Backbone Provider Bridging Networks

A Highly Scalable VLAN (Multicast) Architecture

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

Version 2.0

October 4, 2004 Ottawa

NORTEL NETWORKS™

A Provider Bridge Scaling Solution "Backbone Provider Bridging"

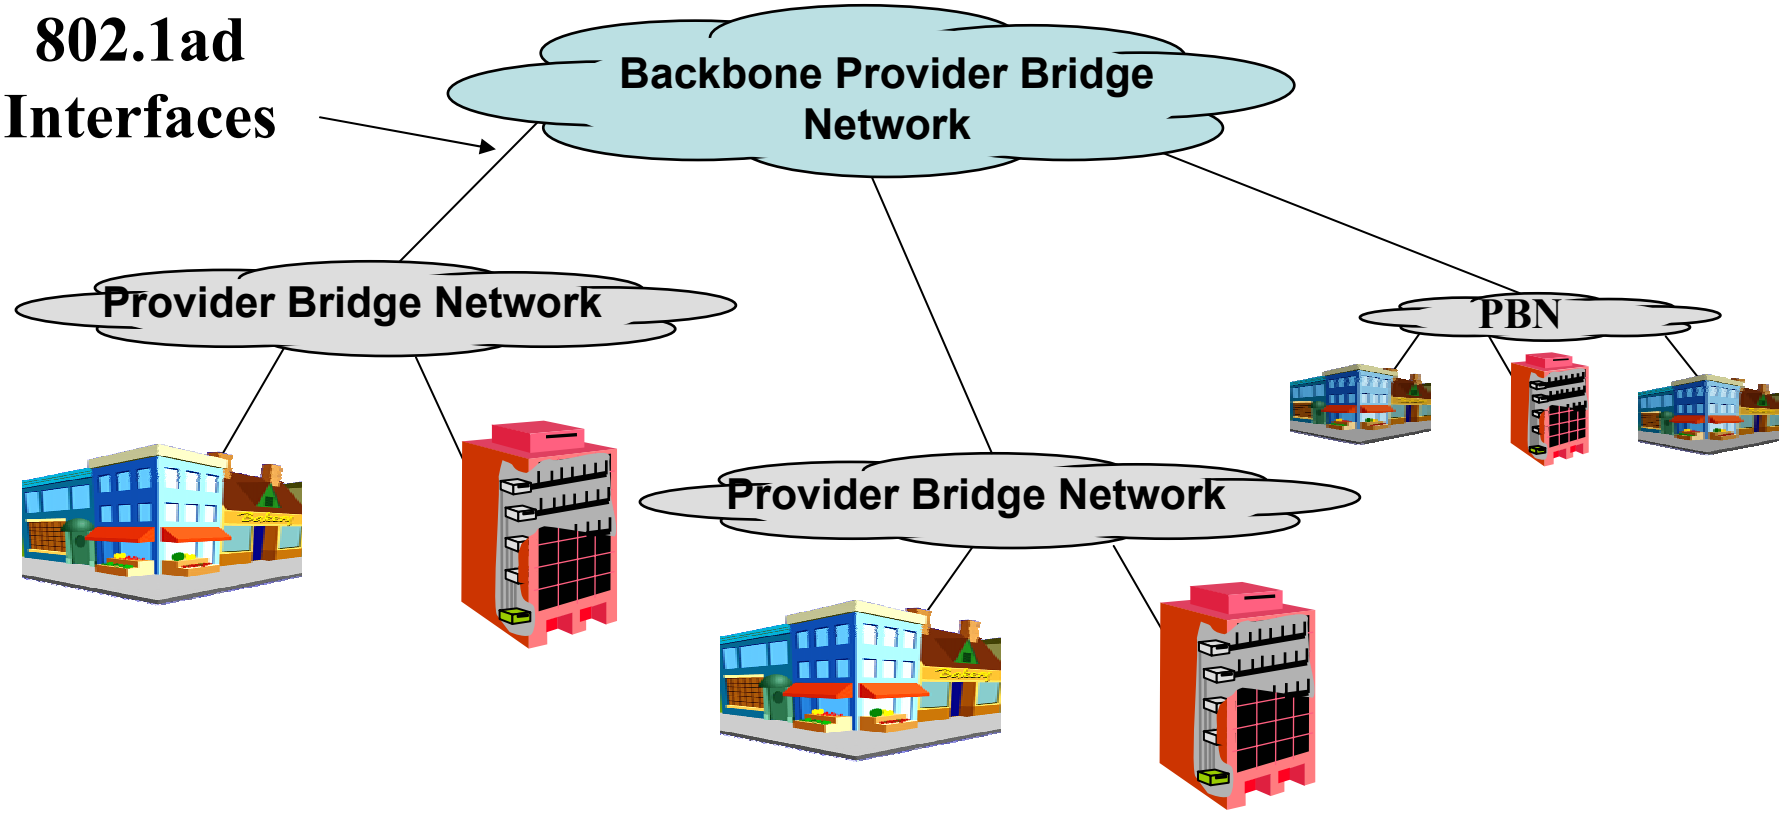

NORTEL NETWORKS

Ethernet Service Types

MEF Ethernet Virtual Connections (EVCs)

