





>THIS IS THE WAY

Provider Backbone Bridging Networks

Paul Bottorff, Mark Holness, Norival Figueira, Michael Chen,
Dinesh Mohan, Glenn Parsons

Version 2.0

November 15, 2004 San Antonio

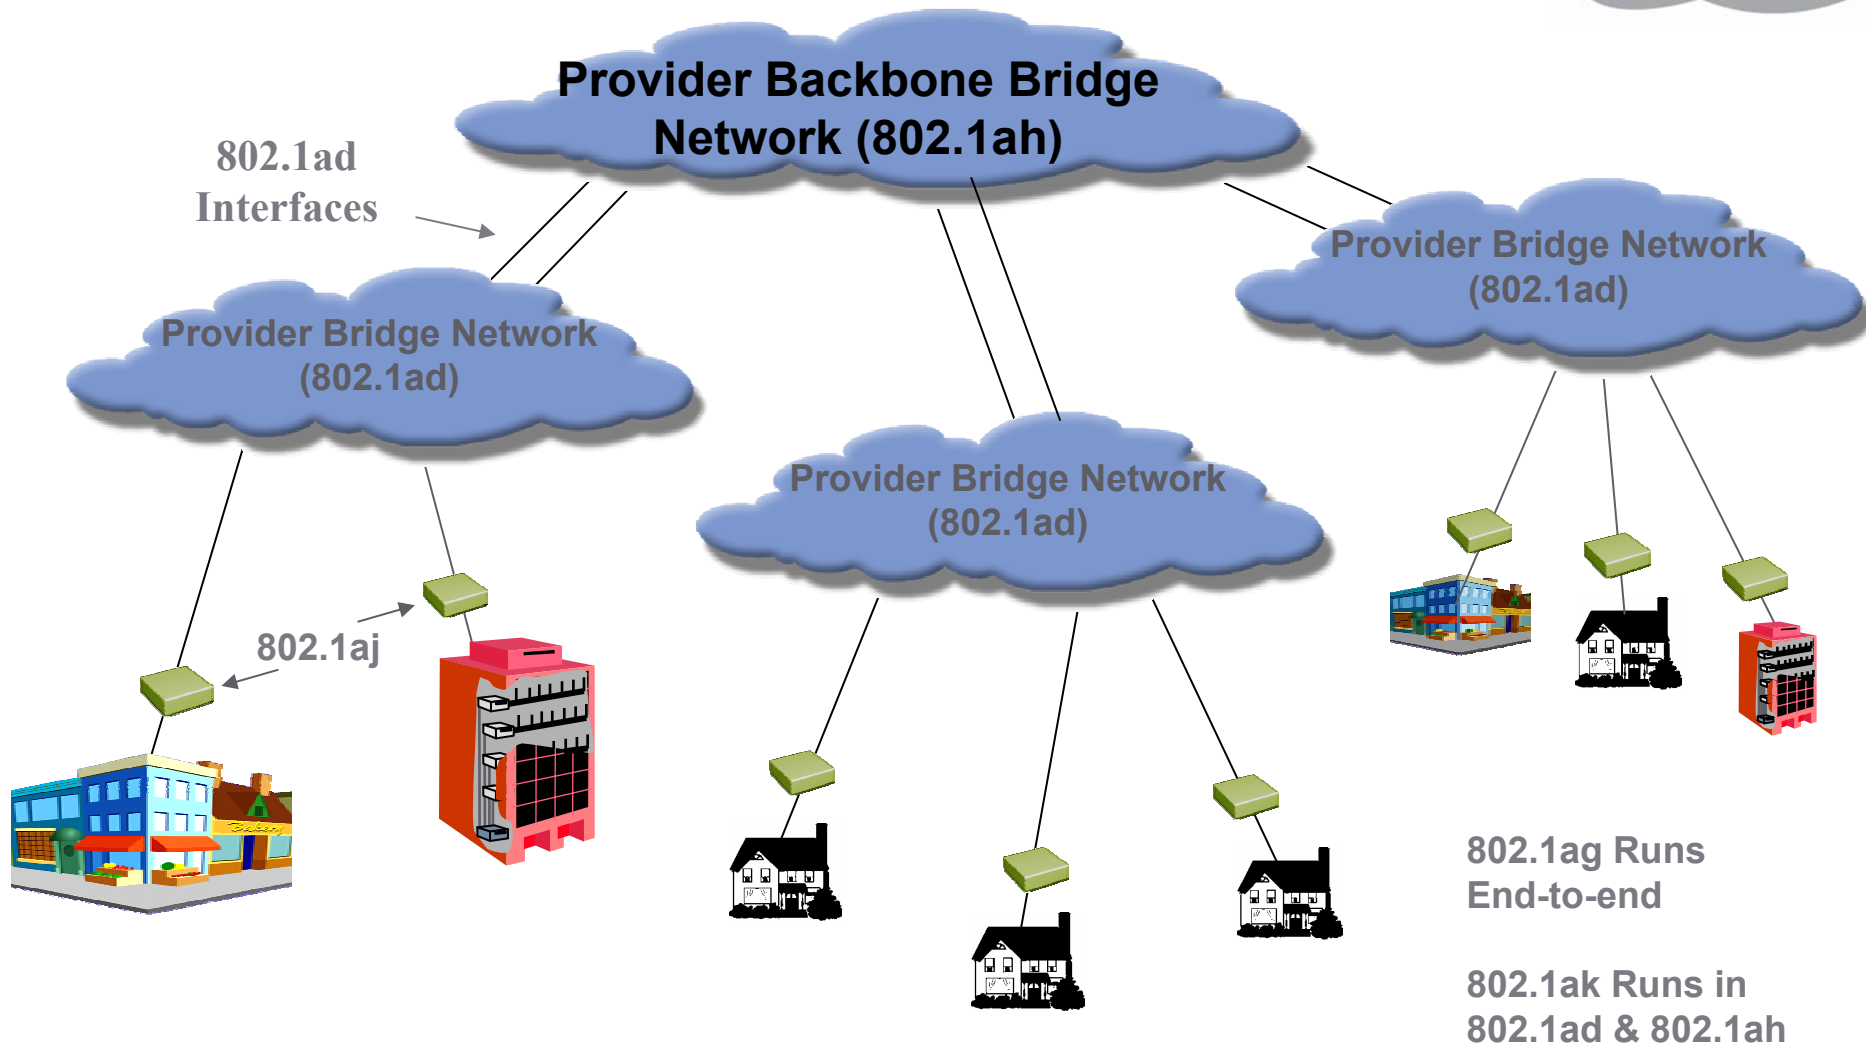
>THIS IS NOT NORTTEL



Where Are We At

- > At Ottawa interim we finished a draft Provider Backbone Bridge PAR (P802.1ah).
- > The P802.1ah PAR was forward to the 802 Exec within the 30 day rule as authorized by the Portland meeting.
- > At this meeting the 802.1 needs to approve the PAR.
- > If approved by 802.1 the Exec will then also be able to approve it at this meeting.

A Provider Bridge Scaling Solution “Provider Backbone Bridging”





802.1ah PAR

- > **1. ASSIGNED PROJECT NUMBER:** 802.1ah
- > **2. SPONSOR DATE OF REQUEST:** 2004-01-01
- > **3. TYPE OF DOCUMENT:** Standard
- > **4. TITLE OF DOCUMENT:** Standard for Local and Metropolitan Area Networks – Virtual Bridged Local Area Networks - Amendment 6: Provider Backbone Bridges
- > **5. LIFE CYCLE:** Full-Use
- > **6. TYPE OF PROJECT:** Amendment P802.1Q
- > **11. TYPE OF SPONSOR BALLOT:** Individual
Expected Date of Submission for Initial Sponsor Ballot: 2006-12-31
- > **12. PROJECTED COMPLETION DATE FOR SUBMITTAL TO REVCOM:** 2007-09-31

Scope



The scope of this standard is to define an architecture and bridge protocols compatible and interoperable with Provider Bridged(1) Network protocols and equipment allowing interconnection of multiple Provider Bridged Networks, to allow scaling to at least 2^{20} Service Virtual LANs, and to support management including SNMP.

Completion of this document contingent? Yes

1)This standard is designed to support Provider Bridges (IEEE P802.1ad).



Purpose & Reason

Purpose: This standard will complete the future work identified by P802.1ad, by providing a specific means for interconnecting Provider Bridged Networks. It will enable a Service Provider to scale the number of Service VLANs in a Provider Network by interconnecting the Service VLANs, and provide for interoperability and consistent standards based management.

Reason: This project is intended to facilitate the scaling of Provider Bridged P802.1ad networks using existing Bridged and Virtual Bridged LAN technologies. Despite user demand and initial deployment of LAN-based backbones for connecting P802.1ad networks, there is currently no interoperability between different vendors, nor a coherent management framework for different techniques. Most major carriers, who will be the users of this standard, are currently deploying LAN-based service networks that need to be scaled to meet the demands both of transition from existing leased line service and expansion of multipoint services.

Broad Market Potential



A standards project authorized by IEEE 802 shall have a broad market potential. Specifically, it shall have the potential for:

- a) Broad sets of applicability.
- b) Multiple vendors and numerous users.
- c) Balanced costs (LAN versus attached stations).

This project is intended to facilitate the scaling of Provider Bridged P802.1ad networks using existing Bridged and Virtual Bridged LAN technologies. Despite user demand and initial deployment of LAN-based backbones for connecting P802.1ad networks, there is currently no interoperability between different vendors, nor a coherent management framework for different techniques.

Most major carriers are currently deploying LAN-based service networks which need to be scaled to meet the demands both of transition from existing leased line service and expansion of multipoint services.

The costs related to this technology should be broadly similar to those of existing Bridging technology based on 802.1D/802.1w/802.1Q/802.1s.

Compatibility



IEEE 802 defines a family of standards. All standards shall be in conformance with the IEEE 802.1 Architecture, Management and Interworking documents as follows: 802. Overview and Architecture, 802.1D, 802.1Q, and parts of 802.1f. If any variances in conformance emerge, they shall be thoroughly disclosed and reviewed with 802. Each standard in the IEEE 802 family of standards shall include a definition of managed objects which are compatible with systems management standards.

This standard will be compatible with 802.1Q as amended by P802.1ad and P802.1ag. This project will be compatible with existing 802.1 Architecture, Management and Interworking standards.

The Provider Backbone Bridge will rely on extensions to 802.3 frame size for additional header space. Work on frame size extension is currently under study at 802.3.



Distinct Identity

Each IEEE 802 standard shall have a distinct identity. To achieve this, each authorized project shall be:

- a) Substantially different from other IEEE 802 standards.
- b) One unique solution per problem (not two solutions to a problem).
- c) Easy for the document reader to select the relevant specification.

There is no other IEEE standard or project that allows scaling of a Provider Bridge network to support large numbers of Service VLANs. No existing solution provides a multipoint LAN backbone for interconnection of Provider Bridges. The document reader will have an easy reference to scaling of Provider Bridge networks.



Technical Feasibility

For a project to be authorized, it shall be able to show its technical feasibility. At a minimum, the proposed project shall show:

- a) Demonstrated system feasibility.
- b) Proven technology, reasonable testing.
- c) Confidence in reliability.

The proposed standard will be based on existing, proven, standardized, Bridged LAN and Virtual Bridged LAN technology. These technologies are widely implemented and highly reliable.



