p802.1ad Provider Bridging

Stephen Haddock

December 3, 2003

Agenda

- 802.1 Overview
- Basic Concepts of p802.1ad Provider Bridging
- Provider Bridge Architecture
- Scalability Issues
- Summary

802.1 (Bridging) projects (Nov 03)

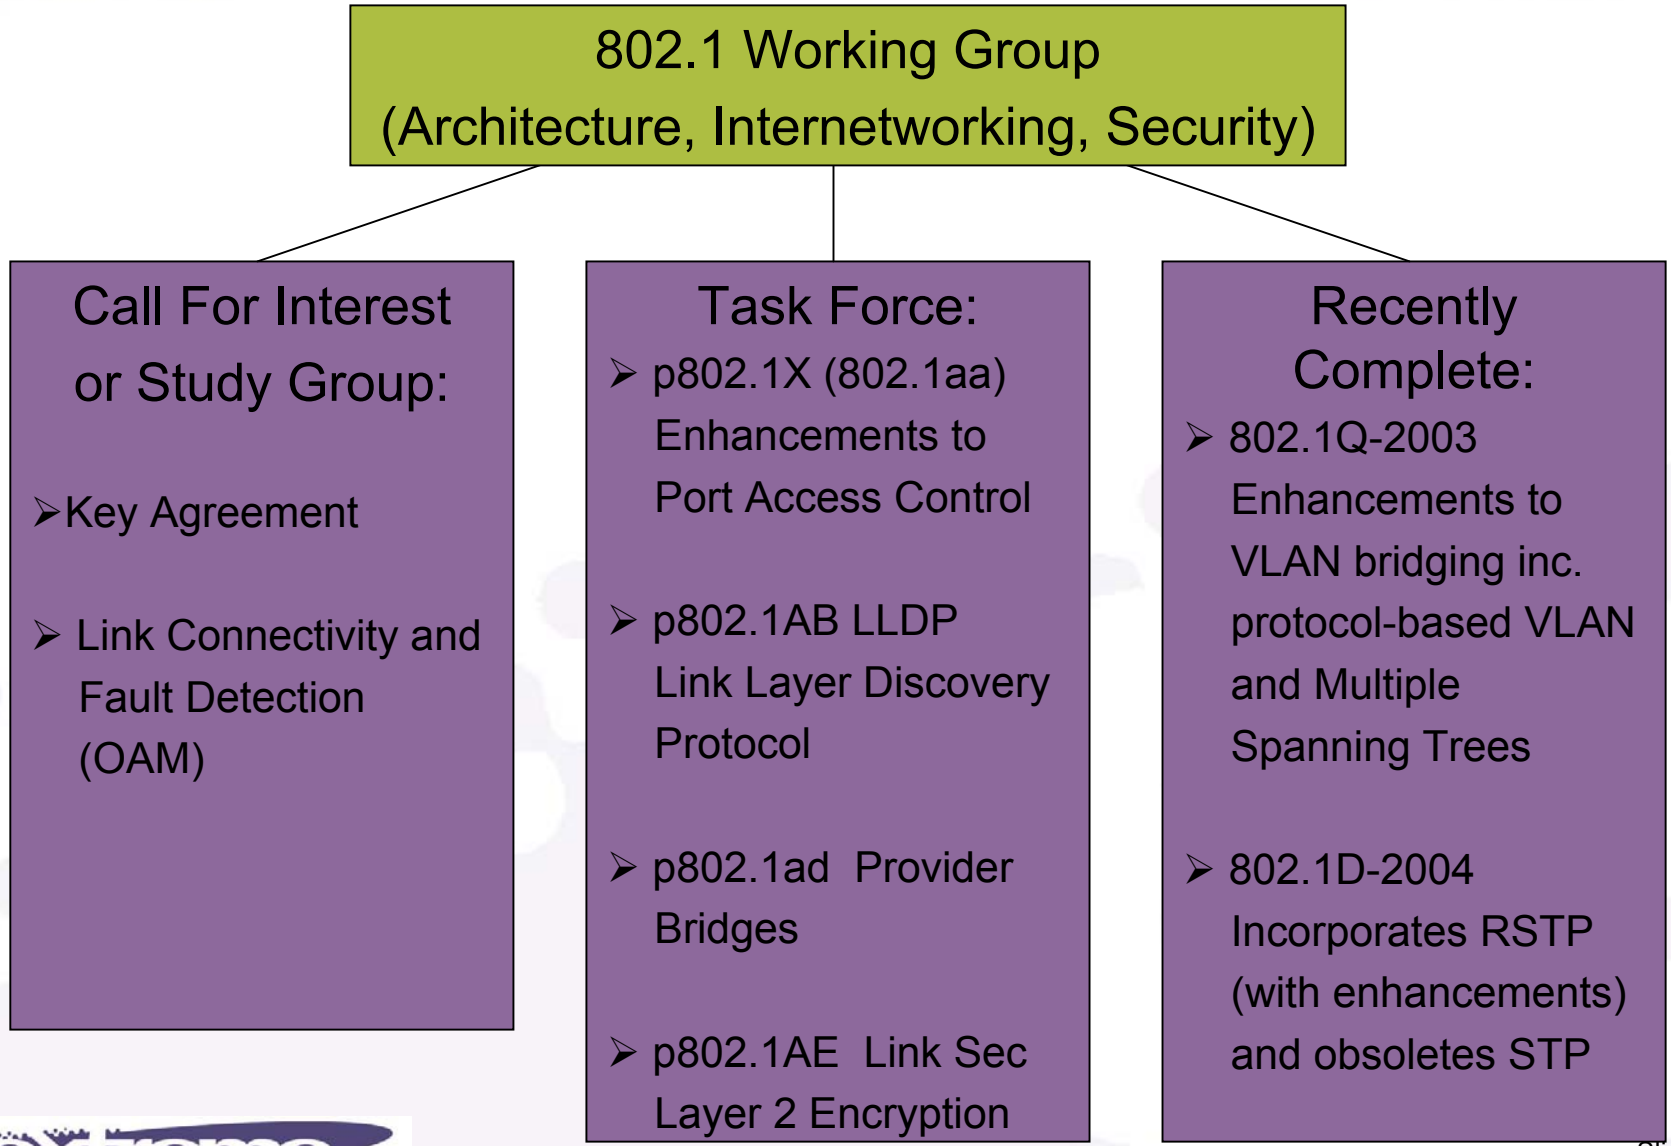

p802.1ad Provider Bridging PAR

- 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

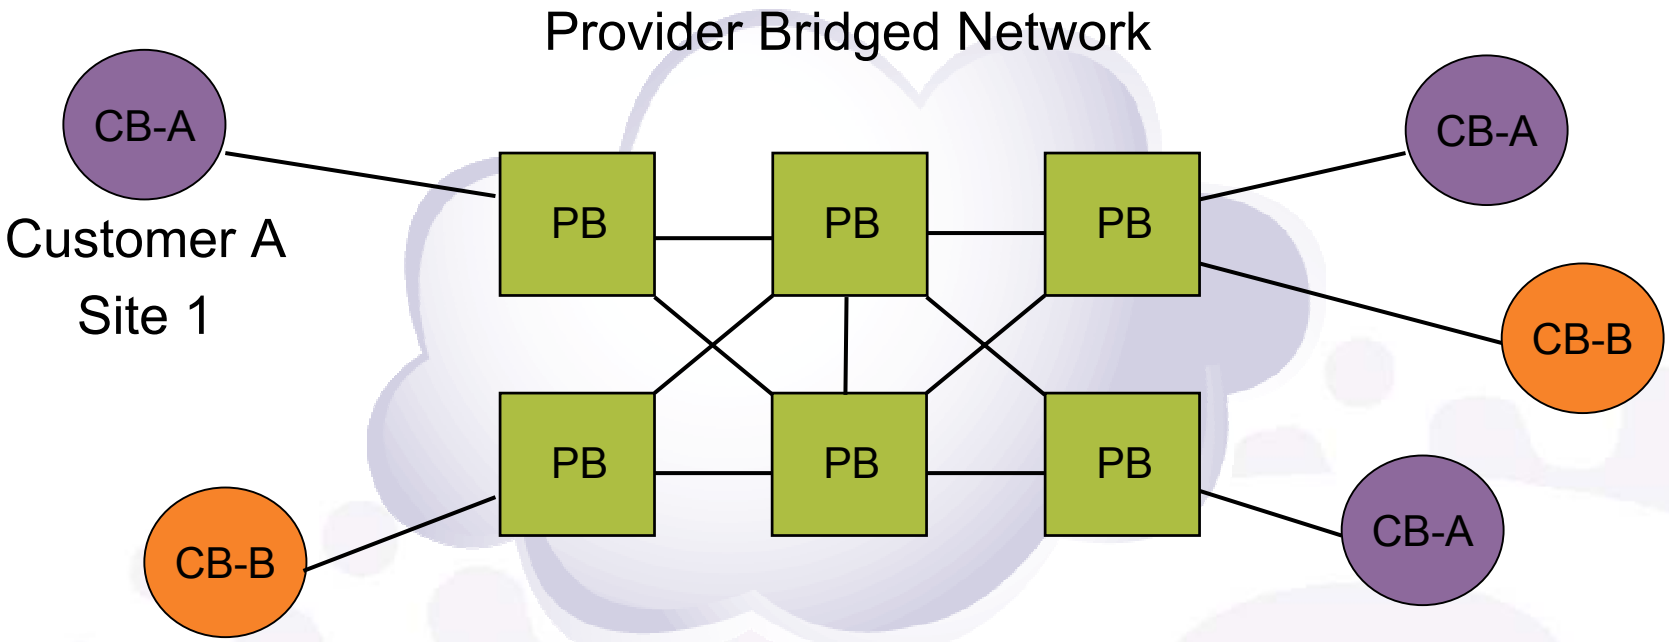

Goal: Transparently interconnect all of Customer A sites and all of Customer B sites while maintaining complete isolation between Customers A and B.

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 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.
     - Customers 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".
2. Use different reserved MAC addresses for Provider Spanning Tree BPDUs and Customer Spanning Tree BPDUs.
   - Provider network will tag Customer Spanning Tree BPDUs with a S-TAG, and transport them transparently.