E-LINE
Router Mesh

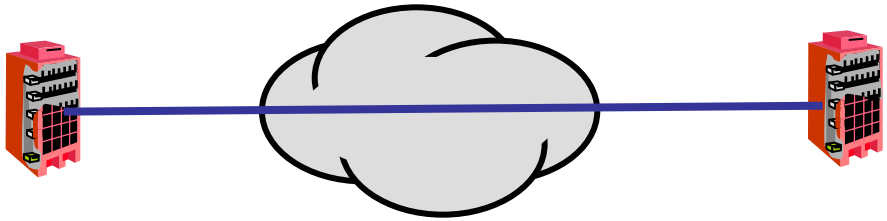

Pt-Pt, Like
Duplex Ethernet
Any-to-any

E-TREE
Hub & Spoke

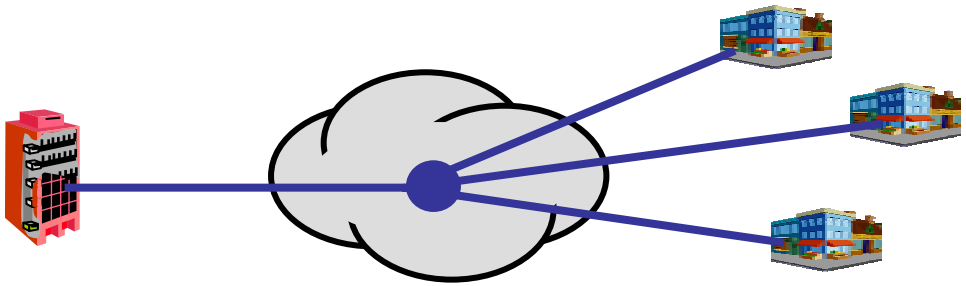

Pt-MPt, Like
EPON Ethernet,
Root-to-Leaf and
Leaf-to-Root

E-LAN
Multi-Site

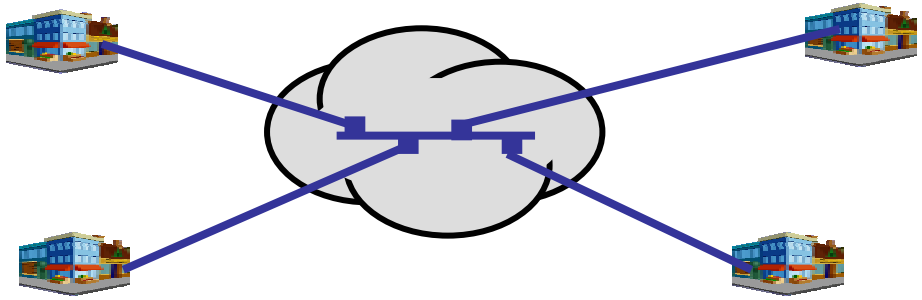

MPt, Like VLAN,
Any-to-any

NORTEL NETWORKS™

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-existant, SLAs can only provide Best Effort

