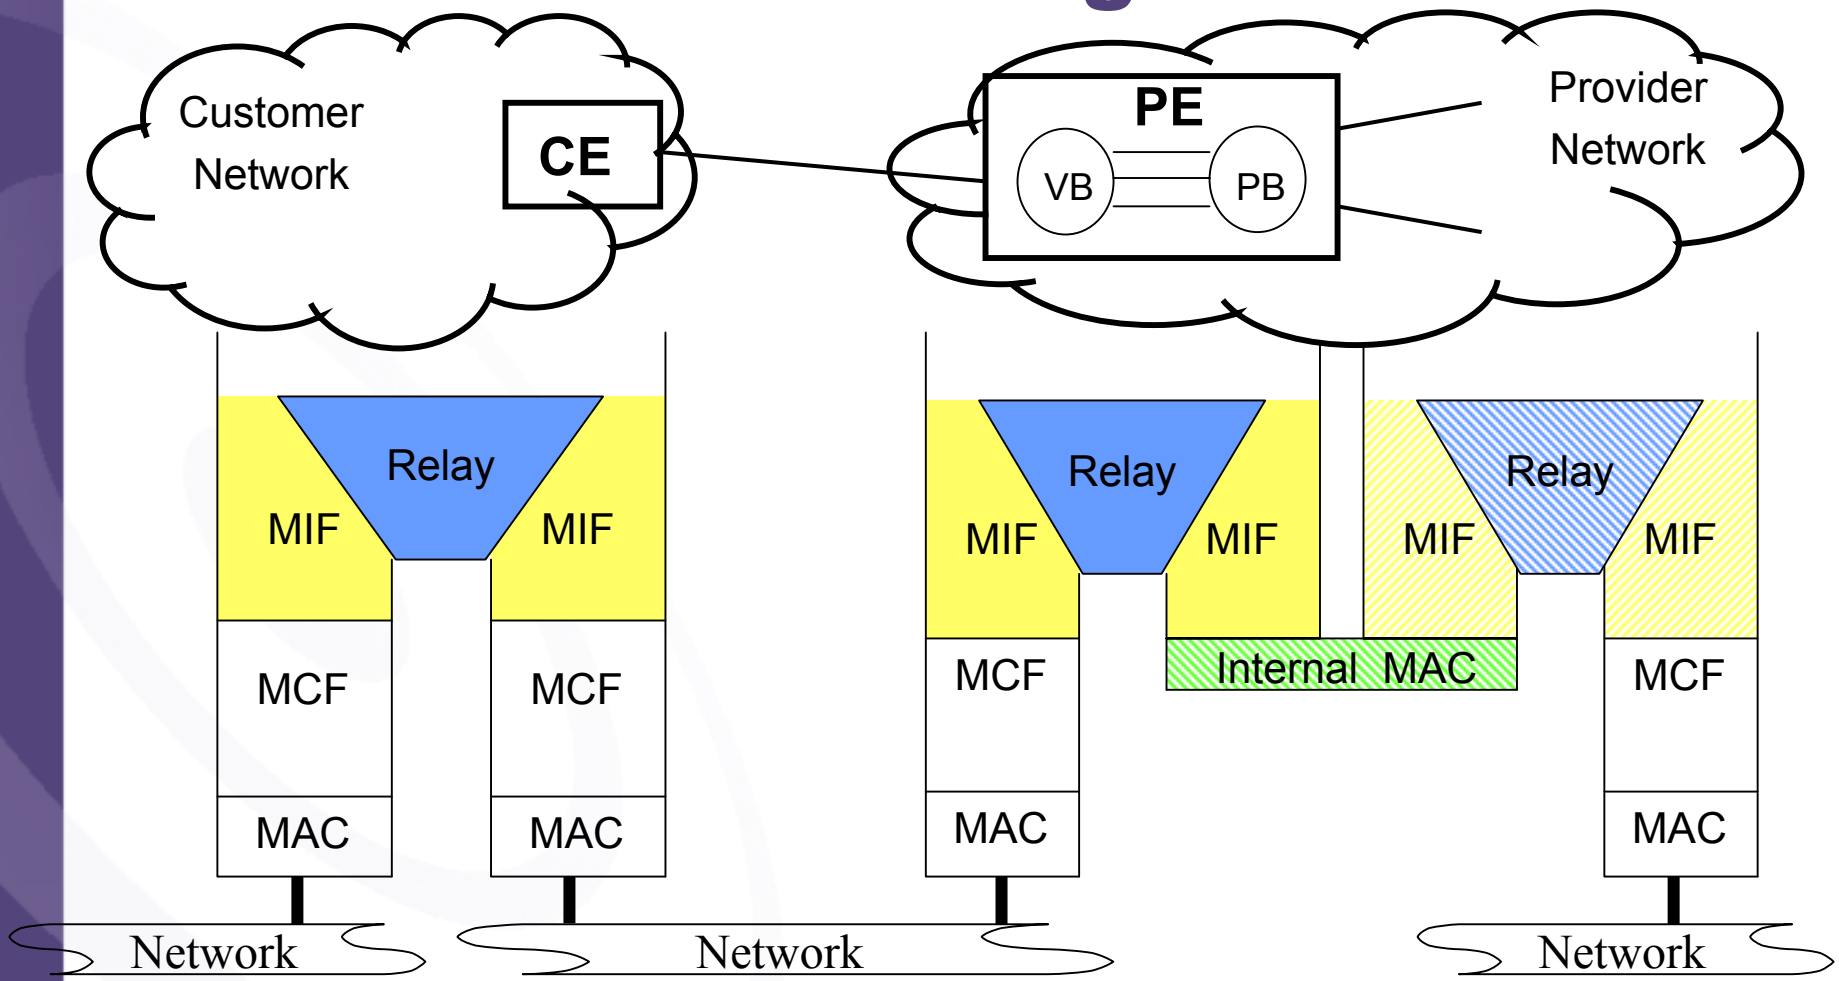

Customer Equipment
(VLAN Bridge)

Provider Equipment
(Provider Edge Bridge)

Service Multiplexing – Data Plane

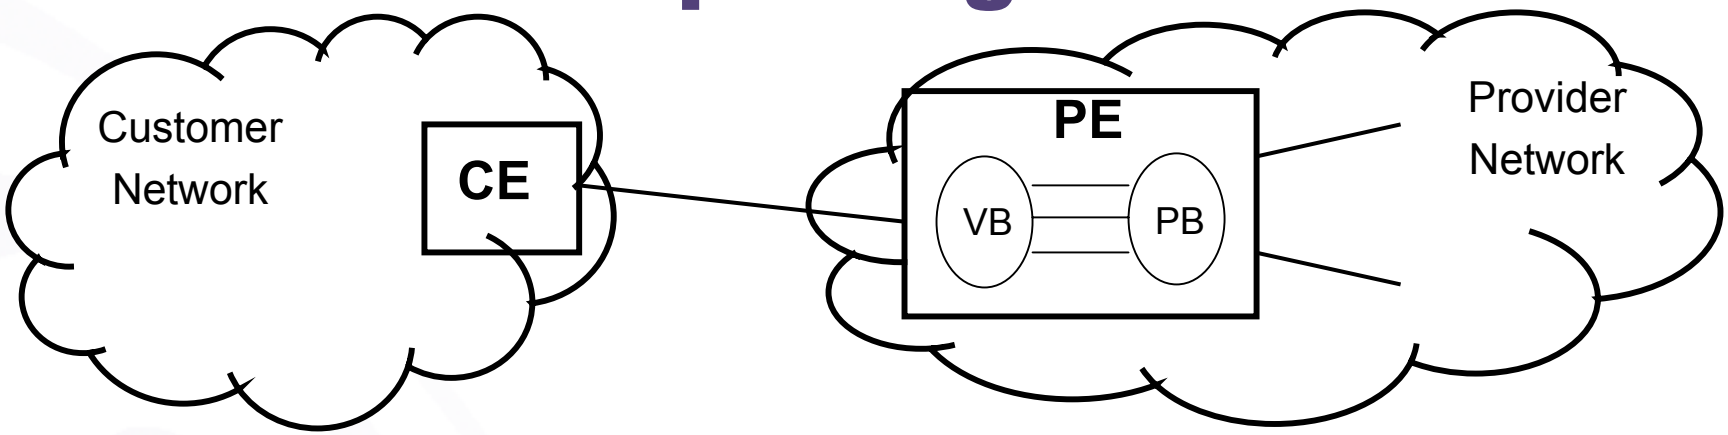

- Customer accesses 3 different Provider Services over a single physical link to the Provider Edge Bridge.
  - VLAN Bridge portion of PE connects to Provider Bridge portion via 3 internal ports – one per service instance (EVC in MEF terminology).
  - VLAN Bridge portion of PE selects service based on Customer VLAN IDs by forwarding packets for each service to the appropriate internal port.
  - Provider Bridge portion of PE creates Provider Tag using a Service VLAN ID assigned to the internal port.