802.1D – 1998 Transparent Bridge

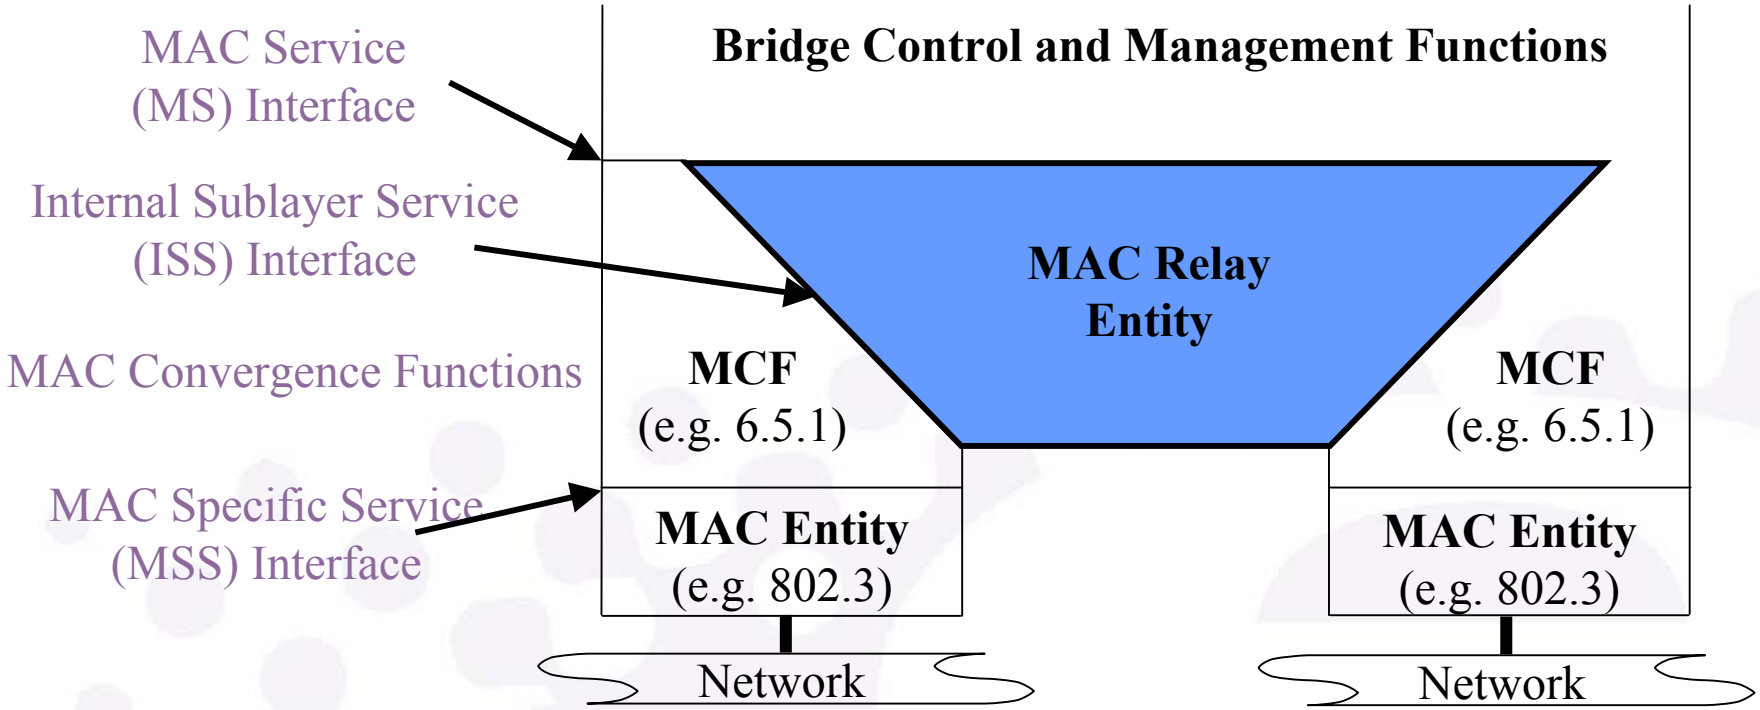

802.1Q – 1998 VLAN Bridge

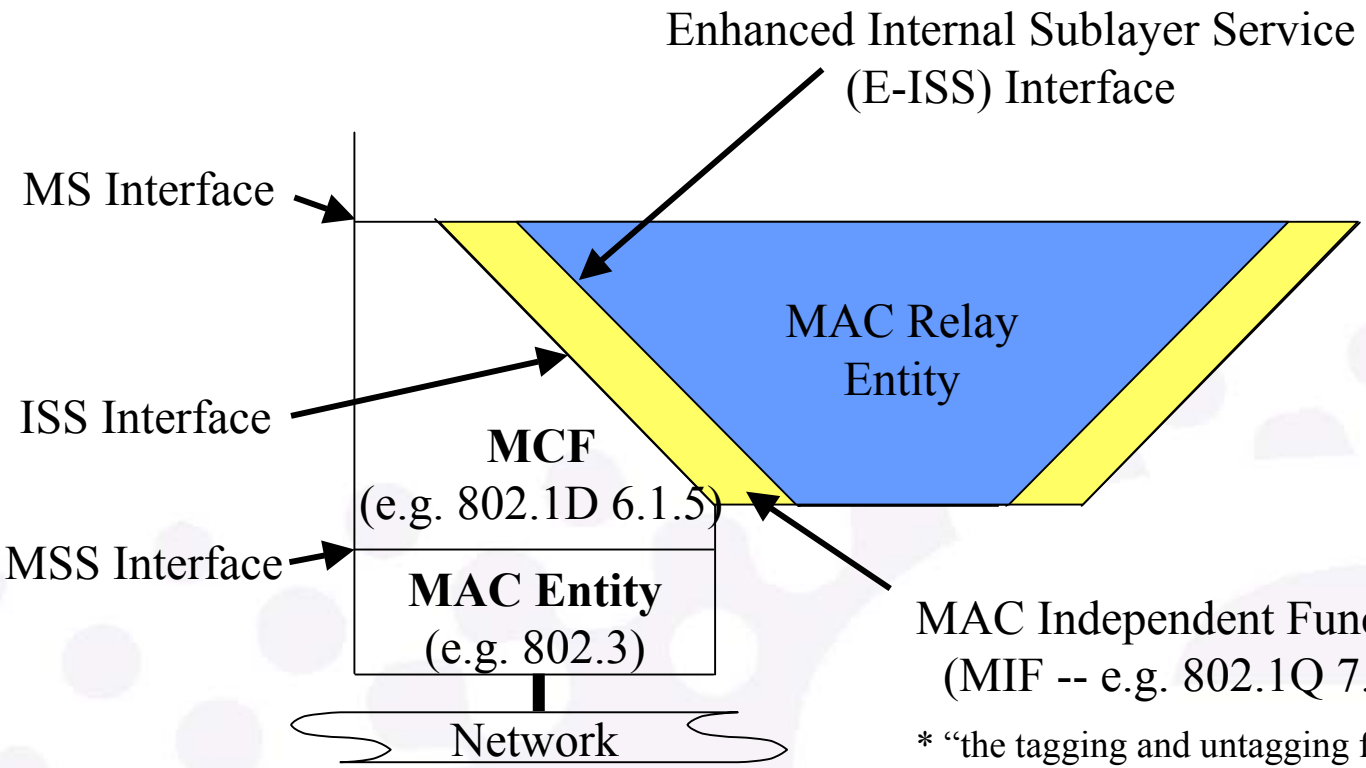

802.1ad – Provider Bridge

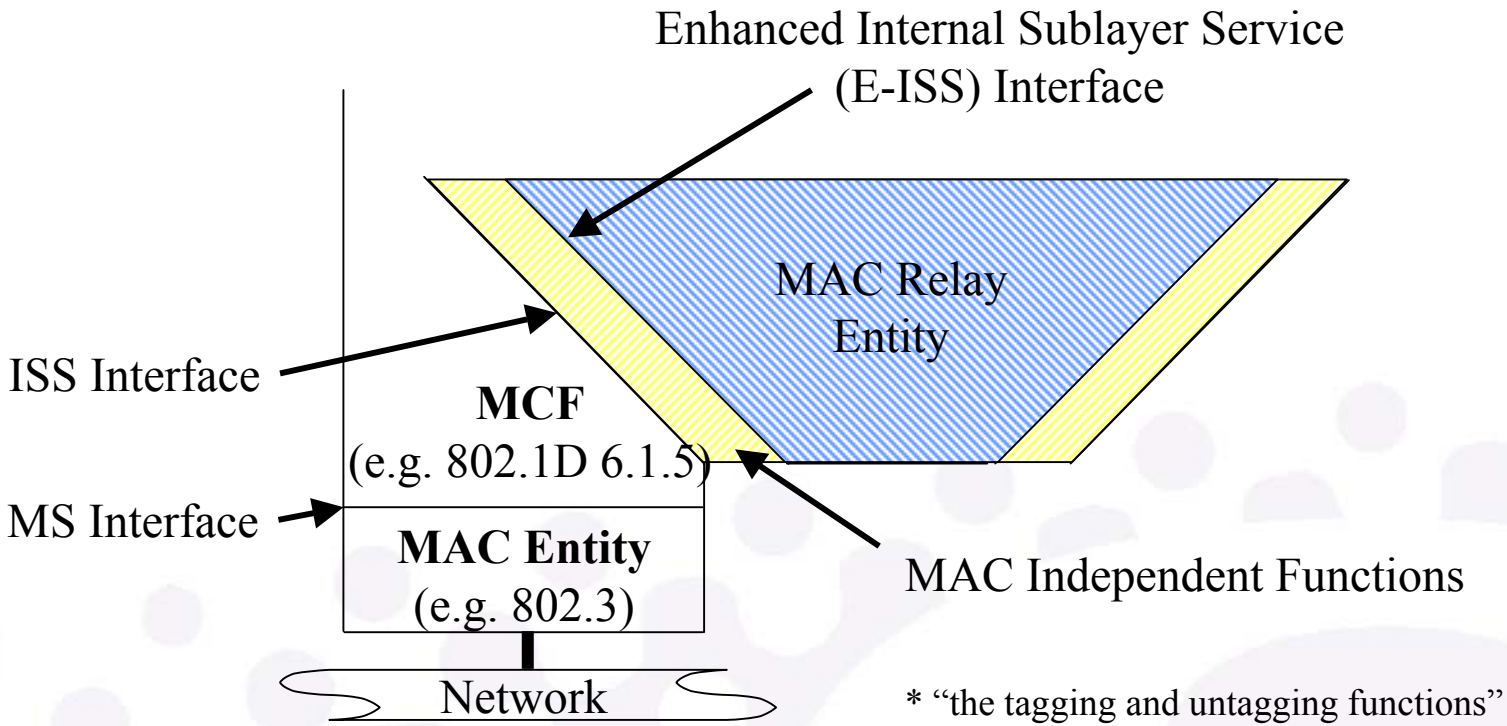

The mapping between the ISS and the E-ISS is the same as in 802.1Q 7.1.2 except that the operations are performed on a different tag – the Service VLAN Tag (S-TAG) rather than the Customer VLAN Tag (C-TAG).