E-TREE S-VLAN Mapping

E-TREE
Hub & Spoke

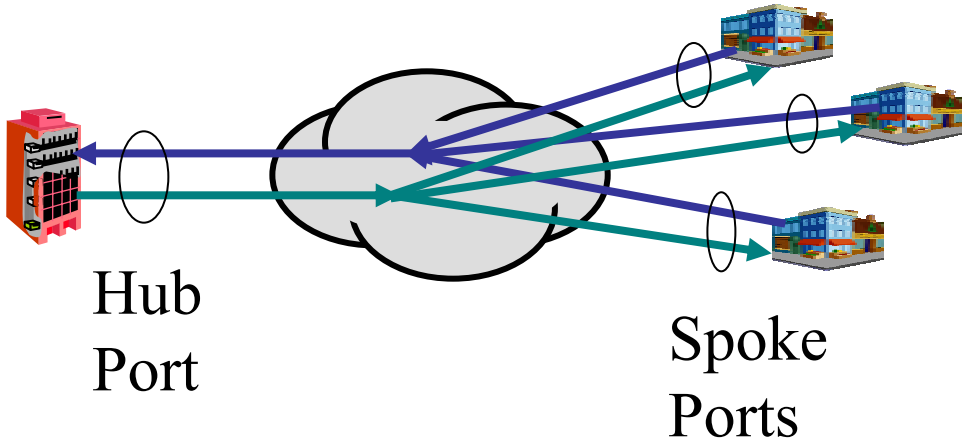

Pt-MPt, Like EPON Ethernet, Root-to-Leaf and Leaf-to-Root

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

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

E-LINE
Hub & Spoke

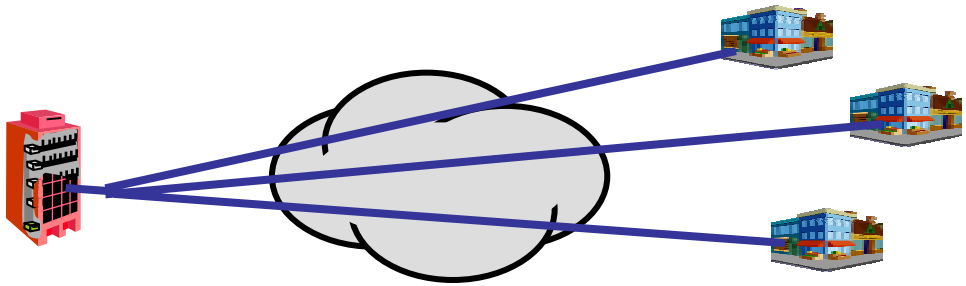

Hub
Port

Spoke
Ports

Pt-Pt
Root-to-Leaf and
Leaf-to-Root

- Hub port would usually be multipexed to allow the multiple Pt-Pt attachments.
- Each E-LINE is a seperate 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 residental