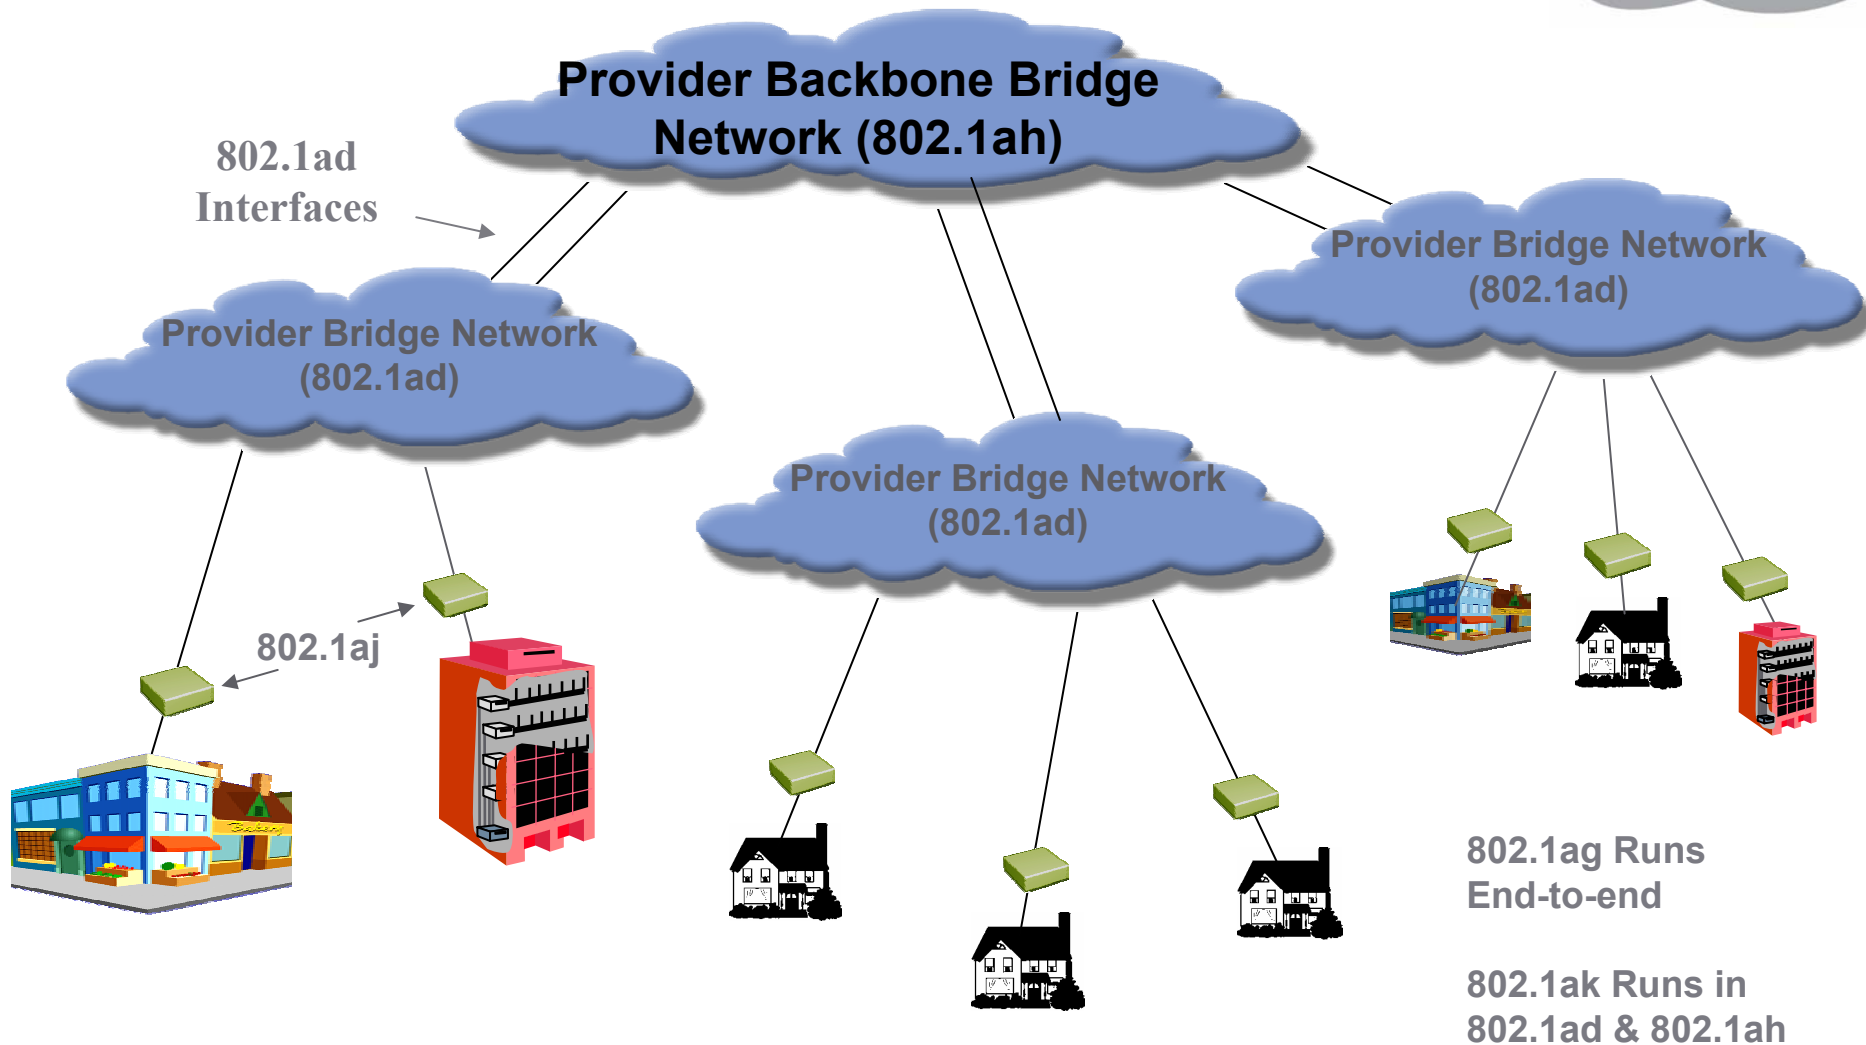
Economic Feasibility

For a project to be authorized, it shall be able to show economic feasibility (so far as can reasonably be estimated), for its intended applications. At a minimum, the proposed project shall show:

- a) Known cost factors, reliable data.
- b) Reasonable cost for performance.
- c) Consideration of installation costs.

The technology that will be developed in the proposed standard will not differ significantly from the economic factors associated with existing Bridged LAN and Virtual Bridged LAN technologies. The cost factors for Virtual Bridged LAN technology are favorable when compared to existing provider networks based on MPLS or SONET.

A Provider Bridge Scaling Solution “Provider Backbone Bridging”



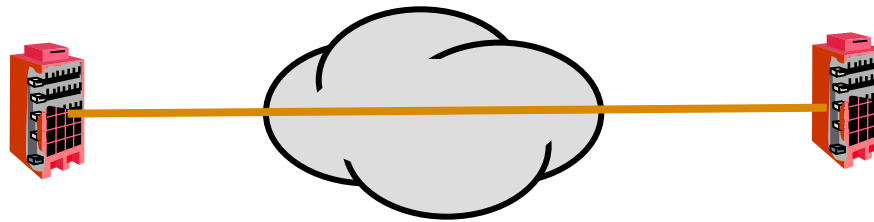


Ethernet Service Types

MEF Ethernet Virtual Connections (EVCs)

E-LINE

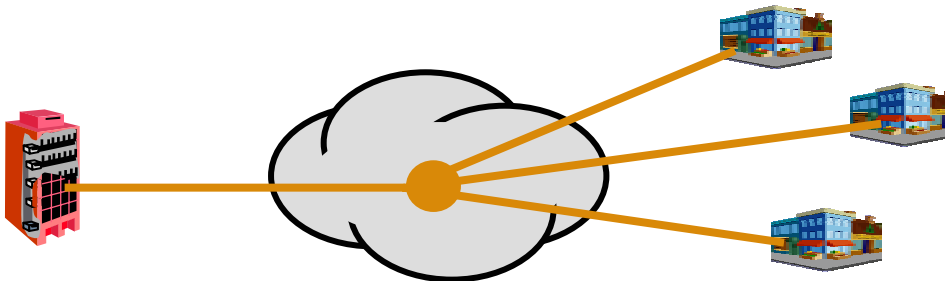
Router Mesh



Pt-Pt, Like
Duplex Ethernet
Any-to-any

E-TREE

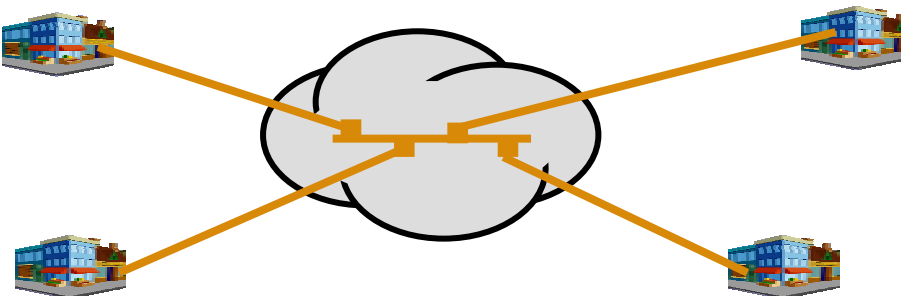
Hub & Spoke



Pt-MPt, Like
EPON Ethernet,
Root-to-Leaf and
Leaf-to-Root

E-LAN

Multi-Site



MPt, Like VLAN,
Any-to-any



E-LINE Dominates Today

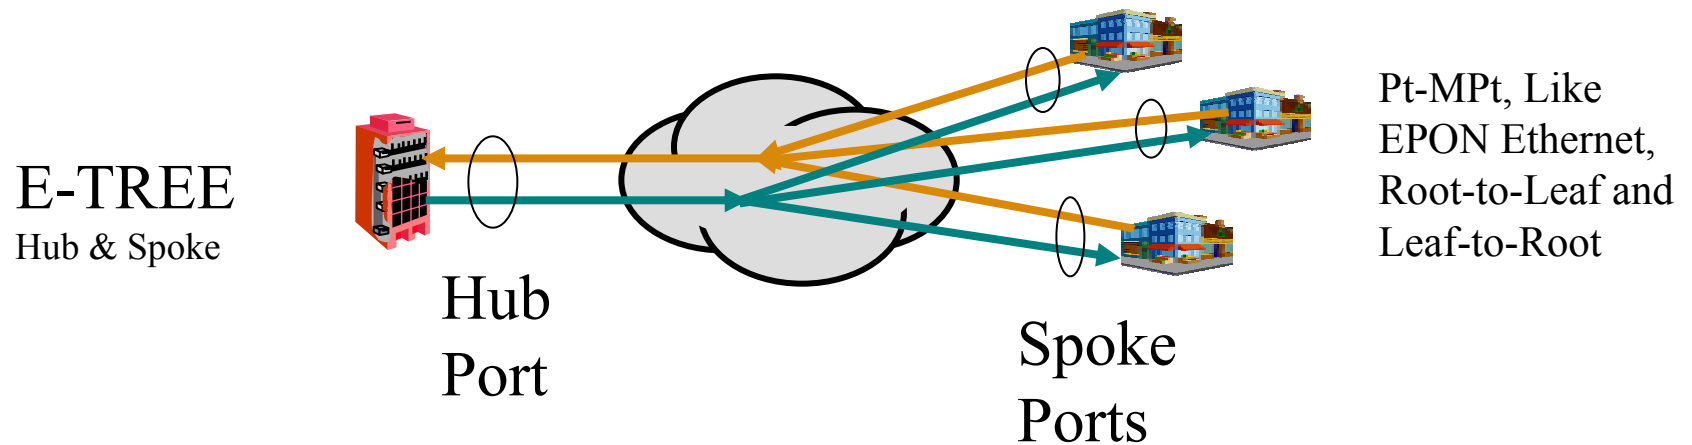
- > E-LINE is a natural leased line replacement for subscribers
 - Ethernet leased lines offer high bandwidth
 - Lines provide bandwidth on demand
 - Interfaces are compatible with off the shelf Ethernet switches/routers
 - Best for router mesh
- > E-LINE provides natural migration for carriers
 - Consistent with current operations model
 - Allows carrier equipment reductions
 - Bill models can follow well understood FR services
 - Current QoS models allow both traffic control and service monitoring of E-LINE service offerings
 - Service OAM models for E-LINE are relatively straightforward
- > Each E-LINE service instance requires 1 S-VLAN



E-TREE Ideal For ISP Connect

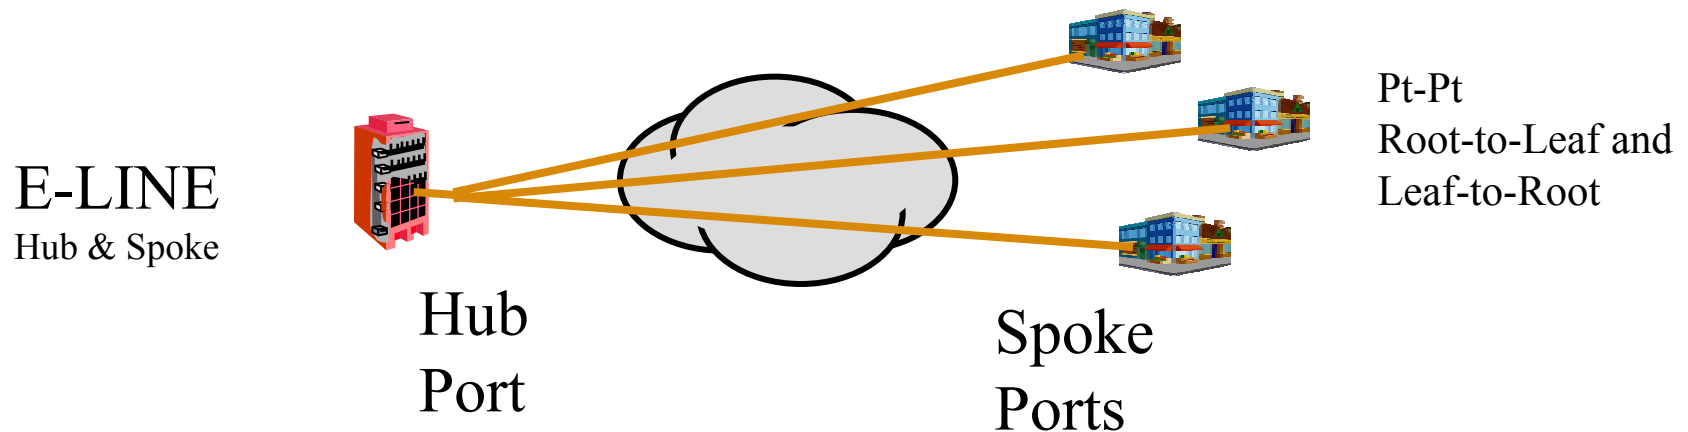
- > E-TREE Future Service With Great Promise
 - Useful as a multiplexed connection to an application service provider like an ISP
 - Service is unlike traditional Ethernet since leaf nodes can not talk with each other
- > E-TREE has deployment issues
 - No clear billing model
 - For instance if one leaf is disconnected is the circuit down?
 - What is the distance of the tree?
 - OAM management not fully understood
 - QoS model non-existent, SLAs can only provide Best Effort

E-TREE S-VLAN Mapping



- > Each E-TREE service instance requires 2 S-VLANs
- > Both S-VLANs comprising an E-TREE S-VLANs are unidirectional
- > The S-VLANs of an E-TREE service instance are typically multiplexed on the same port