Simple Provider Service

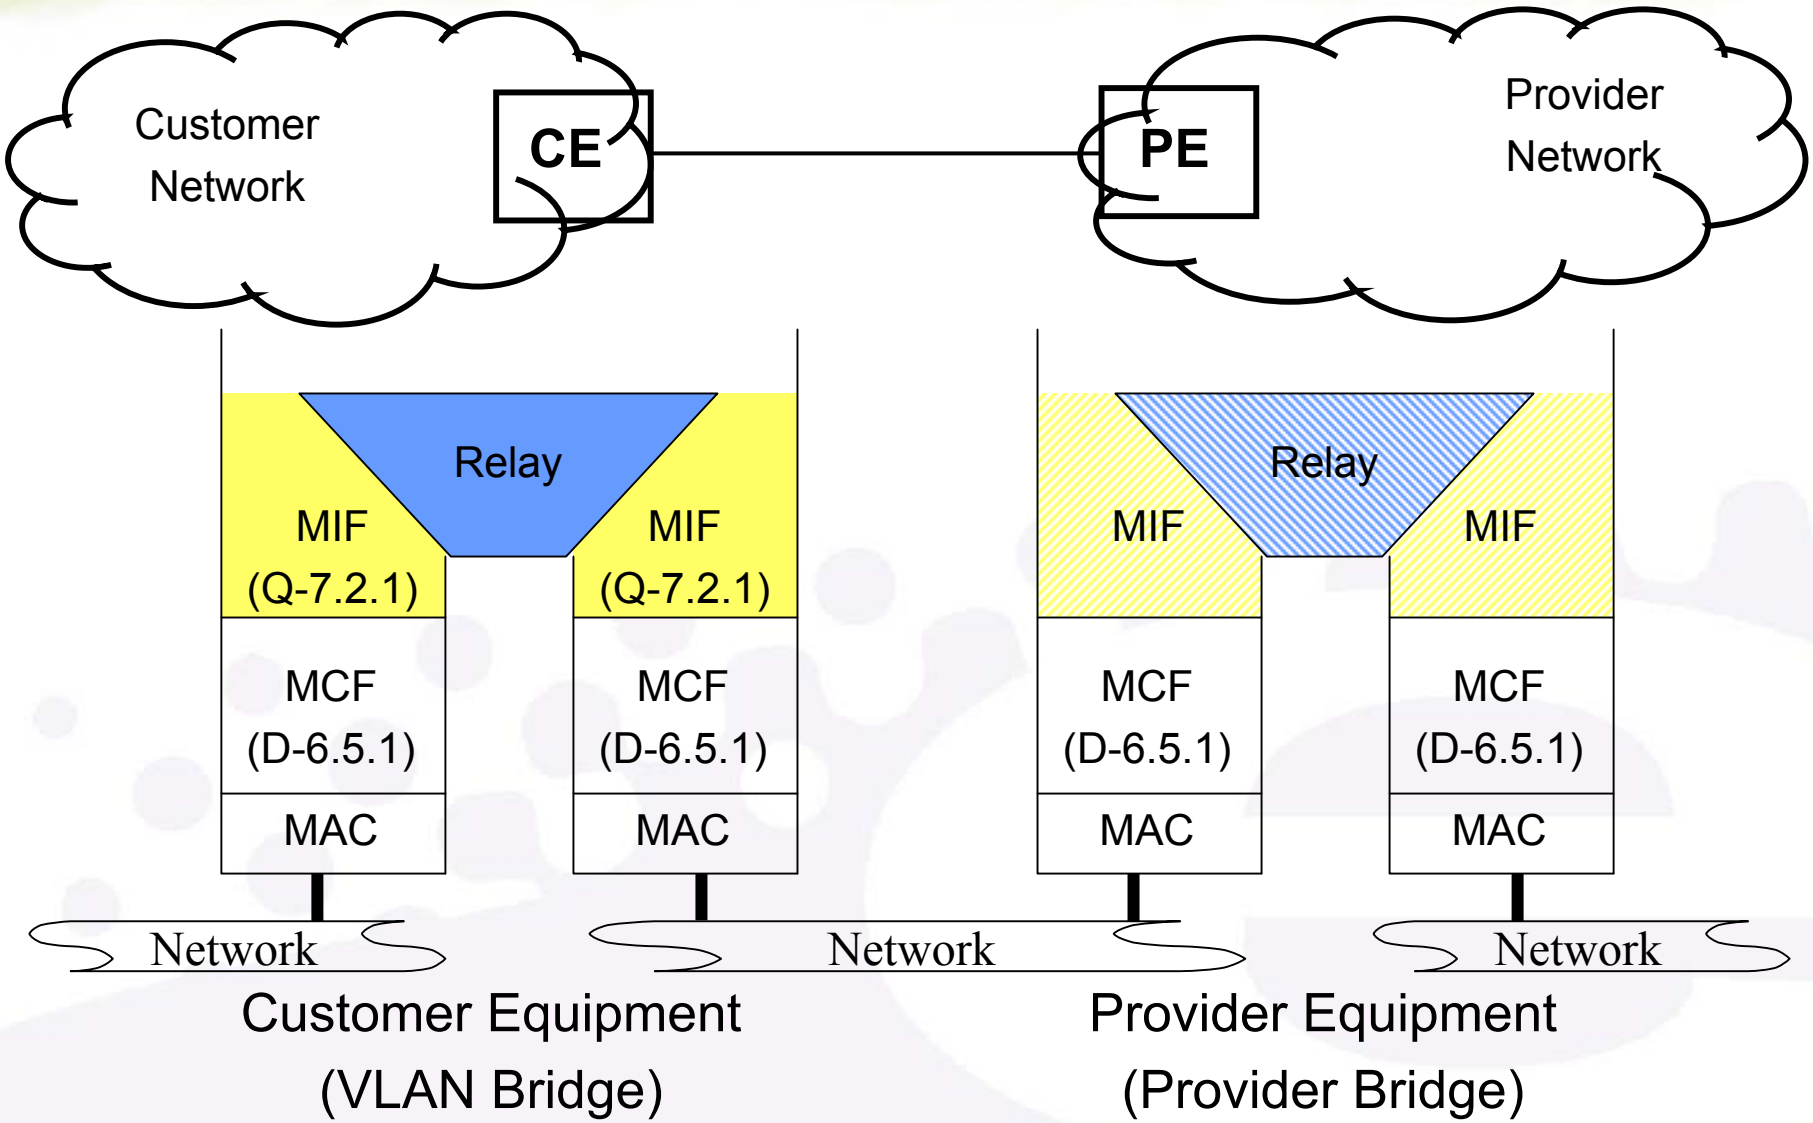

Simple Provider Service

- All the Provider Bridge does is insert a Service Tag in all frames received from the Customer Equipment.
- No changes are required to convert 802.1Q to 802.1ad beyond assigning a new Provider BPDU Address and a Service Tag Ethertype.
- 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?