Major MSA Networks

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

Typical Metropolitan Serving Area – MSA

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

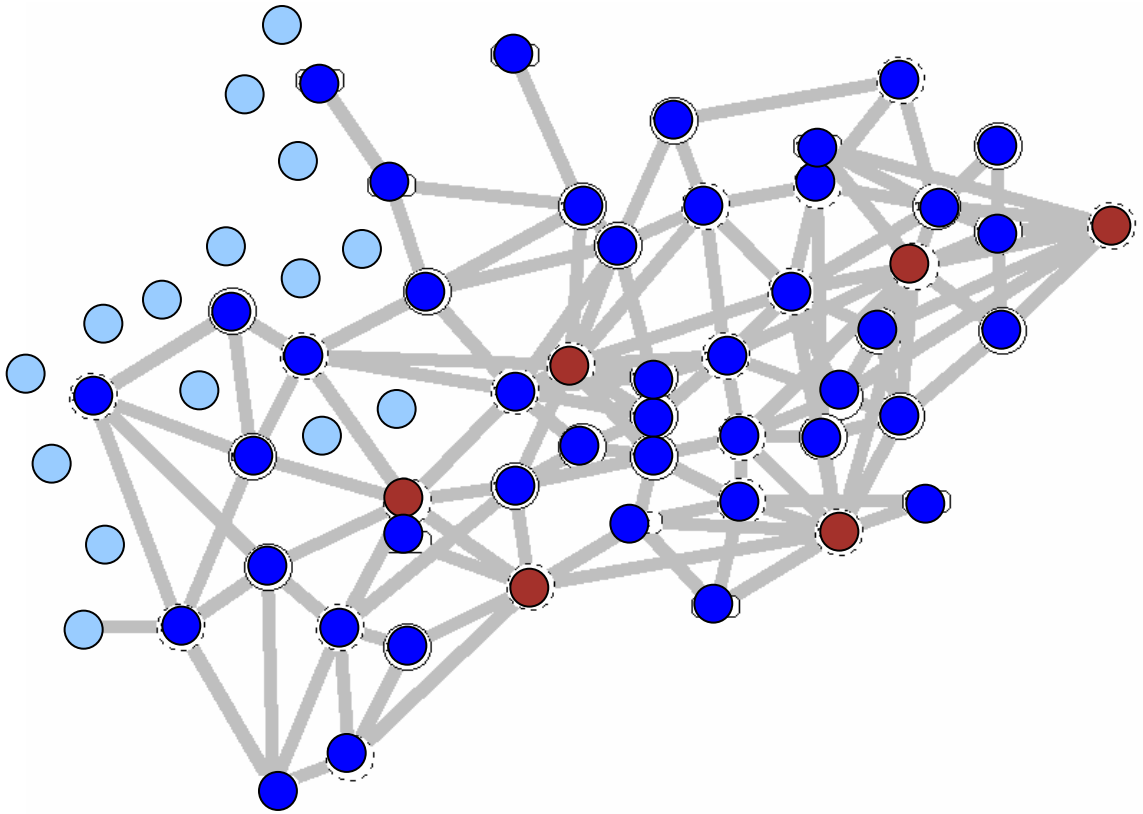

Support 1,000,000 Service Instances

- 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

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

Large Service Address Space Needed