Some Carriers Will Use E-LINE in Hub and Spoke Arrangement



- > Hub port would usually be multiplexed to allow the multiple Pt-Pt attachments.
- > Each E-LINE is a separate managed S-VLAN
- > This arrangement allows use of E-LINE management, billing, and QoS
- > Many more S-VLANs are required



E-LAN Many Future Applications

- > E-LAN is deployed for broad connectivity in select network
 - Interconnect of multiple corporate sites
 - Multi-player gaming
 - Ubiquitous any-to-any connectivity
 - E-LAN has many future applications
- > E-LAN has deployment issues
 - Deployments are very spotty
 - Unclear billing model
 - How is availability defined?
 - No definitions for QoS or performance measurement
 - What is the distance of a E-LAN
 - Unclear management models
 - Unlike existing carrier service offerings
- > Each E-LAN service instance is a single S-VLAN



Prototypical Major Metro Area

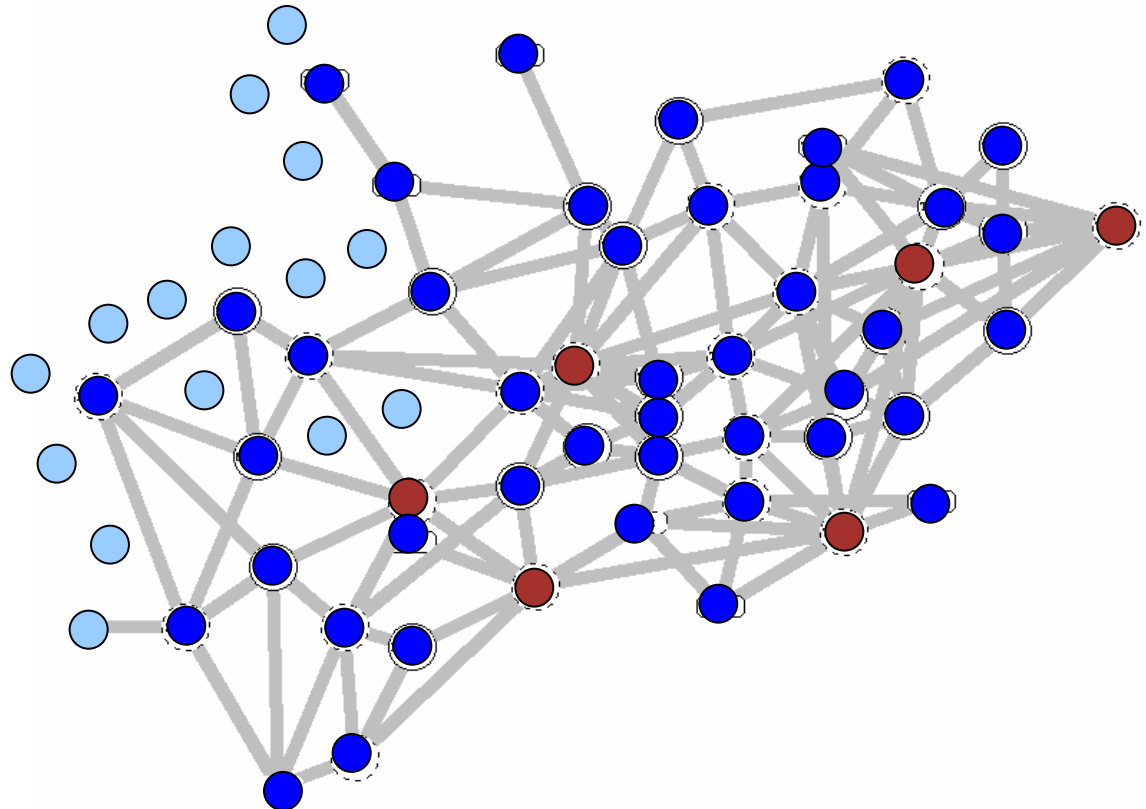
- > Business Subscriber Population 100K-2M
 - San Jose Yellow Pages ~100K businesses
 - The SF Bay Area lists ~1M businesses
- > Large Business Sites 500-5,000
- > Residential Subscriber Population 1M-20M
- > Leased Line Density 10K-200K
 - Roughly 1/10 Yellow Page Listings
- > Application Service Provider Sites 100-2000
 - Large APSPs sites may service residential

Major MSA Networks

Typical SP	Access 	Business CLE	Small Office 	Medium Office 	Large Office 
Network Scale	>10,000 Remotes	>10,000 CLEs	>500 COs	100-200 COs	10-60 COs
Metro Scale	>4,000 Remotes	>1,000 CLEs	>50 COs	>20 COs	>4 COs

Typical Metropolitan Serving Area – MSA

- > MSA example shown
- > ASIA/PAC more CO/MSA
- > Europe less CO/MSA



Support 1,000,000 Service Instances



- > Must be able to support E-LINE service for leased line replacement for entire MSA
 - This is the way Ethernet is entering the markets
 - The objective is 200K E-LINE instances
- > Must support E-LINE for APSP to Subscribers
 - Not all service providers will allow E-TREE because of deployment problems
 - The objective of an additional 200K E-LINE is adequate for transition until E-TREE
 - Requirements for around 10K E-TREE instances
 - Requires 20K S-VLANs
- > Must support E-LAN for APSP and B-B
 - Advanced peer applications
 - Number of service instances speculative, however could be large
- > Totals
 - 200K E-LINE S-VLANs for leased line replacement
 - 200K E-LINE S-VLANs for APSP
 - 20K E-TREE S-VLANs
 - ? E-LAN Service Instances
- > Designing Into A Corner Will Not Instill Confidence In Future
 - Set Objectives to at least 1,000,000 service instances E-LINE, E-TREE, E-LAN
 - E-LAN service will eventually become important for coupling small groups
 - Allow E-TREE and E-LAN service scaling to at least 100,000 for future growth



Proposed Project Objectives

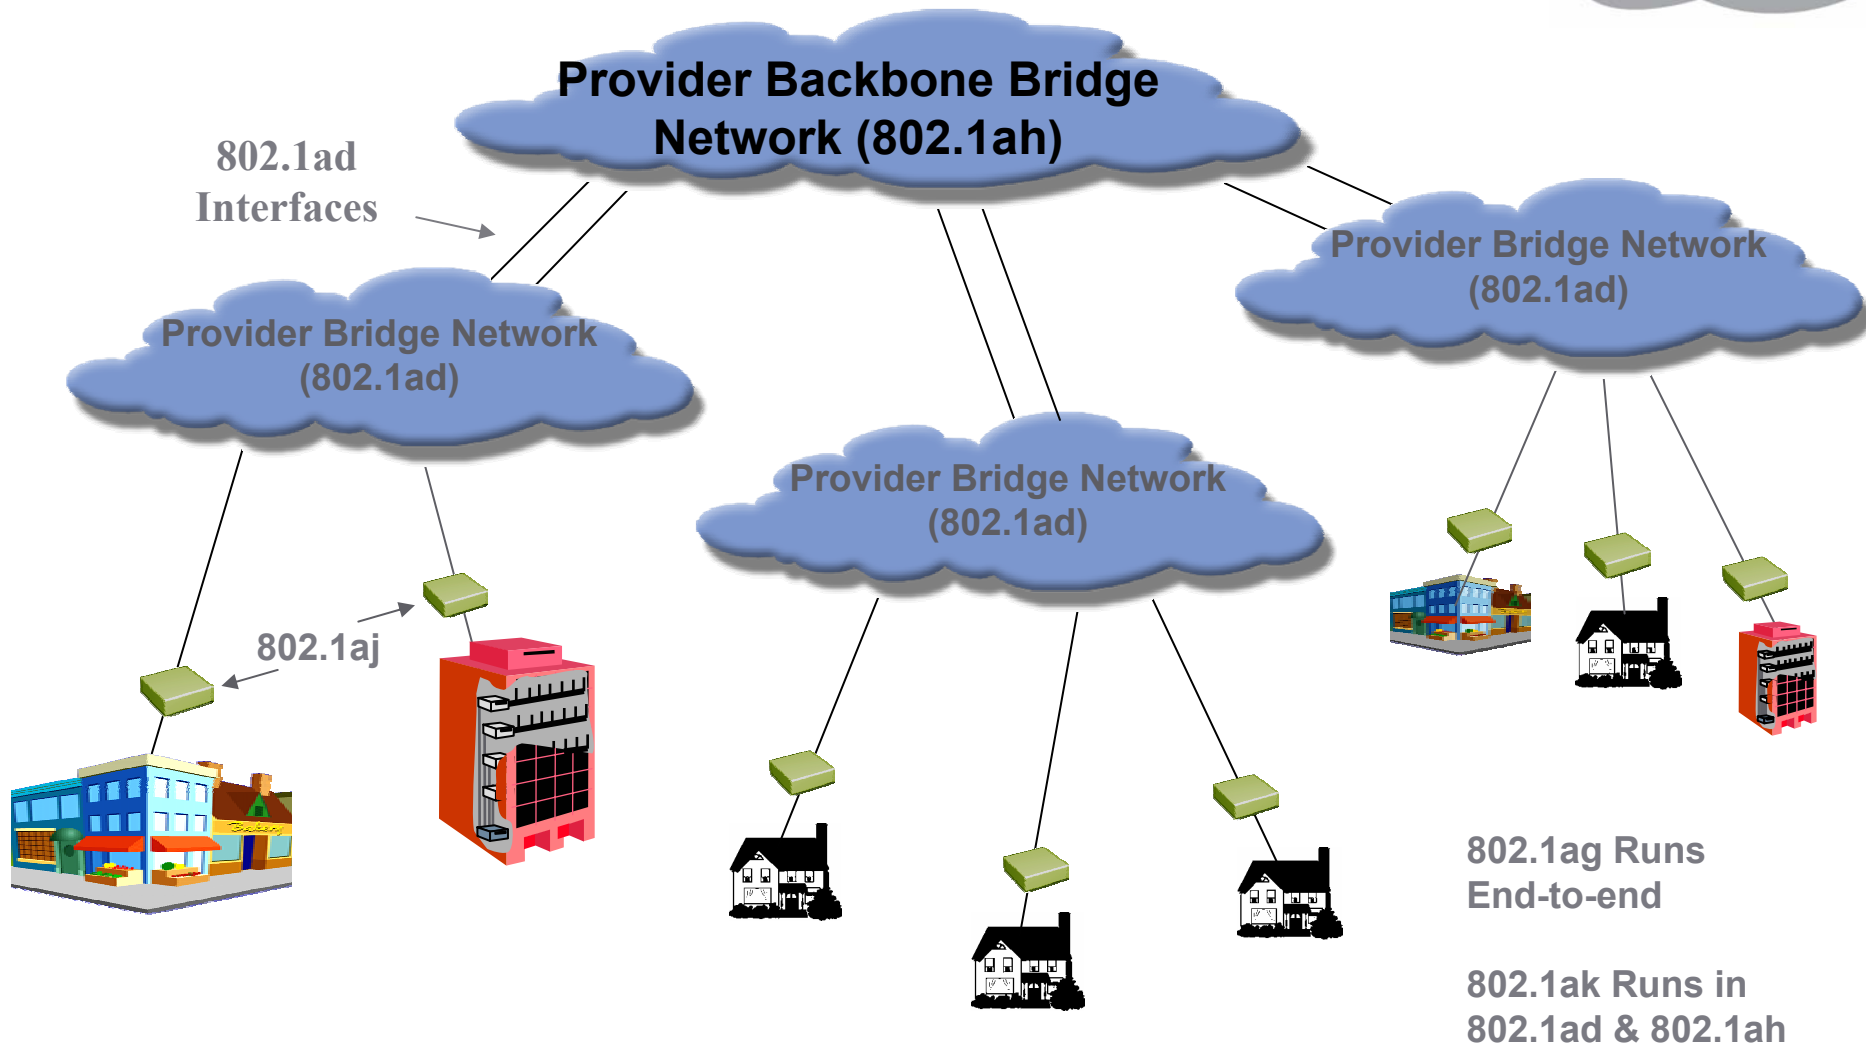
- > Interconnect Provider Bridge (802.1ad) Networks in a manner that allows scaling of the Carrier Bridged Network to support at least 2^{20} S-VLANs
- > Support at least 2^{16} multipoint S-VLANs
- > Interconnect at least 256 Provider Bridged Networks