The "Dual Bridge" Provider Edge Model

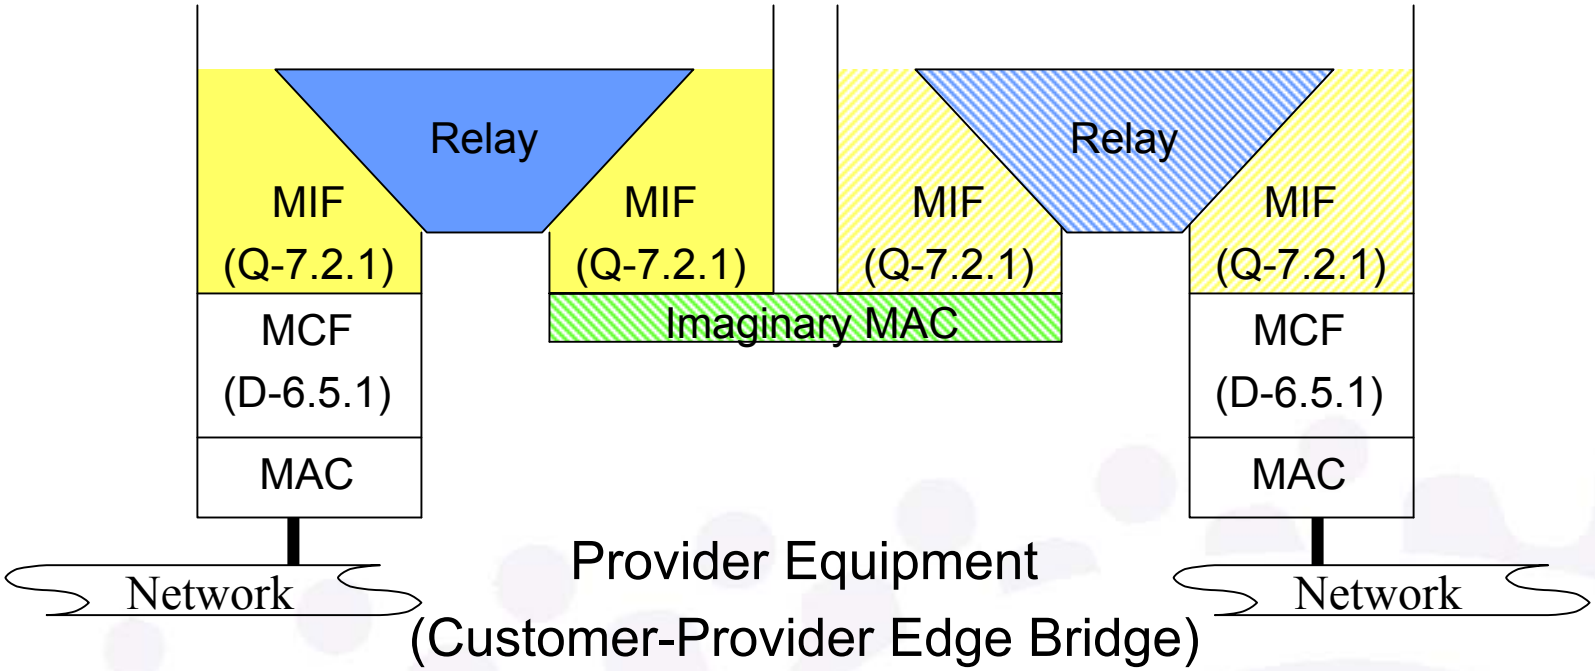

- Specify behavior of a Provider Edge as two bridges in one box.
  - Customer facing side operates on Customer VLAN Tags and BPDUs
  - Provider facing side operates on Provider Tags and BPDUs
  - Interconnect with an "imaginary port" per service instance

Service Multiplexing – Data Plane

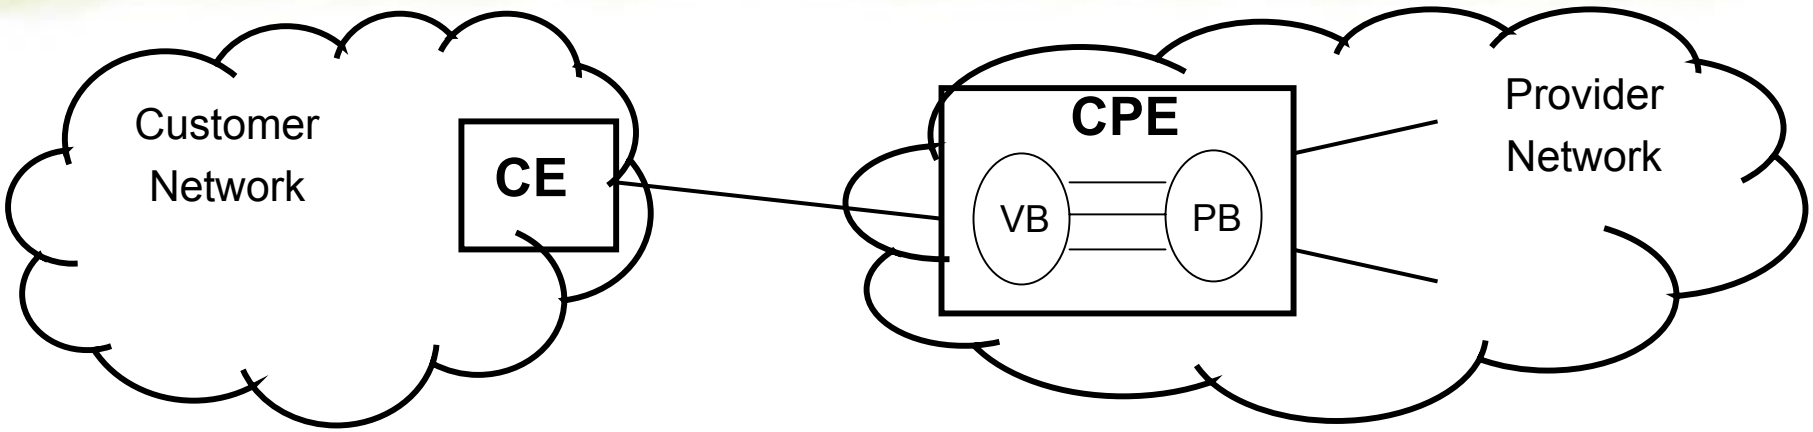

- Customer accesses 3 different Provider Services over a single physical link to the Customer-Provider Edge Bridge.
  - VLAN Bridge portion of CPE connects to Provider Bridge portion via 3 imaginary ports – one per service instance.
  - VLAN Bridge portion of CPE selects service based on Customer VLAN IDs by forwarding packets for each service to the appropriate imaginary port.
  - Provider Bridge portion creates Provider Tag using PVID assigned to the imaginary port.