- 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

Backbone Provider Bridge Technology Principles

Backbone Provider Bridge Project

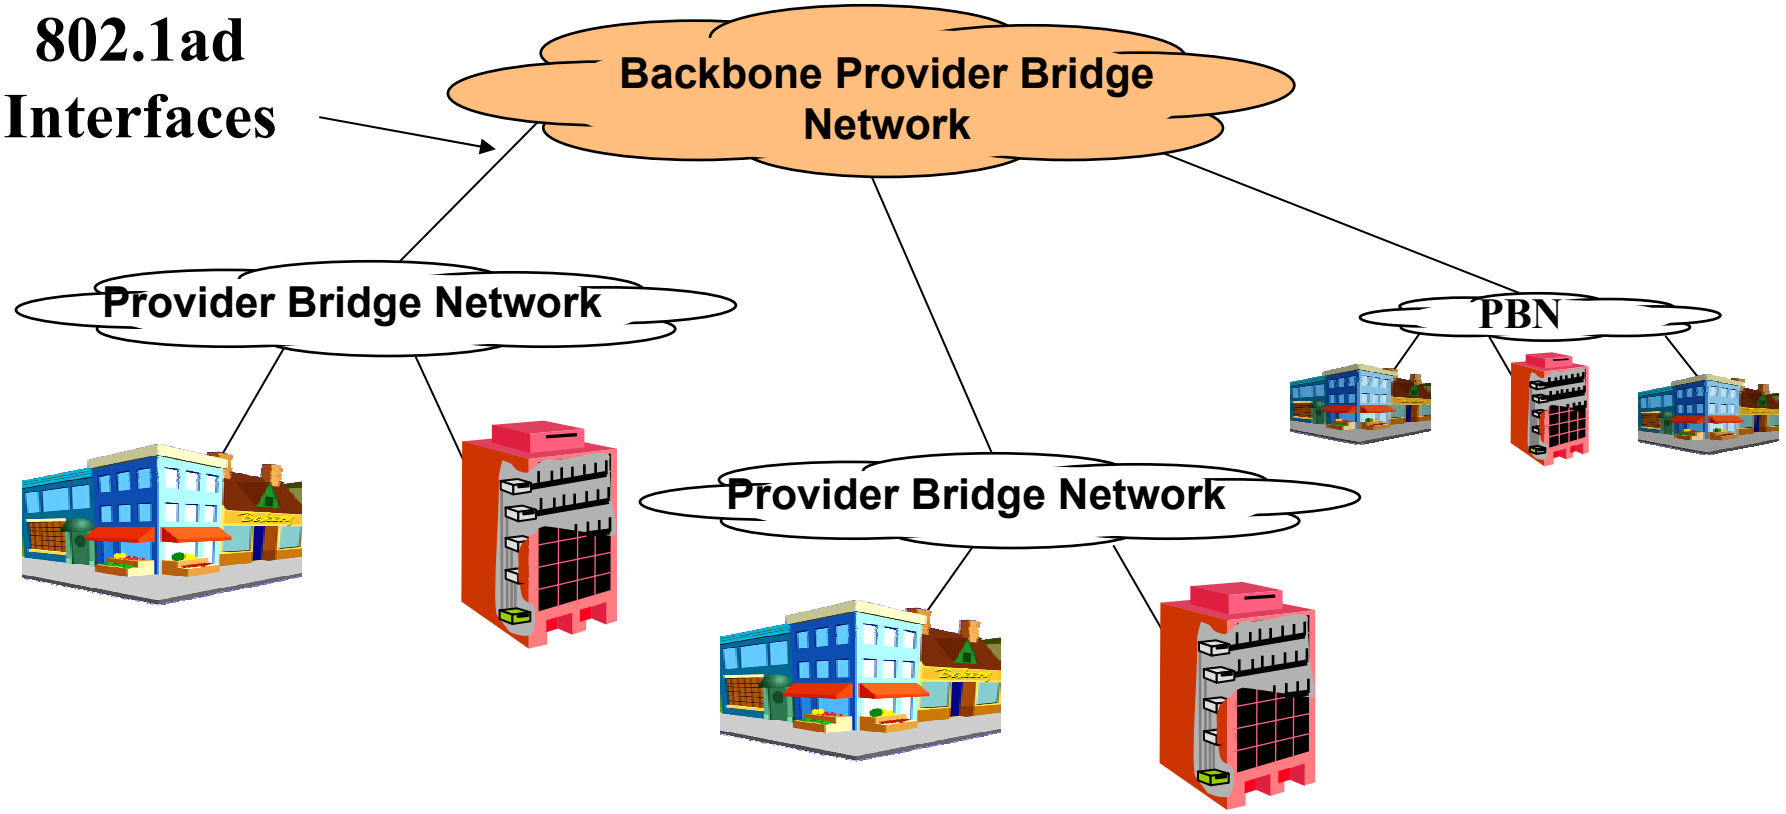

Backbone Provider Bridge Network

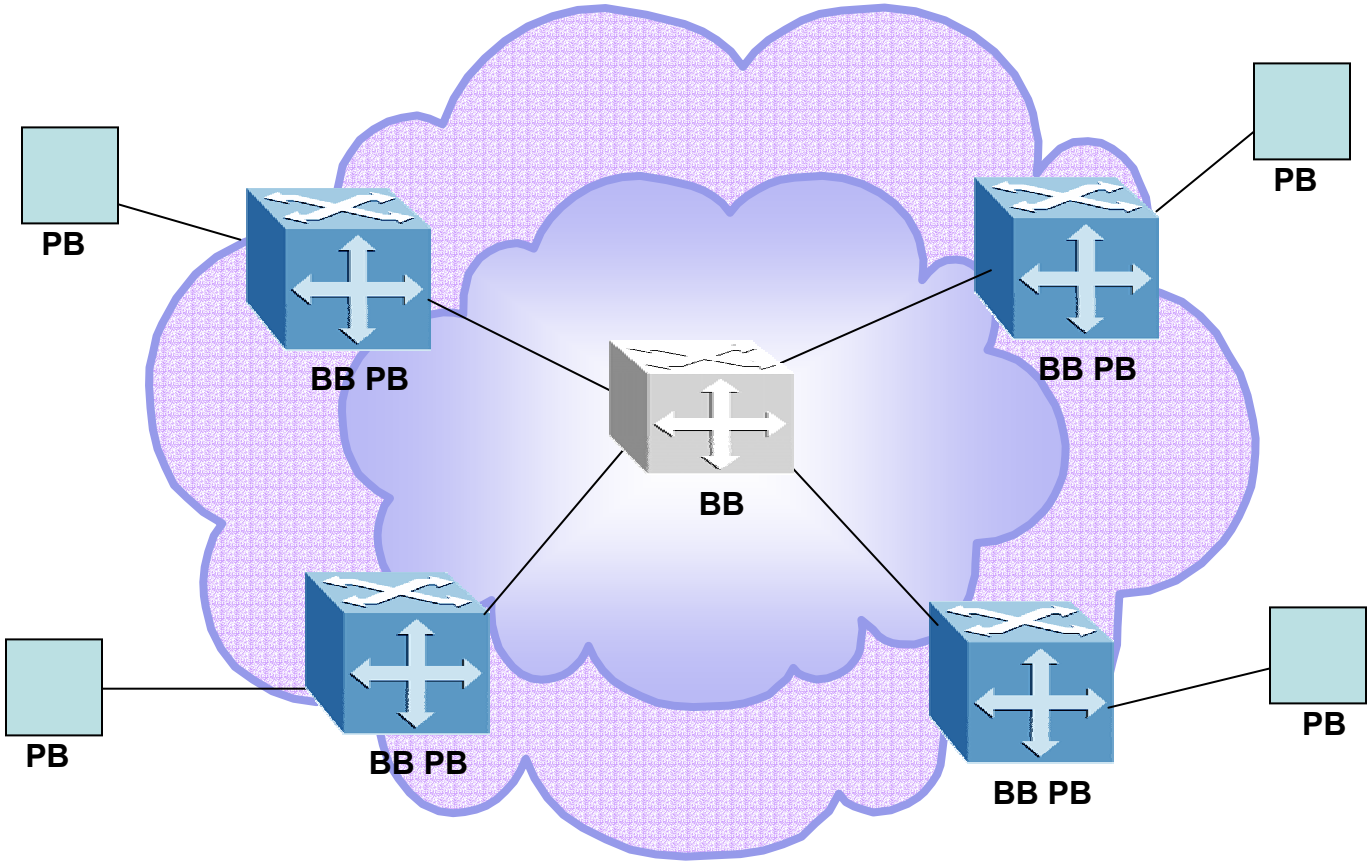

- **PB**: Provider Bridge (as defined by 802.1ad)
- **BB PB**: Backbone Provider Bridge Edge
- **BB**: Backbone Provider Bridge