Service Multiplexing – Control Plane

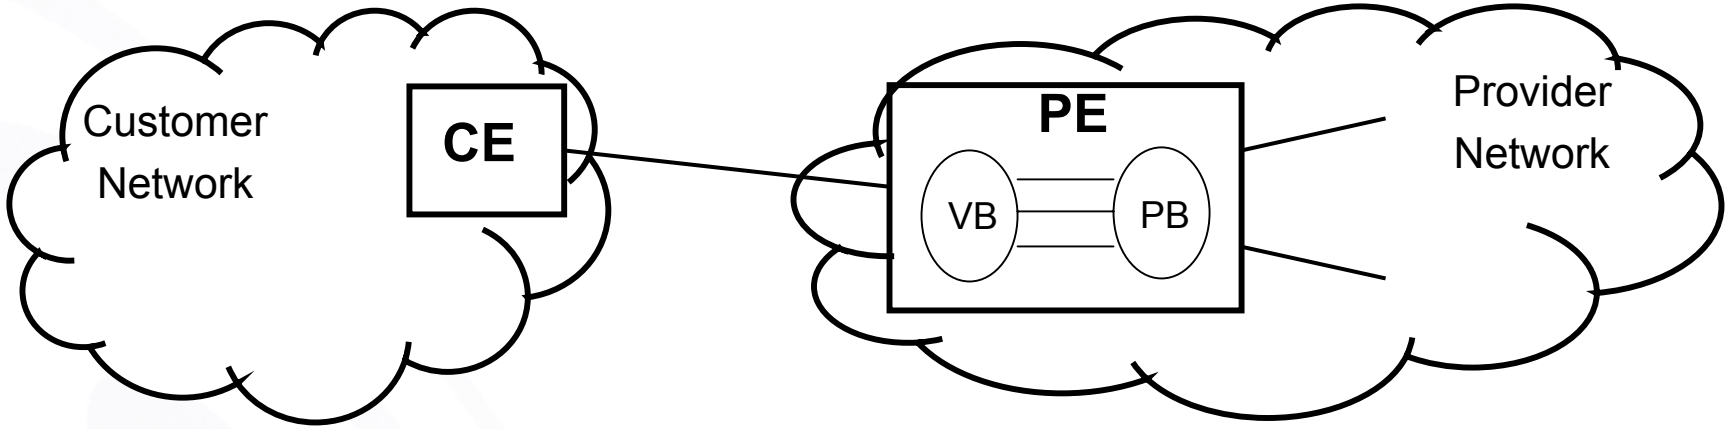

- Customer BPDUs must be transported across each service instance.
  - VLAN Bridge portion of PE participates in Customer Spanning Tree – receives, processes, and transmits Customer BPDUs on each customer facing port and each internal port.
  - Provider Bridge portion takes Customer BPDUs received on the internal ports, tags them with the Service VLAN ID, and “tunnels” them across the Provider Network.
  - Provider Bridge portion may participate in Provider Spanning Tree which is completely isolated from Customer Spanning Tree.