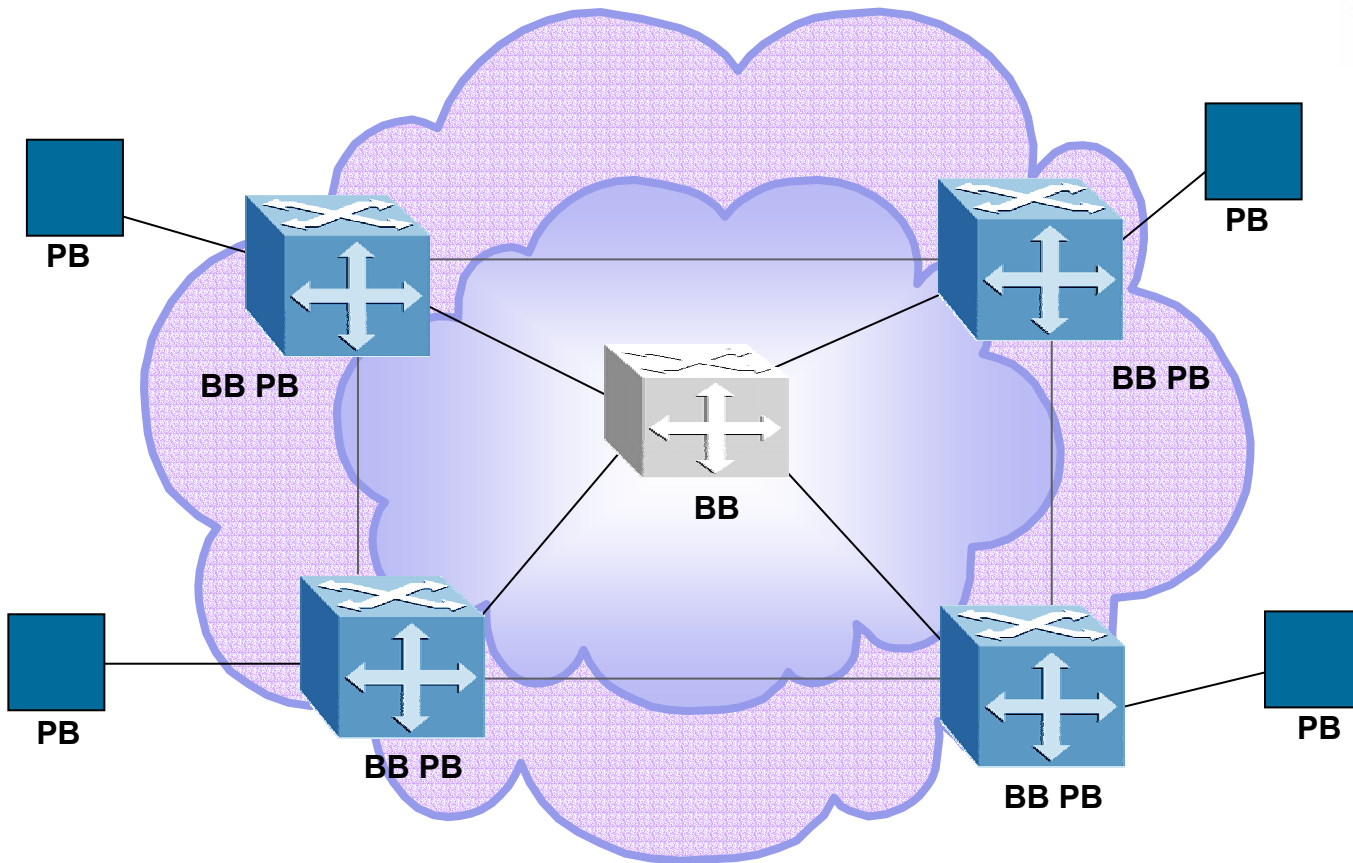
Provider Backbone Bridge Technology Principles

A Provider Bridge Scaling Solution

“Provider Backbone Bridging”



Provider Backbone Bridge Network



- **PB**: Provider Bridge (as defined by 802.1ad)
- **BB PB**: Provider Backbone Bridge Edge
- **BB**: Provider Backbone Bridge



Terminology

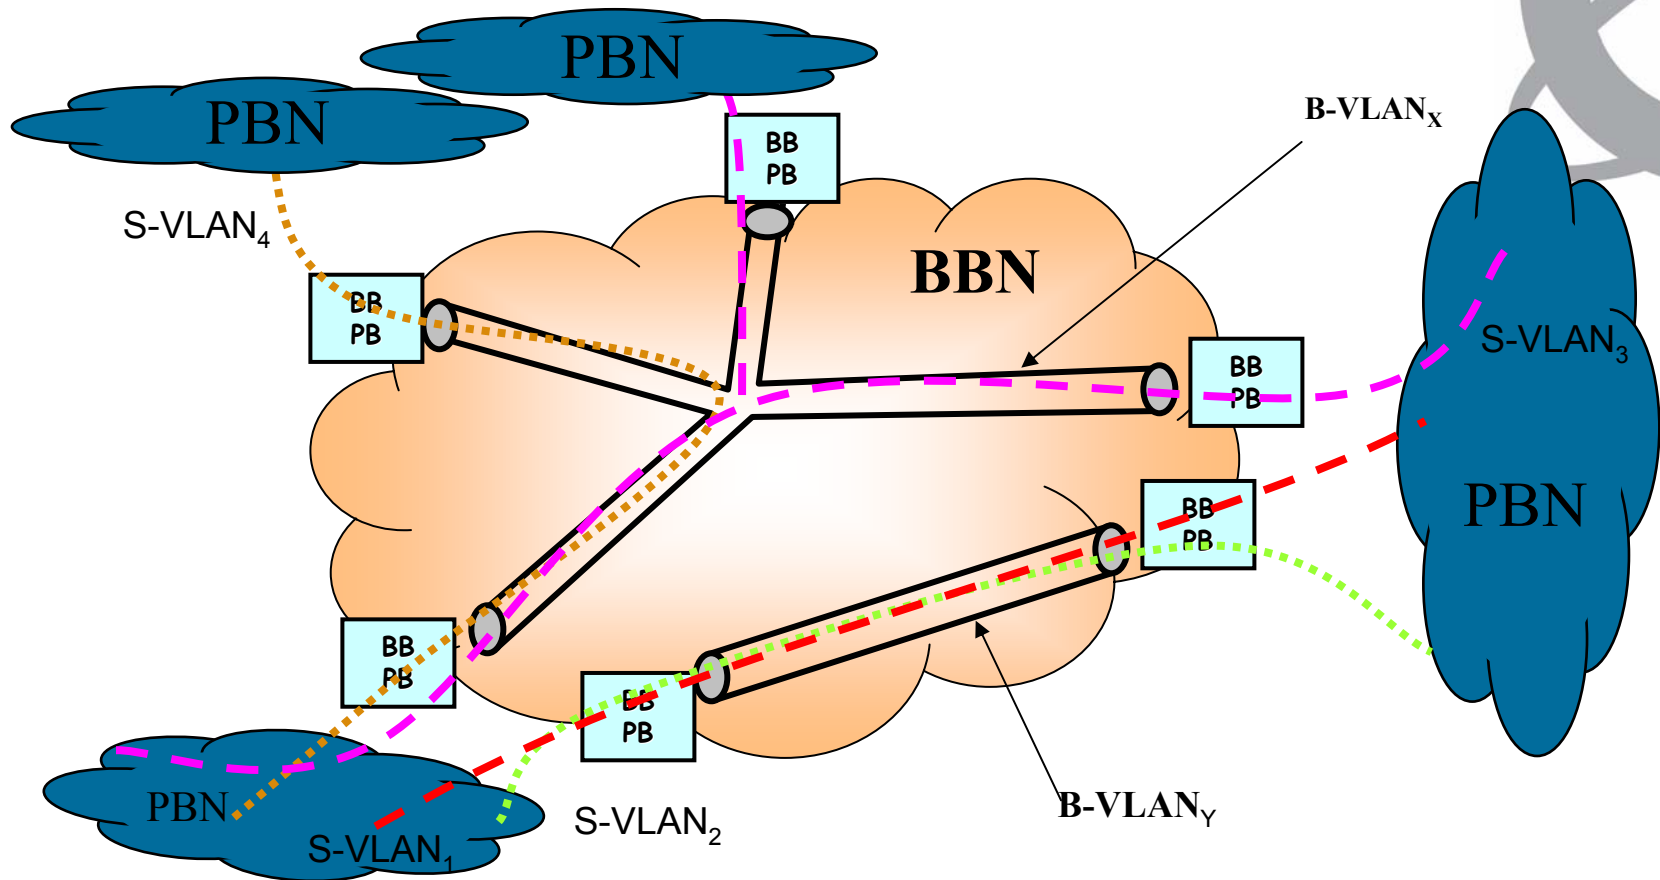
> IEEE 802.1ad Terminology

- C-TAG Customer VLAN TAG
- C-VLAN Customer VLAN
- C-VID Customer VLAN ID
- S-TAG Service VLAN TAG
- S-VLAN Service VLAN
- S-VID Service VLAN ID

> Additional Provider Backbone Bridge Terminology

- XS-TAG Extended Service VLAN TAG (I-TAG/ES-TAG)
- XS-VID Extended Service VLAN ID (SID/ES-VID)
- C-MAC Customer MAC Address
- B-MAC Backbone MAC Address
- B-VLAN Backbone VLAN (tunnel)
- B-TAG Backbone TAG Field
- B-VID Backbone VLAN ID (tunnel)

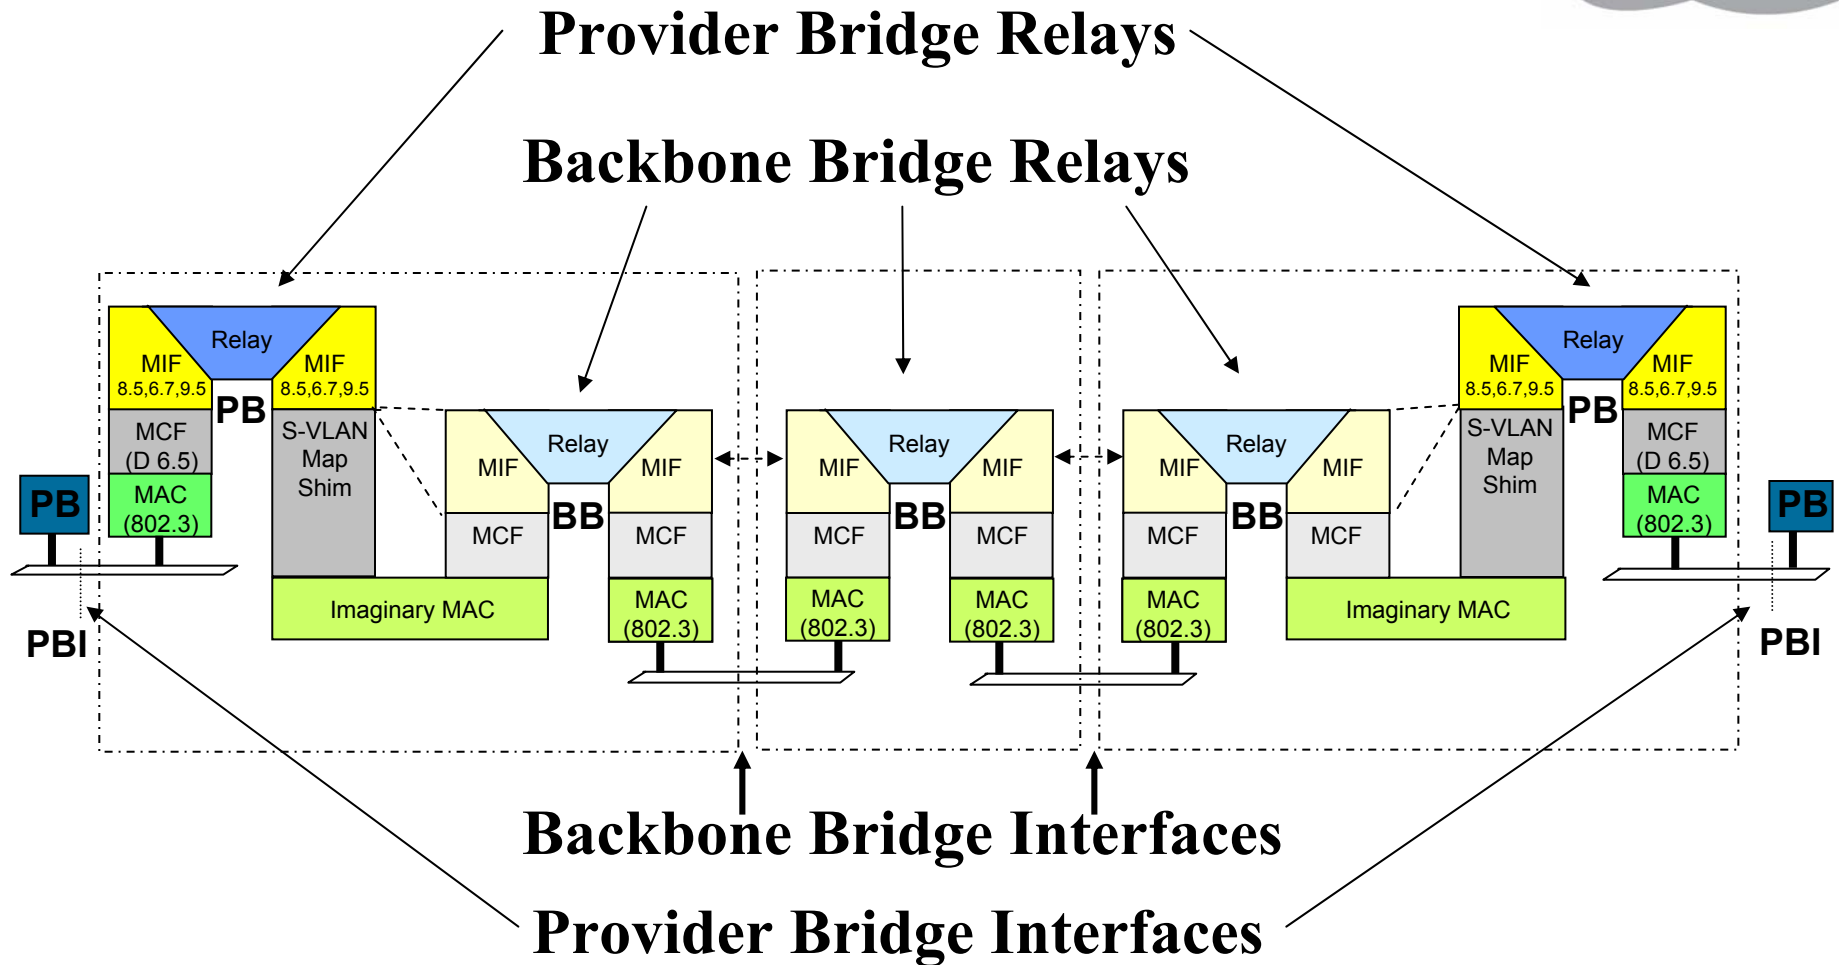
BBN Provides Multi-Point B-VLANs Between PBNs



- **BB PB:** Provider Backbone Bridge Edge
- Each B-VLAN carries many S-VLANs
- S-VLANs may be carried on a subset of a B-VLAN (i.e. all P-P S-VLANs could be carried on a single MP B-VLAN providing connection to all end points).