NORTEL NETWORKS

BBN Provides Multi-Point B-VLANs Between PBNs

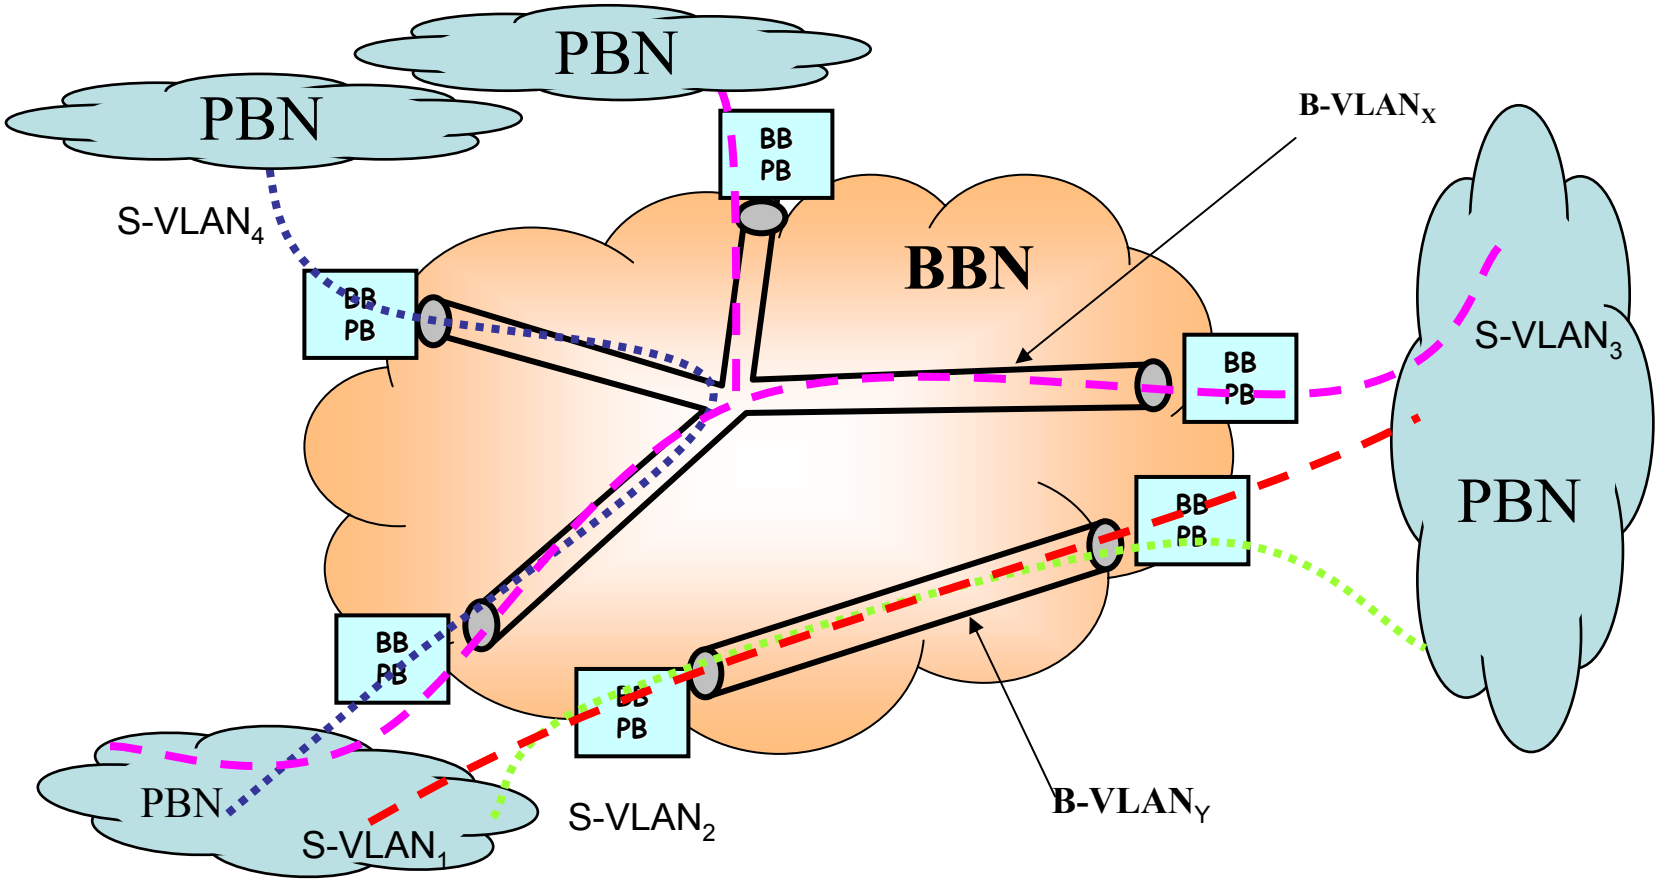

- **BB PB**: Backbone Provider Bridge Edge

- 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.