Multiple Priorities

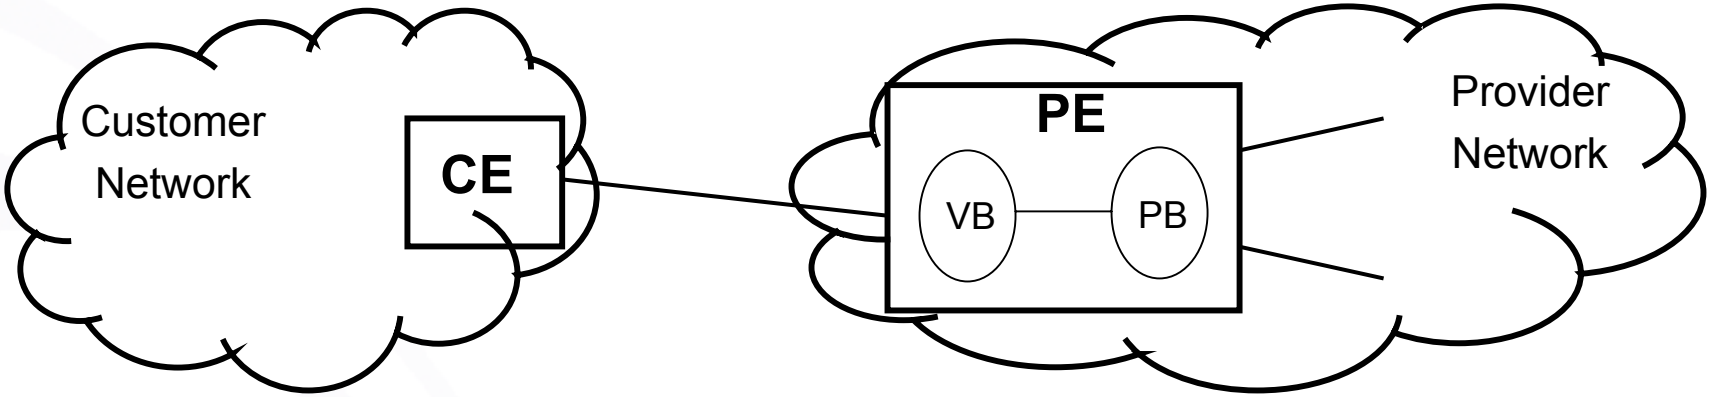

- Customer accesses single Provider Services that handles multiple priorities.
  - VLAN Bridge portion of PE uses the priority field of the Customer VLAN tag to determine the priority for the internal port.
  - Internal port conveys the priority information from the VB to the PB.
  - Provider Bridge portion uses this priority to create the Priority Code Point in the Service VLAN Tag.
    - Provider Bridge may map Customer specified priorities to different priority levels on the Provider network.

Drop Precedence Marking

- Priority Code Point replaces 3-bit priority field in Service VLAN tag

Priority Code Point		7	6	5	4	3	2	1	0
Priority / Drop Eligible	**8P0D** (default)	7	6	5	4	3	2	0	1
	7P1D	7	6	4	4DE	3	2	0	1
	6P2D	7	6	4	4DE	2	2DE	0	1
	5P3D	7	6	4	4DE	2	2DE	0	0DE

Handling Customer Control Protocols

- Spanning Tree BPDUs will be tagged with a Service Tag and transported across the Provider Network
  - Provider Bridges will not recognize the Customer BPDU address as a “reserved” address (“reserved” addresses cannot pass through a bridge).
  - Provider Bridges will use a different reserved address for Provider BPDUs
- Handling of other Layer-2 Protocols is largely determined by the architecture
  - Some Protocols (e.g. 802.3x Pause) are terminated at the MAC and never reach the internal interfaces of the bridge.
  - Other protocols (e.g. 802.3ad Link Aggregation and 802.1X Port Based Access Control) operate between the Customer Bridge and the Provider Bridge when using the current reserved addresses.
    - Considering adding new reserved addresses to allow 802.3ad and 802.1X to operate between Customer Bridges across the Provider Network.

802.1ad: Provider Bridges Summary

- Service Identification
  - Standardize Q-in-Q (VMAN) tags
  - Service Tags will have unique Ethertype
- Service Selection
  - Service ID derived from ingress port and Customer-VID
- Traffic Classification
  - Class of Service in Provider network derived as a function of Service ID and Customer 802.1p bits
  - Priority marking extended to include drop precedence
- Control Protocol
  - Separation of Customer and Provider Control Domains
  - Customer Spanning Tree Protocol packets transported through Provider network

Scaling Solutions

Scalability Issues: Address Learning

- Concern that bridges in the core of a Provider Network will need to learn millions of Customer MAC addresses.
- P802.1ad draft includes “enhanced” learning criteria that MAC addresses only need to be learned for a VLAN if there are more than two ports active on that VLAN.
  - No learning is required for point-to-point services.
  - For multipoint services between N sites, addresses will only need to be learned on at most N-2 bridges.