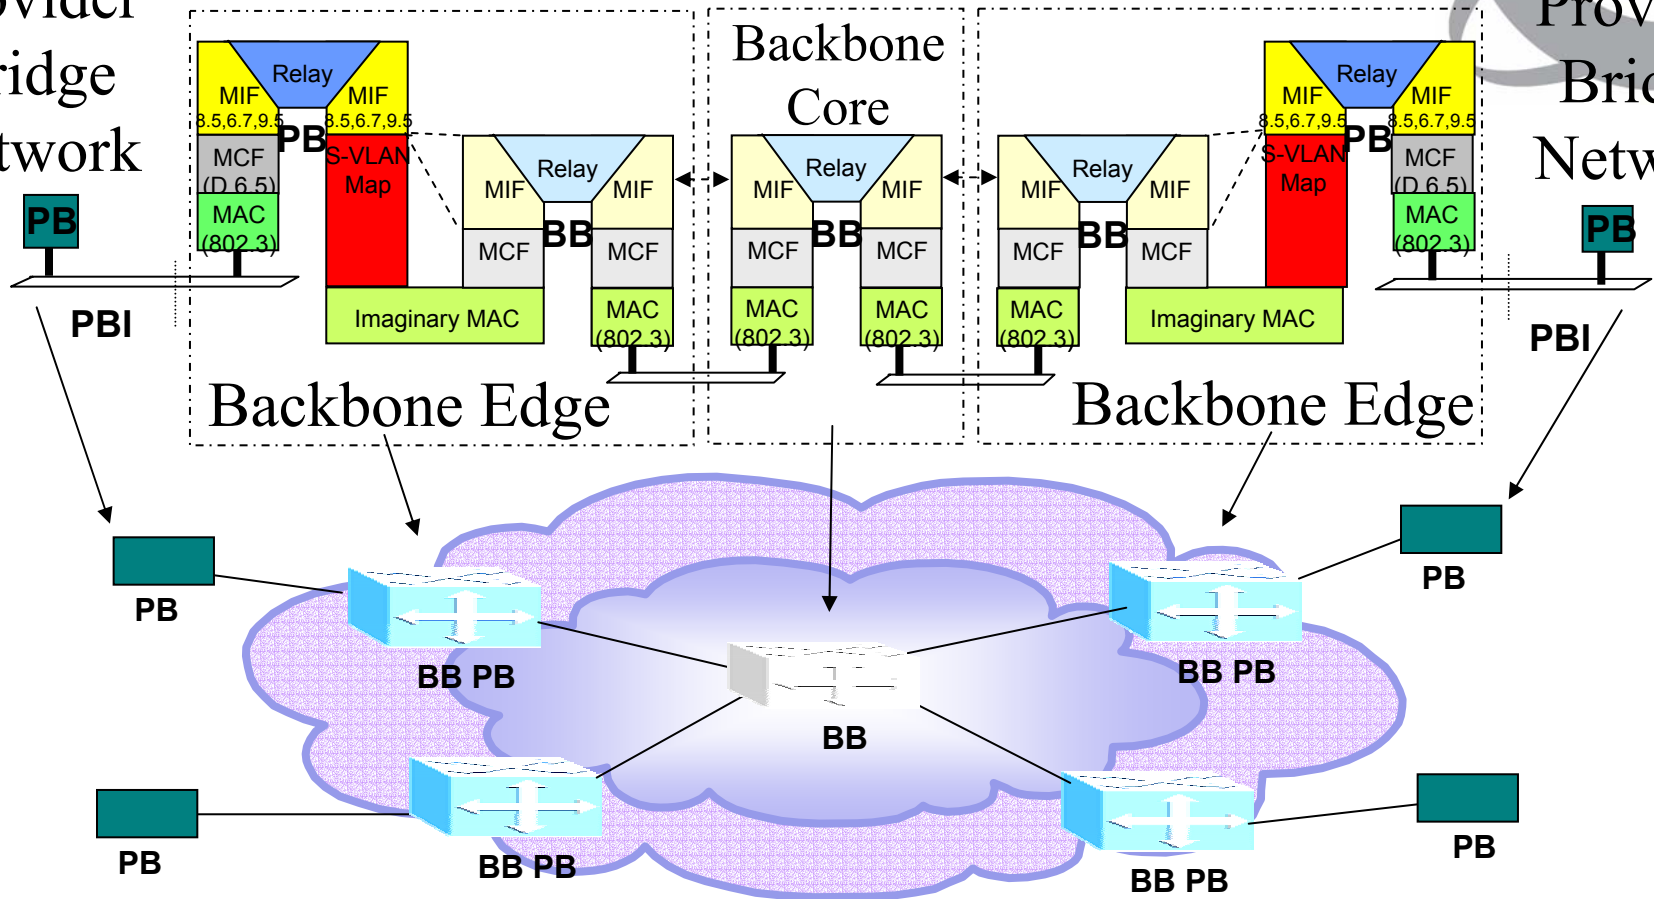
Provider Backbone Bridge Model



Backbone Core Relays Can be 802.1ad

Provider
Bridge
Network

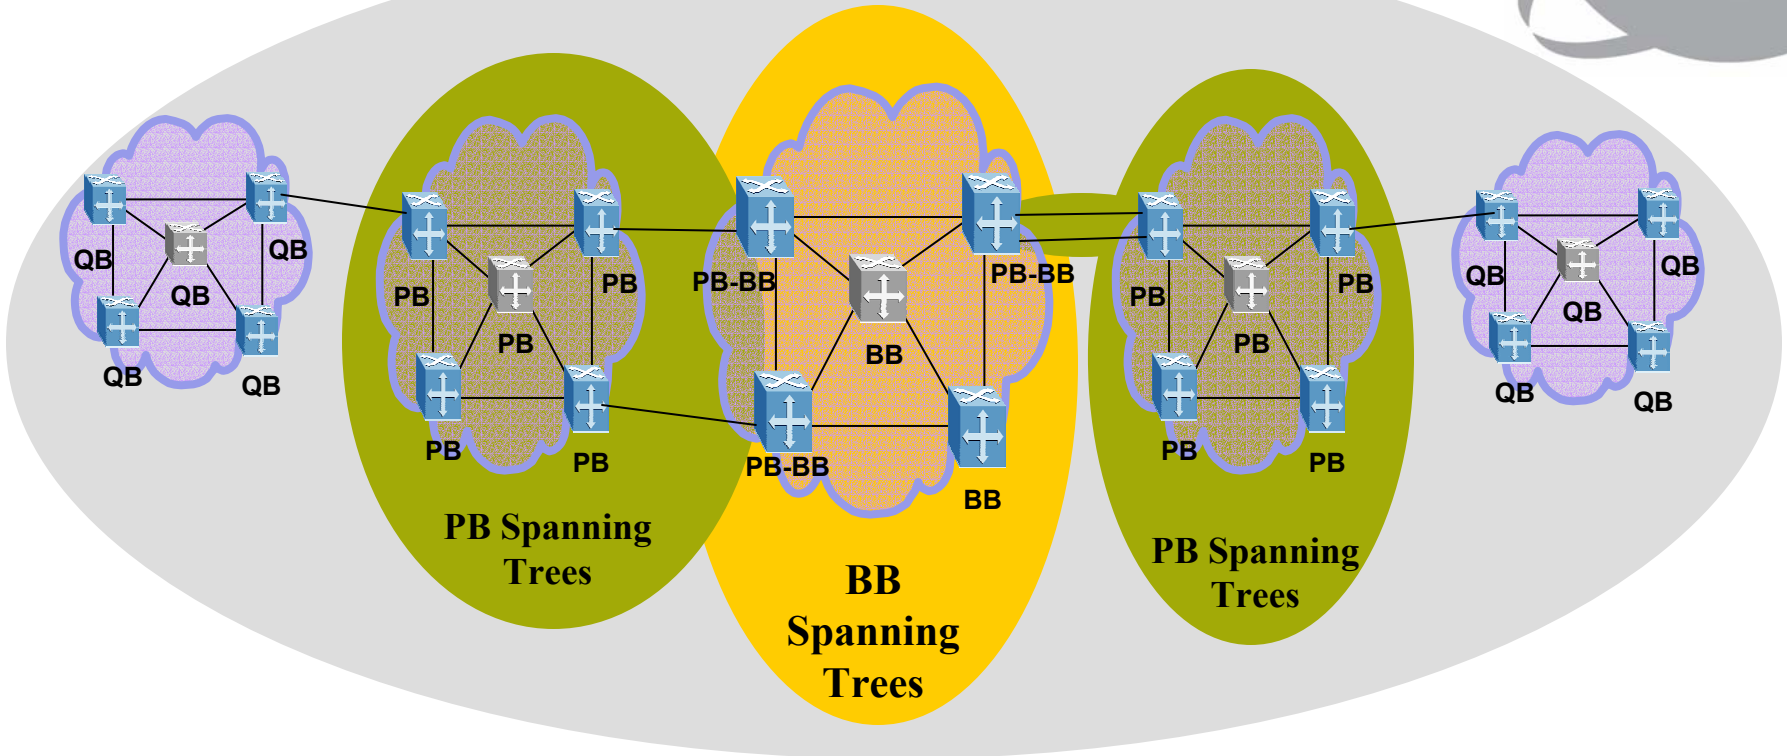
Provider
Bridge
Network



- > Backbone Core can be single 802.1ad relay
- > Backbone Edge is a dual 802.1ad relay and an encap/decap between the two relays.