Service Multiplexing – Control Plane

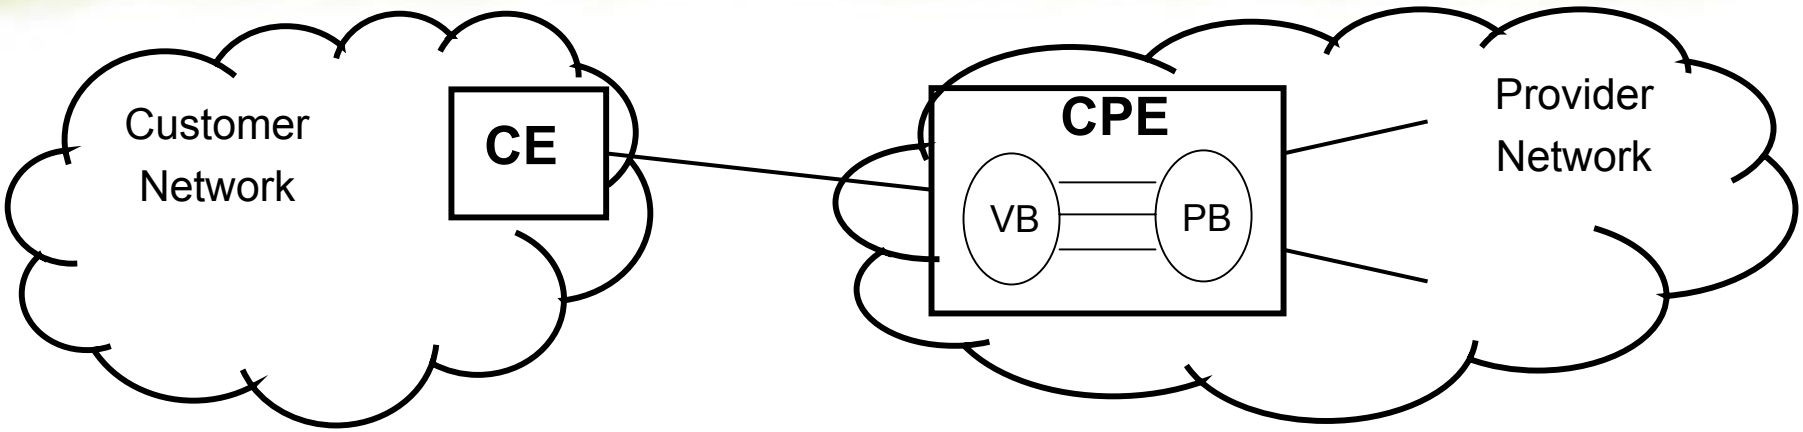

- Customer BPDUs must be transported across each service instance.
  - VLAN Bridge portion of CPE participates in Customer Spanning Tree – receives, processes, and transmits Customer BPDUs on each customer facing port and each imaginary port.
  - Provider Bridge portion "tunnels" Customer BPDUs from imaginary ports across the Provider Network.
  - Provider Bridge portion participates in Provider Spanning Tree.

Multiple Priorities

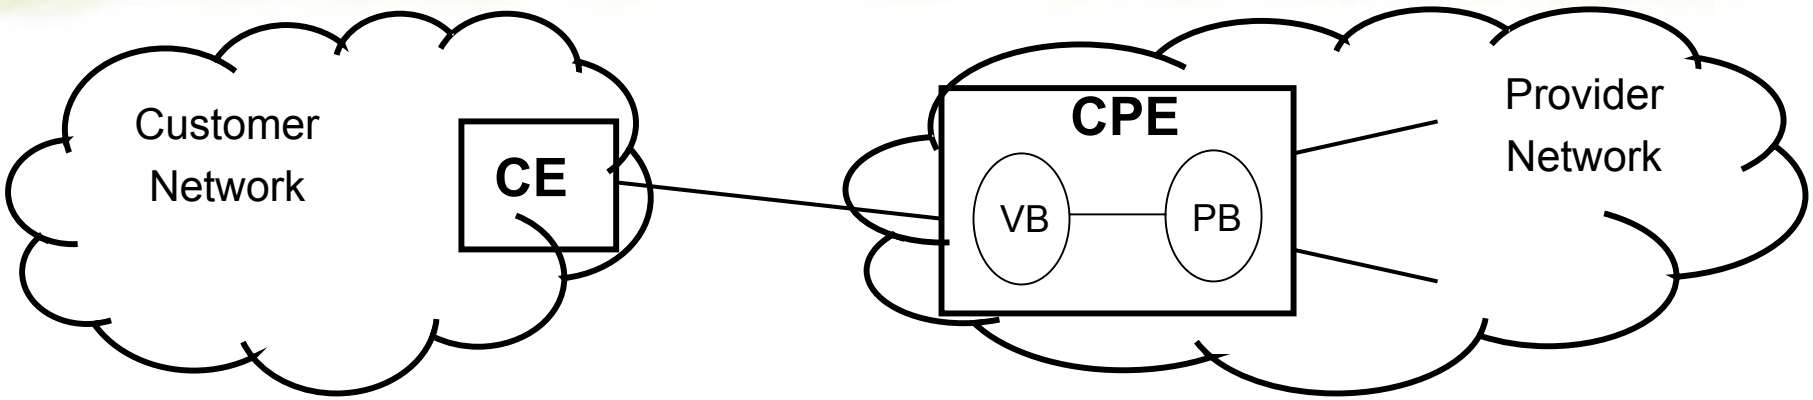

- Customer accesses single Provider Services that handles multiple priorities.
  - VLAN Bridge portion of CPE uses the user_priority field of the Customer VLAN tag to determine the access_priority for the imaginary MAC.
  - Imaginary MAC conveys the priority information from the VB to the PB.
  - Provider Bridge portion creates Provider Tag with a priority field “regenerated” from the priority conveyed by the imaginary MAC.
    - “Regeneration” allows PB to map Customer specified priorities to different priority levels on the Provider network.

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.
  - Still discussion as to whether other protocols (e.g. 802.3ad Link Aggregation) should be extended to work across Provider Networks.

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).

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.

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
  - Will extend CoS marking to include drop precedence
- Control Protocol
  - Separation of Customer and Provider Control Domains
  - Customer Spanning Tree Protocol packets transported through Provider network

Thank You