Scalability Issues: Service ID Space

- Service Tag has a 12 bit ID field
  - Clearly a need to support more than 4096 service instances in a Provider Network
- Simply increasing the ID field ignore significant issues:
  - Control structures for ingress/egress filtering tables, spanning tree state tables, broadcast/flood port lists, etc.
  - Control protocols that have per VLAN fields (such as 802.1s Multiple Spanning Tree and GVRP).
- Other solutions mitigate the scalability issue:
  - Asymmetric or unidirectional VLANs allow creation of a point to multipoint network which can provide Internet Access for thousands of customers using only two Service IDs.
  - Islands of Provider Networks can be interconnected using emulated Ethernets (e.g. IETF VPLS).
  - New 802.1ah standard in development for Provider Backbone Bridging

Unidirectional VLANs

- Bridge Port facing ISP
  - Places all packets from Provider to Subscriber on Blue VLAN
  - Forwards packets on Red VLAN to Provider
- Bridge Port facing Subscribers
  - Places all packets from Subcribers to Provider on Red VLAN
  - Transmits packets on Blue VLAN to Subscriber
  - Blocks any packets on Red VLAN from going to (another) Subscriber

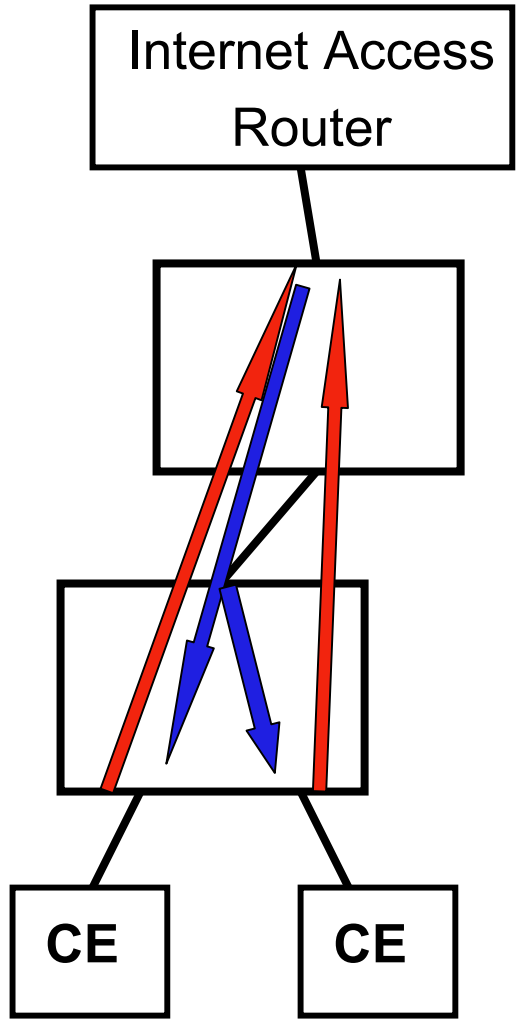

Provider Bridging Access with MPLS Core

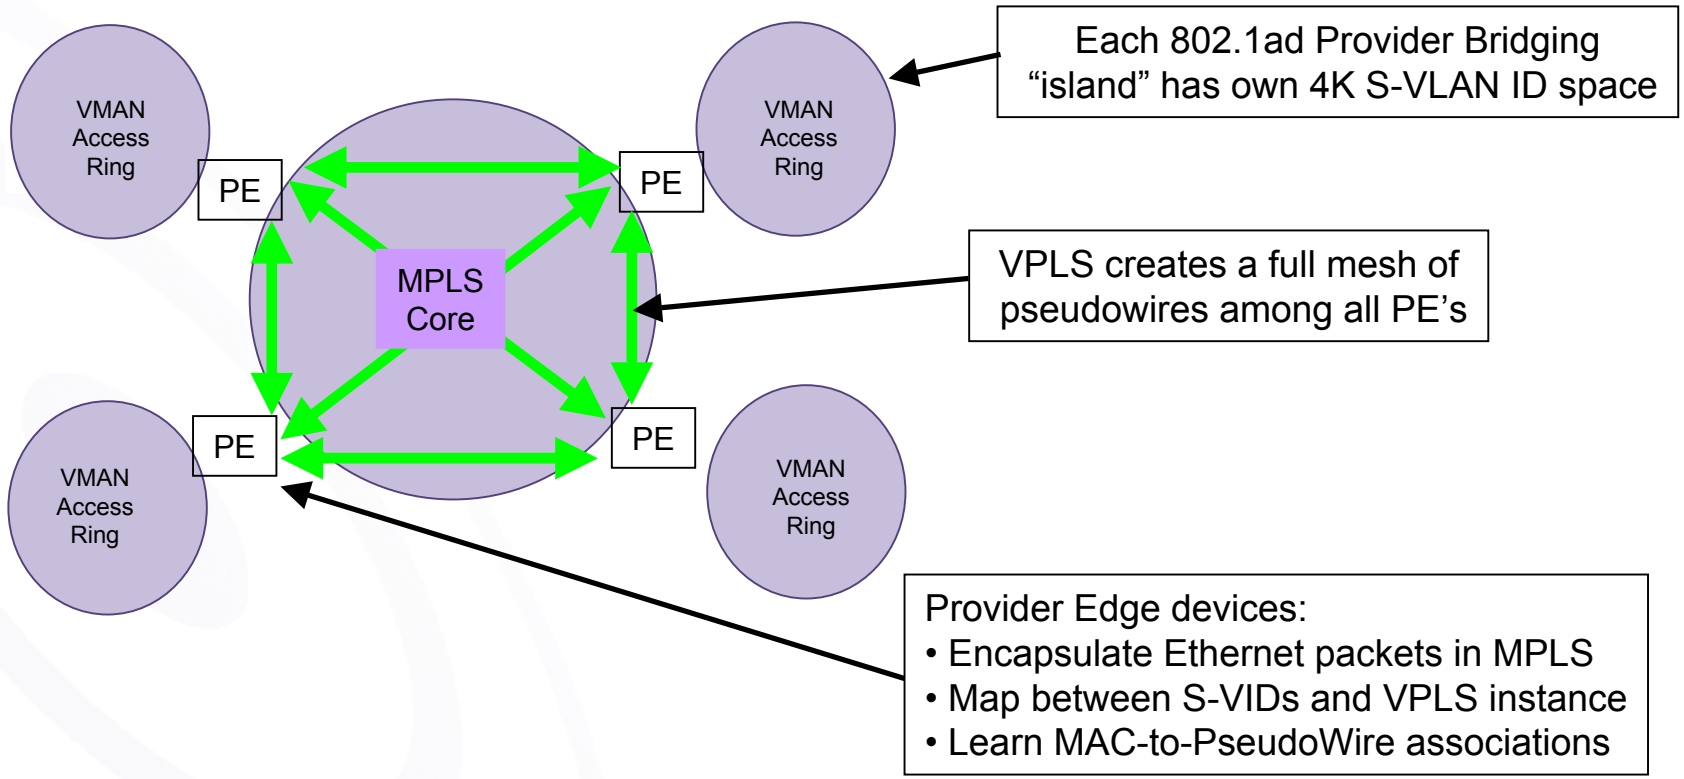

IETF L2VPN Virtual Private LAN Service (VPLS) over MPLS

802.1ah Provider Backbone Bridging