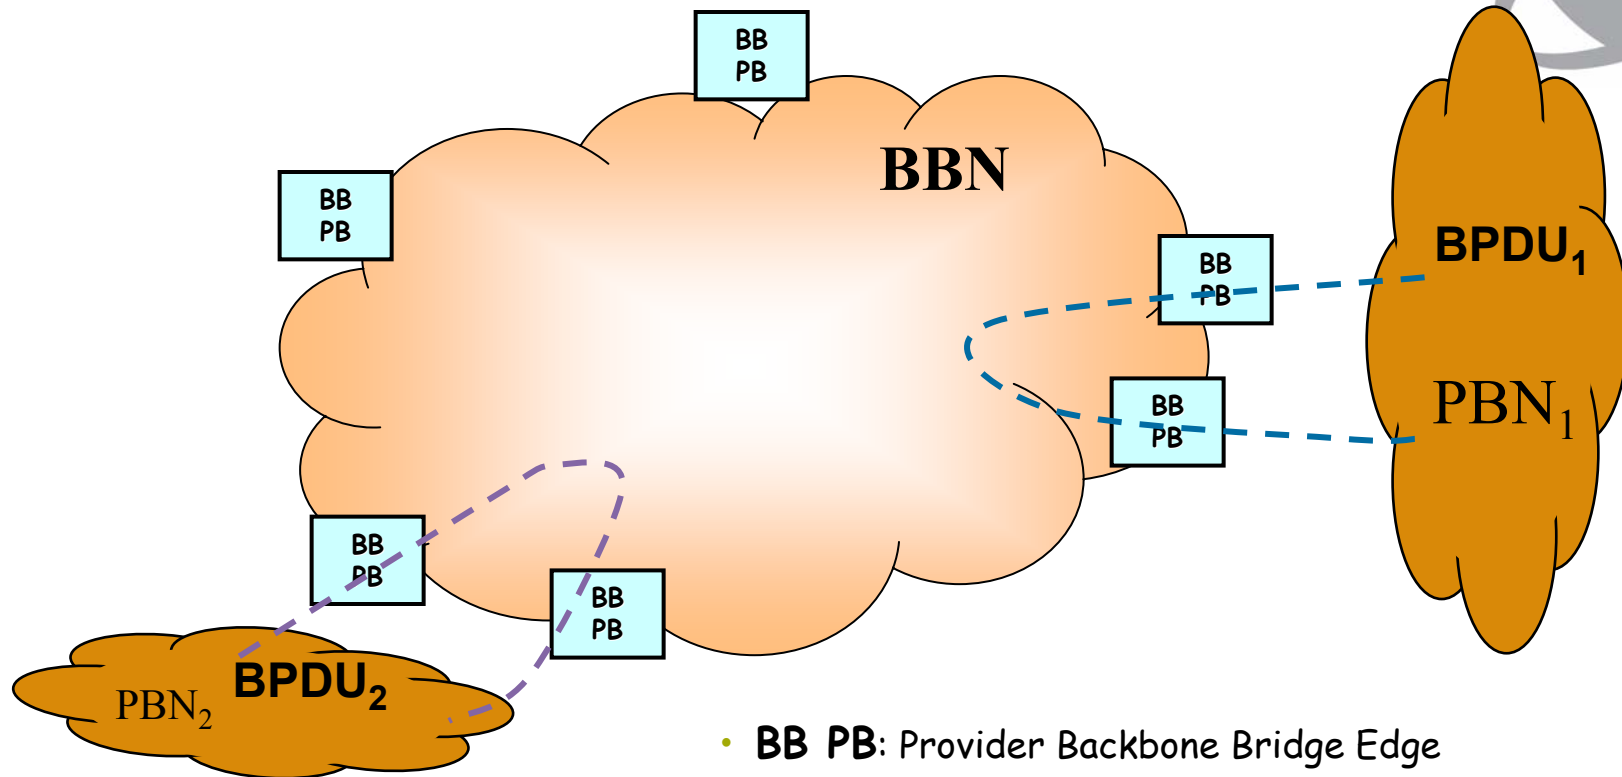
Customer, PB, BB Spanning Trees

Customer Spanning Trees



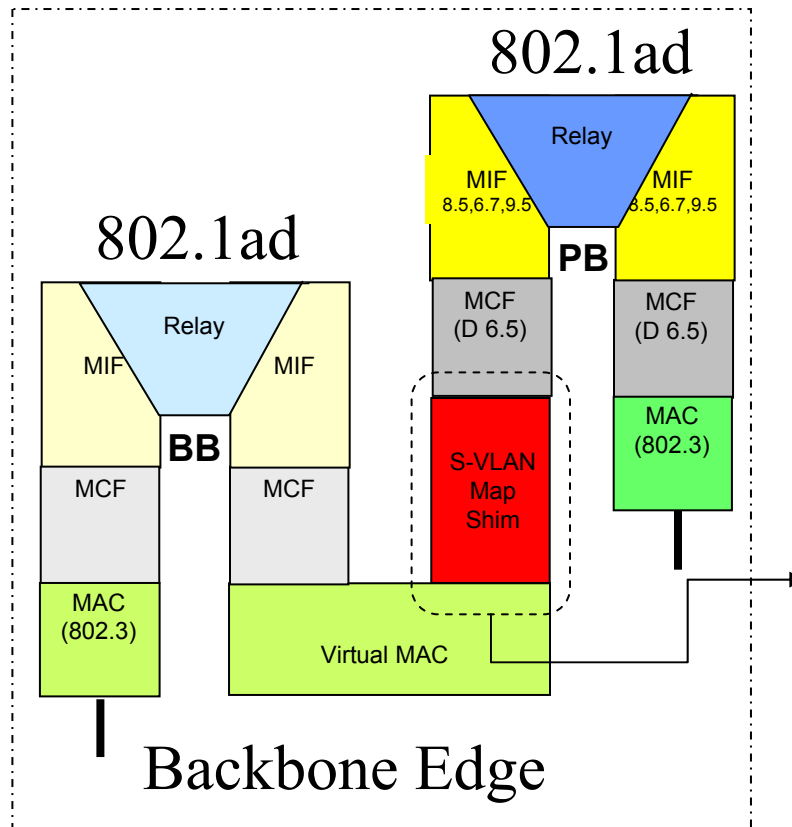
- > Customer spanning trees may extend over Provider Network
- > PB Network and BB Network spanning trees must be decoupled to scale the provider network
- > Provider Backbone Bridge may conform to the requirements for an Interconnect Medium

Provider Bridge Island BPDUs Delivery Inside Provider Backbone Bridge



- Each Provider Bridge Island may be connected to multiple Provider Backbone Bridge ports
- Provider Bridge Islands may not connect directly to each other
- Provider Backbone Bridge delivers Island BPDUs to all ports of that Island
- Island BPDUs are never delivered to other Islands by the BBN

BB Functions In Map Shim



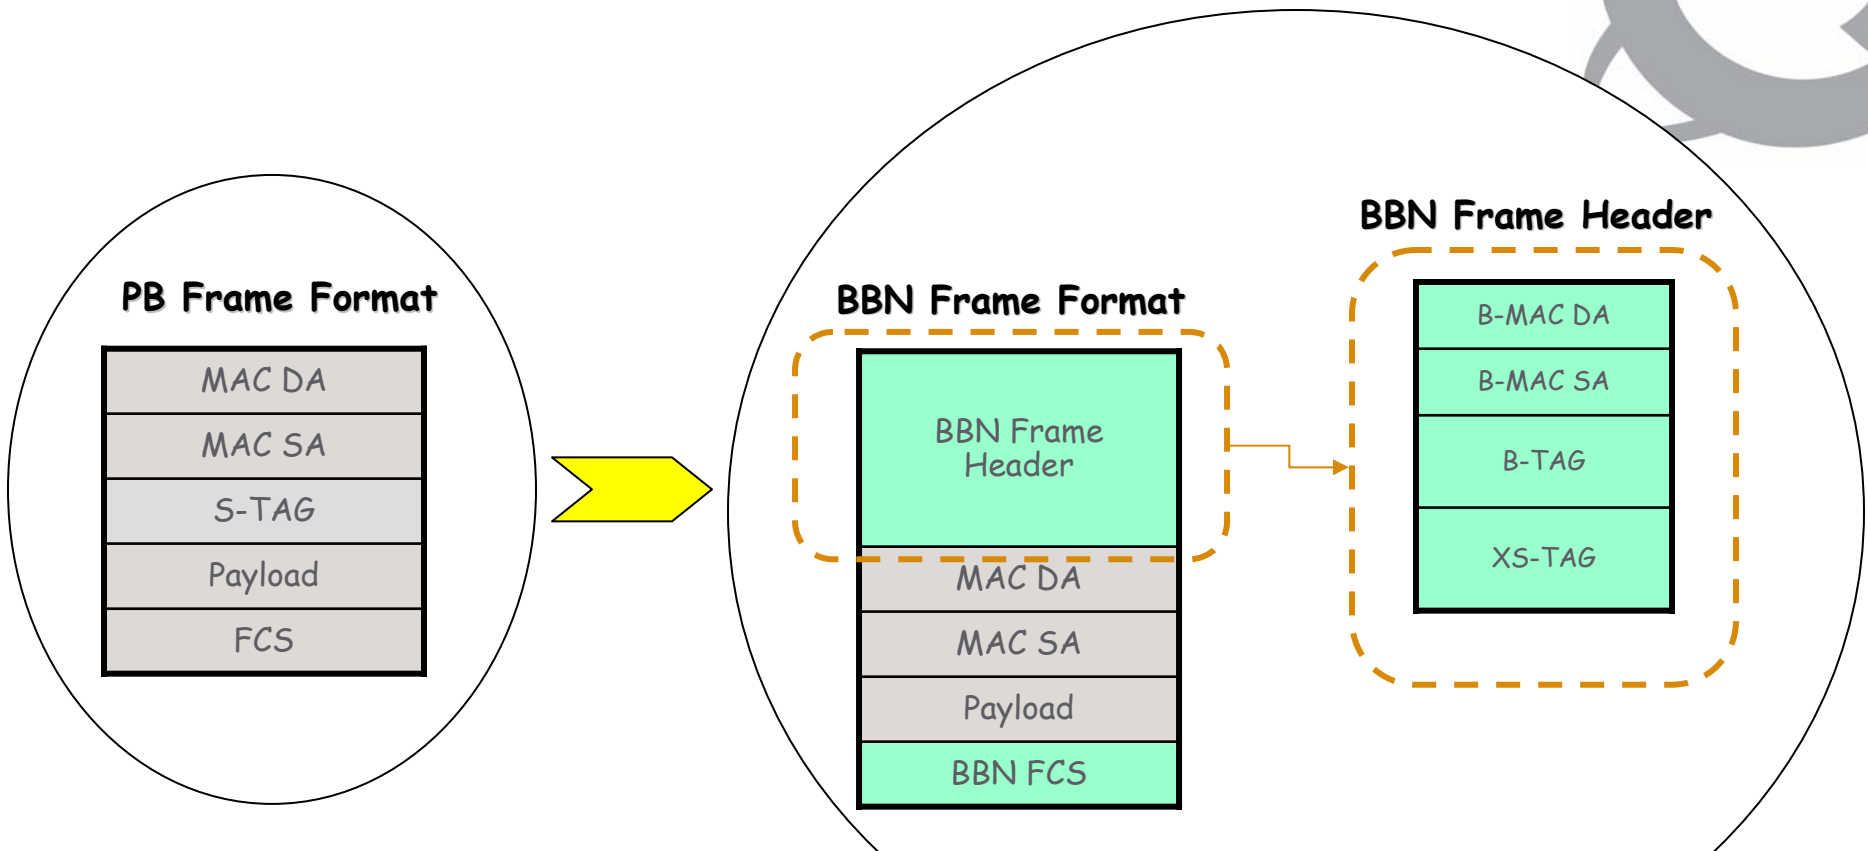
- > Does encaps/decap of 802.1ad frame
- > Maps S-VID from 802.1ad into larger Extended Service VID (XS-VID)
- > Learns and Correlates Backbone POP and Customer MAC addresses
- > Filters L2 control packets sourced by core relays or by provider bridge relays (divides spanning trees)

Map Shim Encap



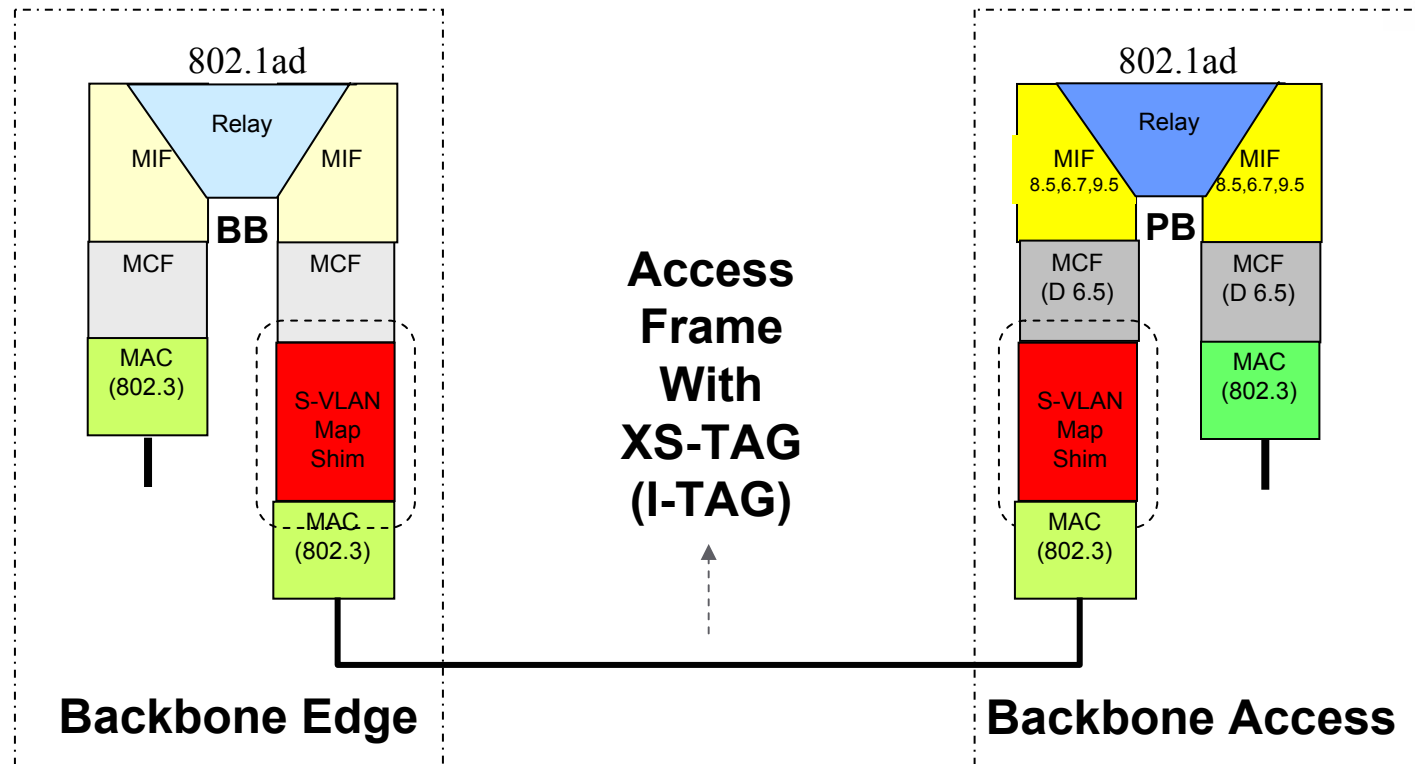
- BBN encapsulates PBN frames with BBN header
- BBN header consists of
 - a) Extended Service VLAN identifier
 - Identifies the Provider Bridge S-VLAN within the BBN
 - Requires 2^{20} bits to identify 1M services
 - b) Site Connectivity identifier
 - Identifies a B-VLAN (or tunnel) that is used to transport the BBN service instance
 - Site connectivity (i.e., tunnel/domain) can be point-to-point or multi-point in nature
 - c) Backbone POP Address
MAC Address for POPs within Site Connectivity
- PBN Service VLAN IDs (S-VIDs) map to BBN Extended Service VLAN IDs (XS-VIDs)
 - PBN S-VIDs are local to the PBN
 - BBN XS-VIDs are local to the BBN

Backbone Frame Format



- > Removing S-Tag is most efficient encode
- > Since FCS is also most efficient encode
- > Un B-Tagged frames could be used
- > B-Tag format should be identical to 802.1ad

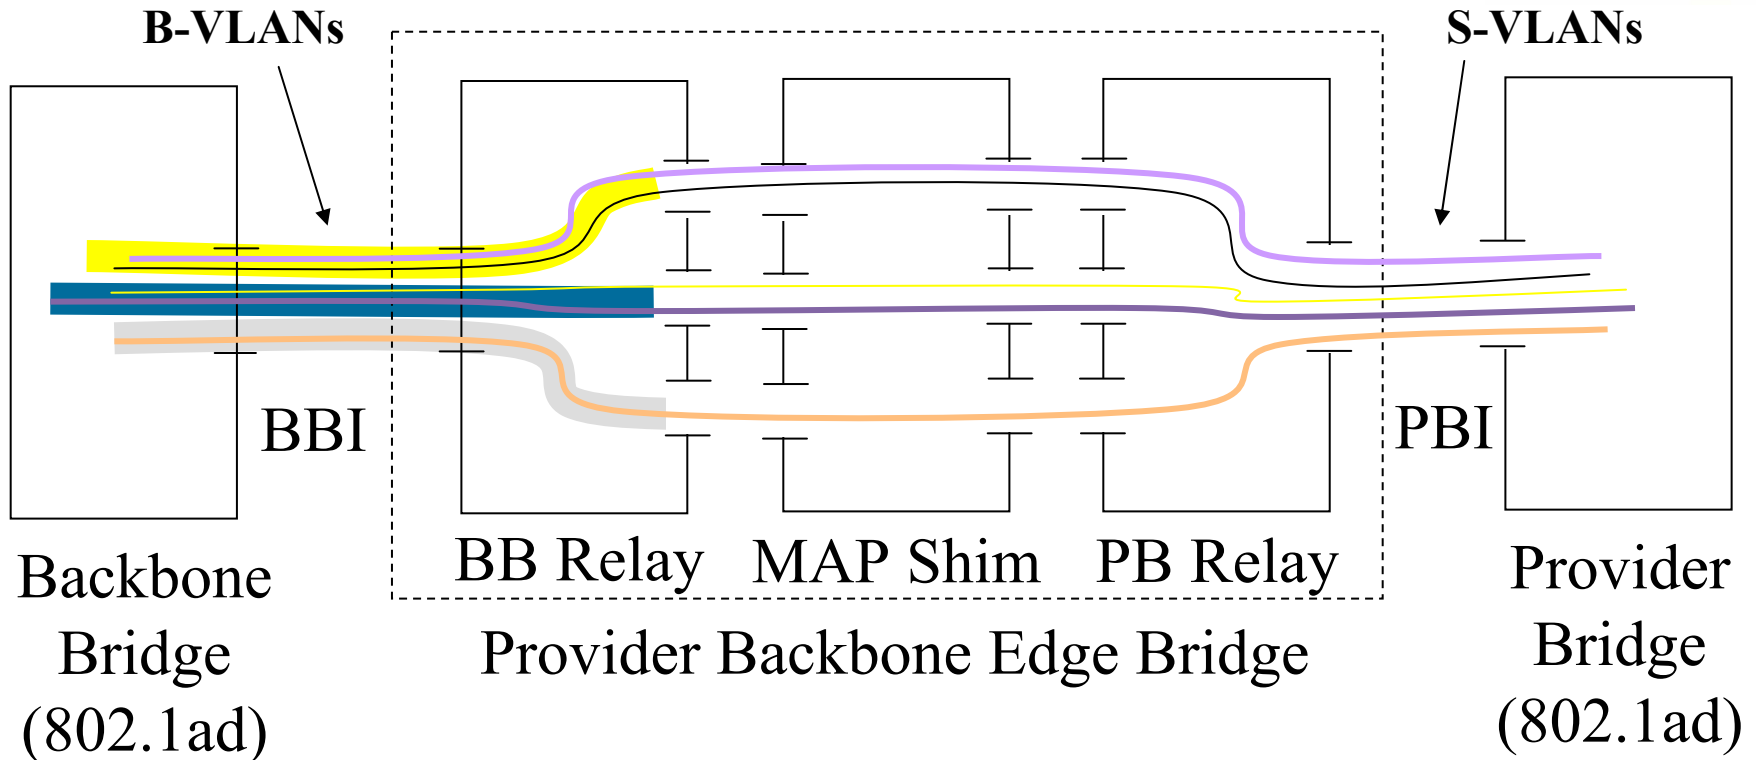
MAP Shim on Ingress or Egress



- > Both work
- > Best if located in one or the other

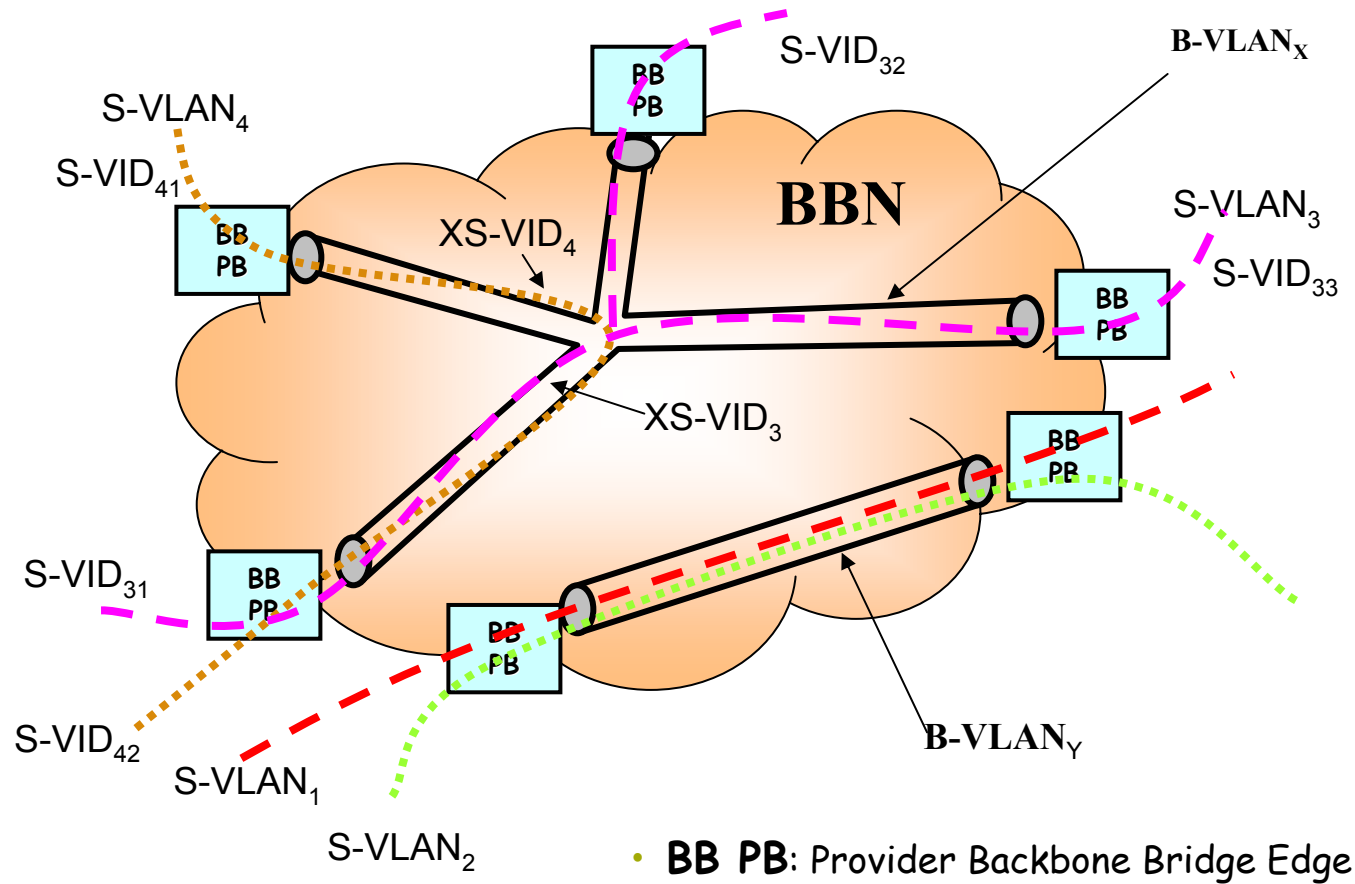


S-VLANs Multiplex into B-VLANs



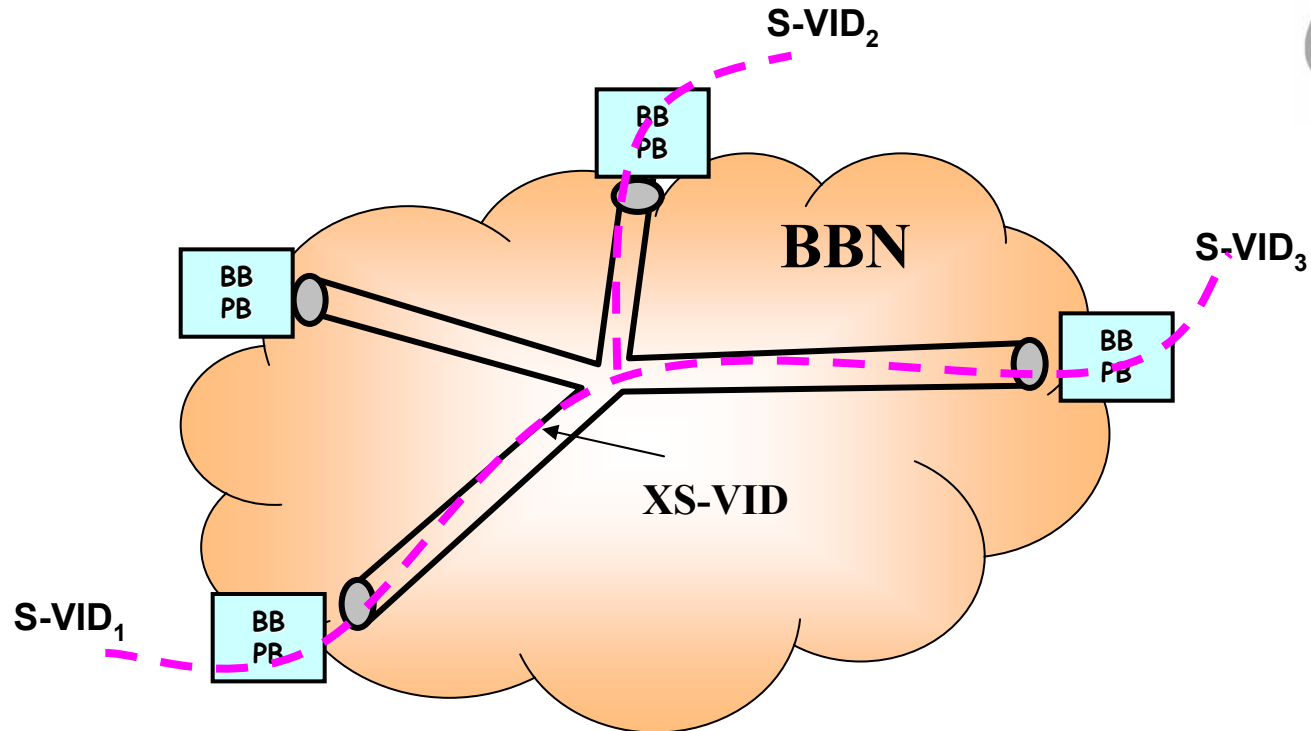
- > MAP Shim performs encaps/decap of frames to/from Provider Bridge Networks

Extended Service VLAN IDs In Backbone



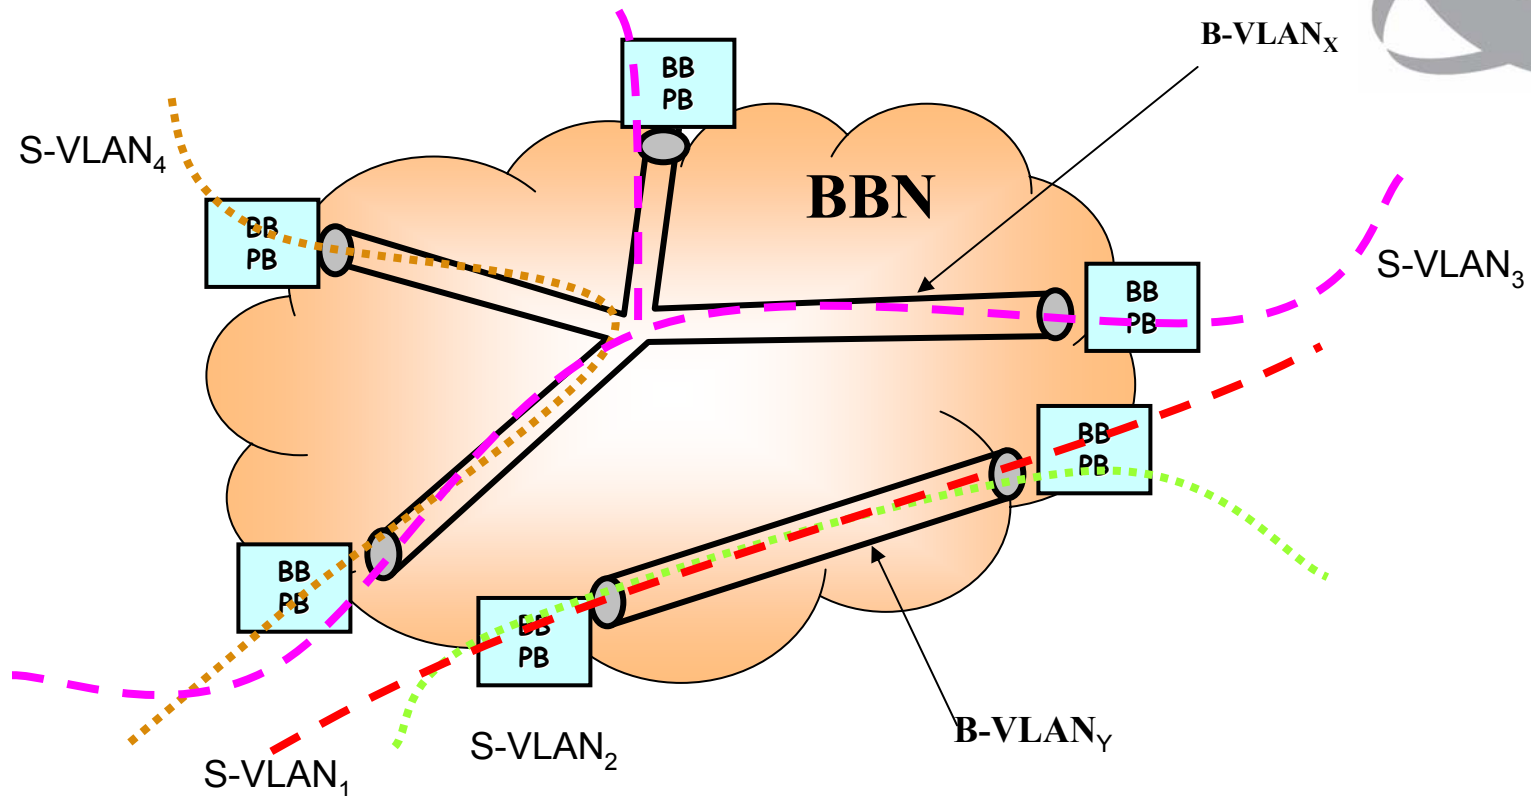
- An XS-VID uniquely identifies a S-VLAN within the Backbone
- The MAP Shim translates between S-VID and XS-VID
- The XS-VID to(from) S-VID mapping is provisioned when a new service instance is created