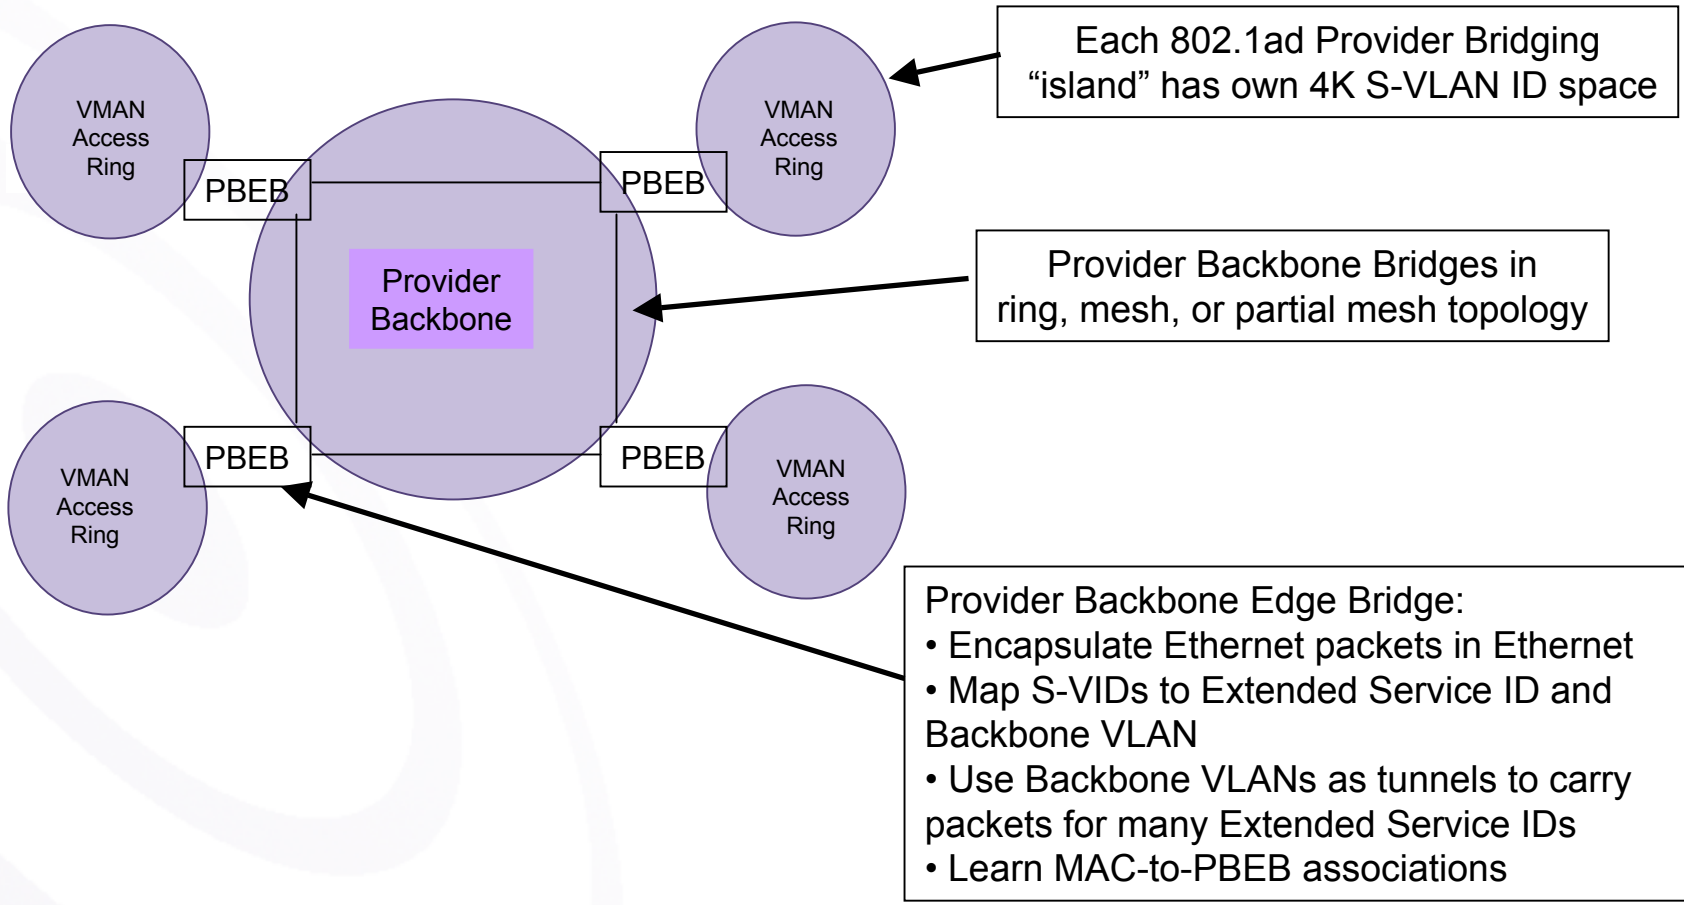

IEEE 802.1ah Provider Backbone Bridging creates an hierarchical Layer 2 topology

Summary

Summary

- Ethernet and Ethernet L2/L3 Switches have been widely adopted for building Next Generation Broadband Networks
- This has driven the development of new industry standards in the MEF, IEEE, and IETF for:
  - Common Ethernet Service Models and Definitions
  - Standard Bridging functionality and packet tagging formats for customer identification and isolation in the control and data planes
  - Quality of Service, Priority, and Traffic Management parameters for meeting Service Level Agreements
  - Scaling Metro Ethernet Networks
  - Monitoring and Managing Metro Ethernet Networks

Thank You