Single XS-VID per S-VLAN



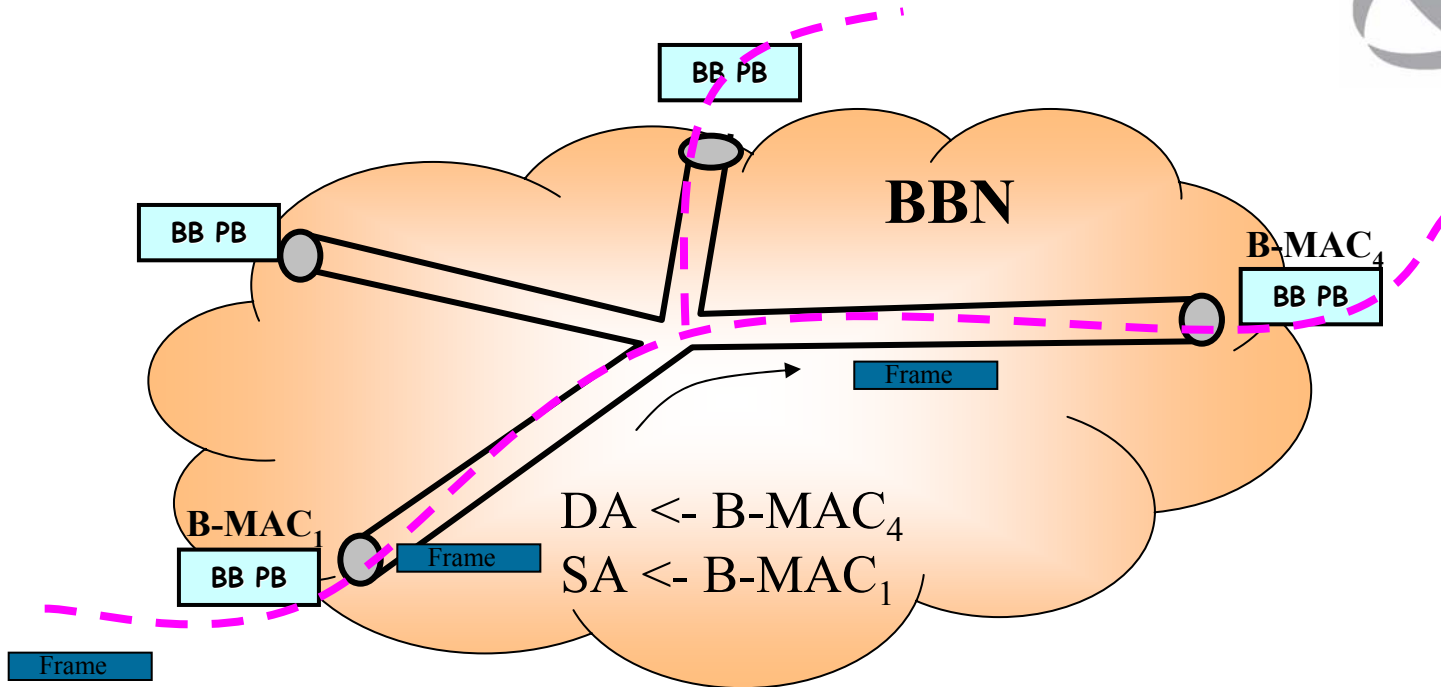
- > Regardless of the XS-VID address size the map tables only have 4096 entries since only one XS-VID exists per S-VLAN and only 4096 S-VLANs exist per Provider Bridge.
- > A different S-VID in each PBN maps to the XS-VID

Site Connectivity B-VLAN ID



- > B-VLANs are addressed like regular VLANs with a 12 bit B-VID
- > B-VID and XS-VID need to be separate ID spaces to allow many S-VLANs to be carried in a single B-VLAN

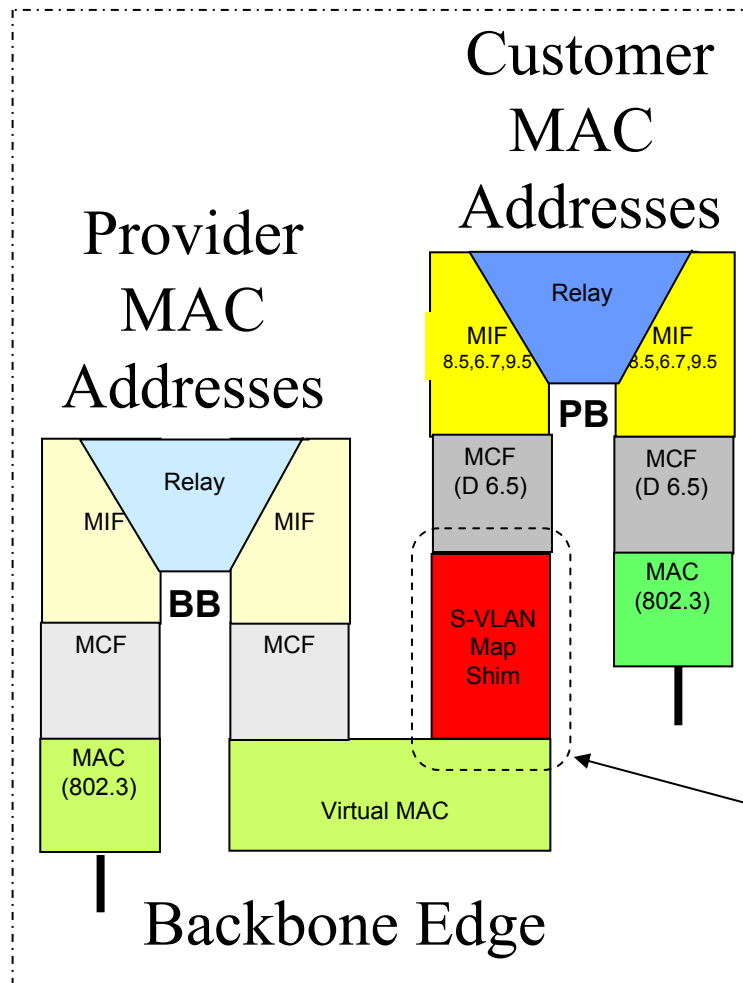
Backbone POP MAC Address



- > B-MAC Addresses identify the Edge Provider Backbone Bridges (BB PB)
- > B-MAC Addresses are learned by other Edge Backbone Edge Bridges
- > The backbone edge MAC address determines which edge on the B-VLAN will receive the frame.
- > Frames may be flooded by sending with broadcast or multicasts DA B-MACs to the B-VLAN.
- > Map shims filter based on the XS-VID removing any misaddressed frames



Customer/Provider Addresses

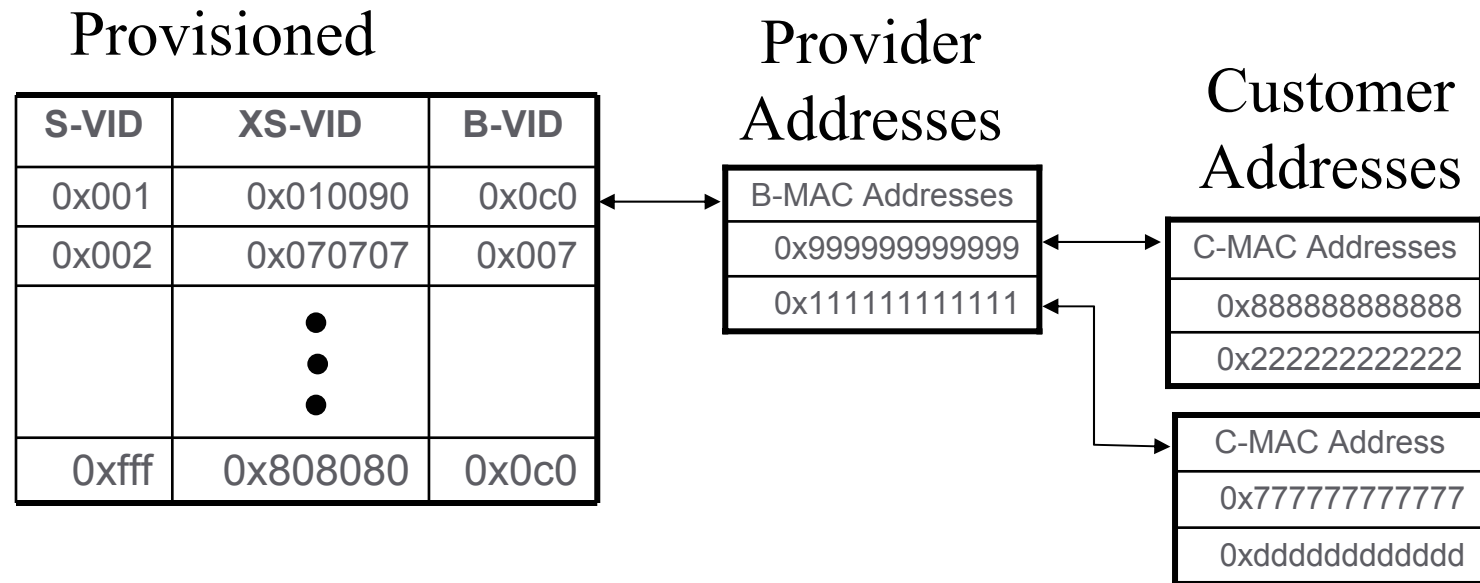


- > PB Relay Learns Customer Address Per S-VLAN
- > BB Relay Learns Provider Addresses Per B-VLAN
- > MAP Shim Learns Correlated Customer and Provider MAC Addresses per S-VLAN

Customer/Provider
MAC Address
Correlation



MAP Shim Correlation Table



- > In the beginning the MAP Shim is provisioned with the correlation between the S-VID, XS-VID, and B-VID
- > During operation the MAP Shim learns both B-MAC addresses and C-MAC addresses
- > The MAP Shim keeps track of which C-MAC addresses are behind which B-MAC
- > The correlation data is used to encapsulate frames from the PBNs



Basic MAP Shim Operation

- > Frames received from PB Relay are encaped
 - S-VID is looked up in correlation table to get XS-VID and B-VID
 - C-DA is looked up in C-MAC table to get B-MAC for encapsulation
 - If C-DA is not present in C-MAC table then multicast to B-VLAN
- > Frames received from BB Relay are de-encaped
 - XS-VID is looked up in correlation table to get a new S-VID
- > B-MAC and C-MAC addresses are learned when frames are received from BB relay
- > B-MAC and C-MAC addresses are aged



Summary

- > A Provider Backbone Bridge standard needs to define the functions of the MAP Shim
- > The 802.1ad control plane may be used on both sides of the MAP Shim
- > Connection Fault Management 802.1ag should be supported by the Provider Backbone Bridges





Backup Material



Terminology

> IEEE 802.1ad Terminology

- C-TAG Customer VLAN TAG
- C-VLAN Customer VLAN
- C-VID Customer VLAN ID
- S-TAG Service VLAN TAG
- S-VLAN Service VLAN
- S-VID Service VLAN ID

> Additional Provider Backbone Bridge Terminology

- XS-TAG Extended Service VLAN TAG Field (I-TAG)
- XS-VID Extended Service VLAN ID (SID)
- C-MAC Customer MAC Address
- B-MAC Backbone MAC Address
- B-VLAN Backbone VLAN (tunnel)
- B-TAG Backbone TAG Field
- B-VID Backbone VLAN ID (tunnel)



Service Instance Address Space Size Options

- > Carriers need to separate the service address space to allow administration of networks
 - Allocation of address blocks to offices
 - Merging network elements
- > The address space usually needs to be 10-100 times larger than the number of services supported
- > Should have an address space around 2^{24}
- > Use of 2^{20} address space would match MPLS
- > Need to resolve this issue

Provider Backbone Bridges May Apply the 5 IM Rules.



1. Each 802.1ad island is responsible for preventing internal forwarding loops.
2. The 802.1ad islands connect to other only through Provider Backbone Bridge.
3. Each 802.1ad island ensures that no customer data frame passes through more than one Provider Backbone Bridge attachment into or out of the island.
4. Each 802.1ad island ensures that it attaches any given S-VLAN to no more than one Provider Backbone Bridge network.
5. A Provider Backbone Bridge network ensures that if an attached port can talk to *any* other attached port, it can talk to *all* of the ports attached to the Backbone network.