Backbone Provider Bridge Model

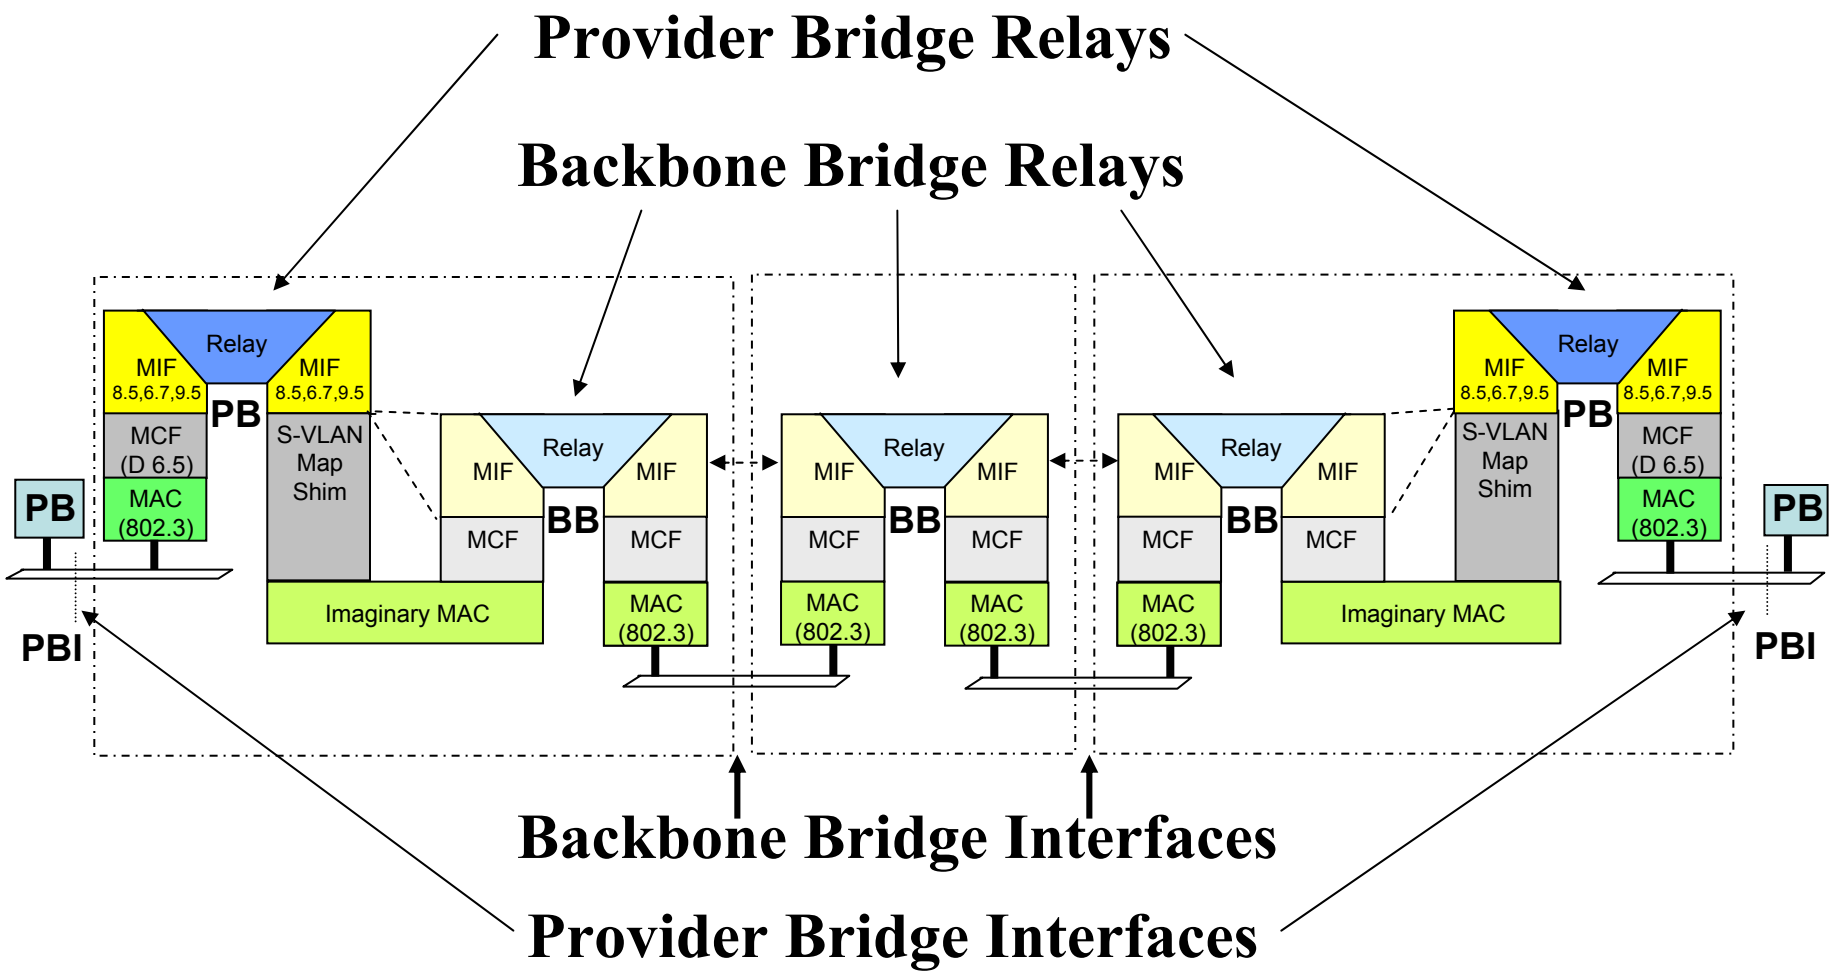

Backbone Core Relays Can be 802.1Q

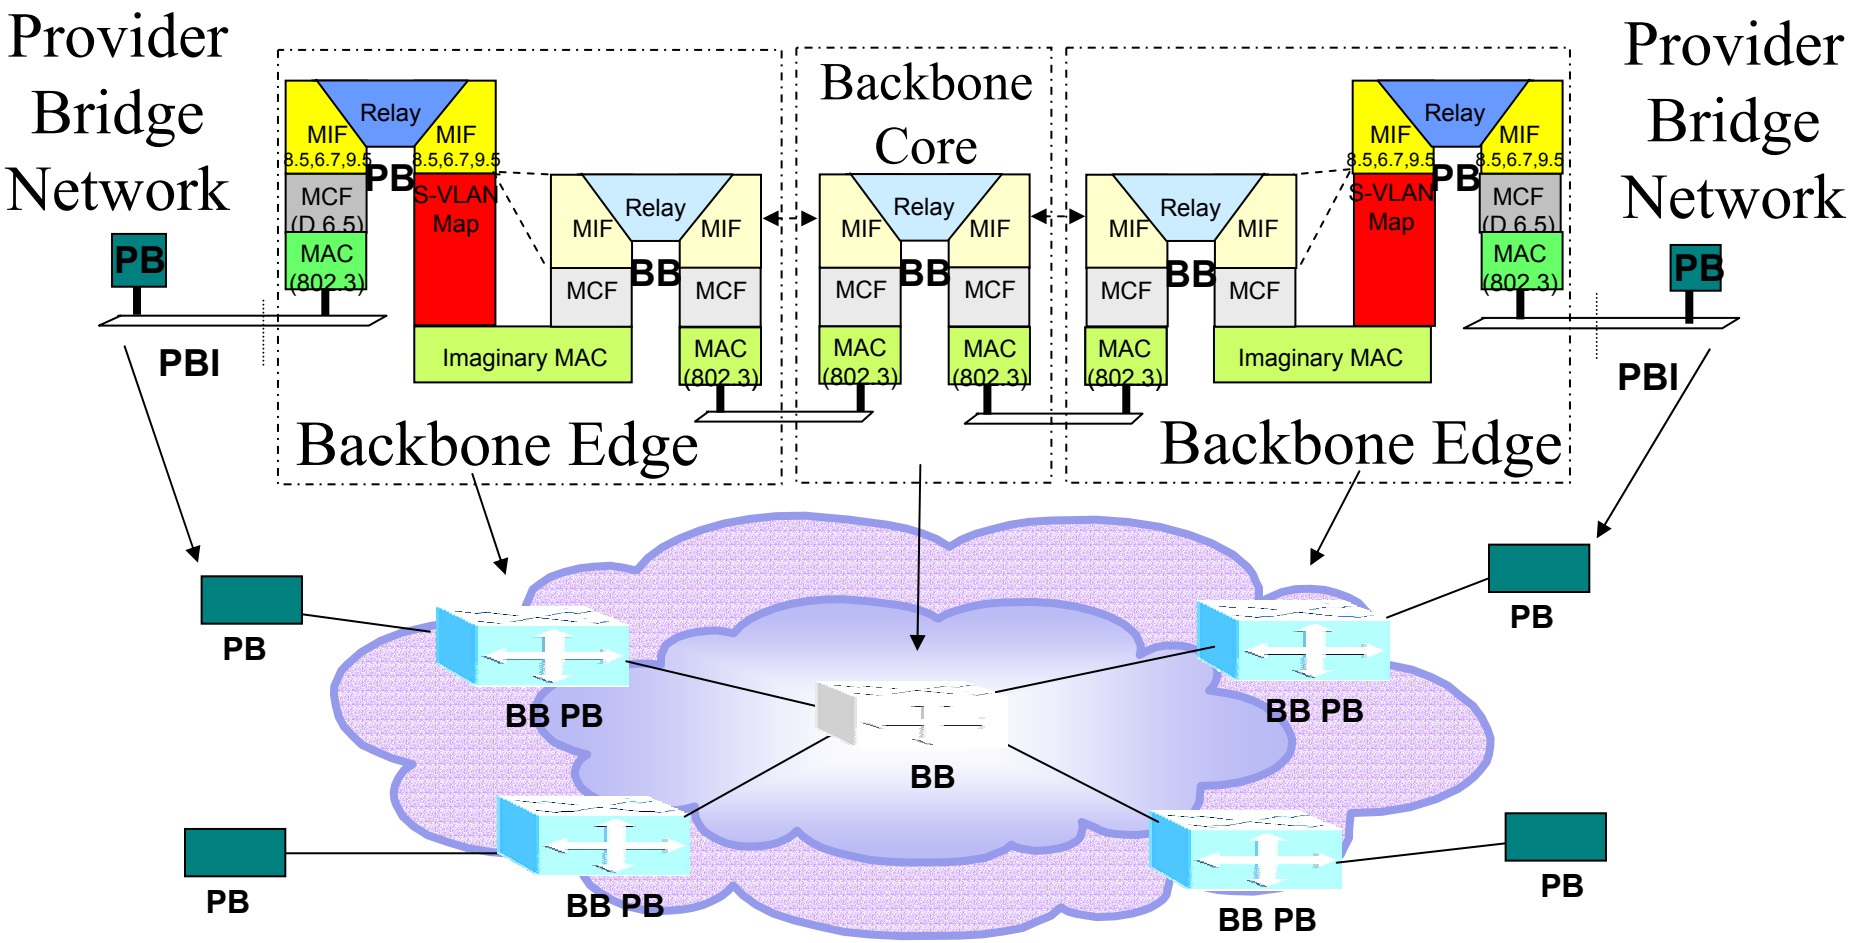

- Backbone Core can be single 802.1ad relay
- Backbone Edge is a dual 802.1ad relay and an encap/decap between the two relays.

BB Functions In Map Shim

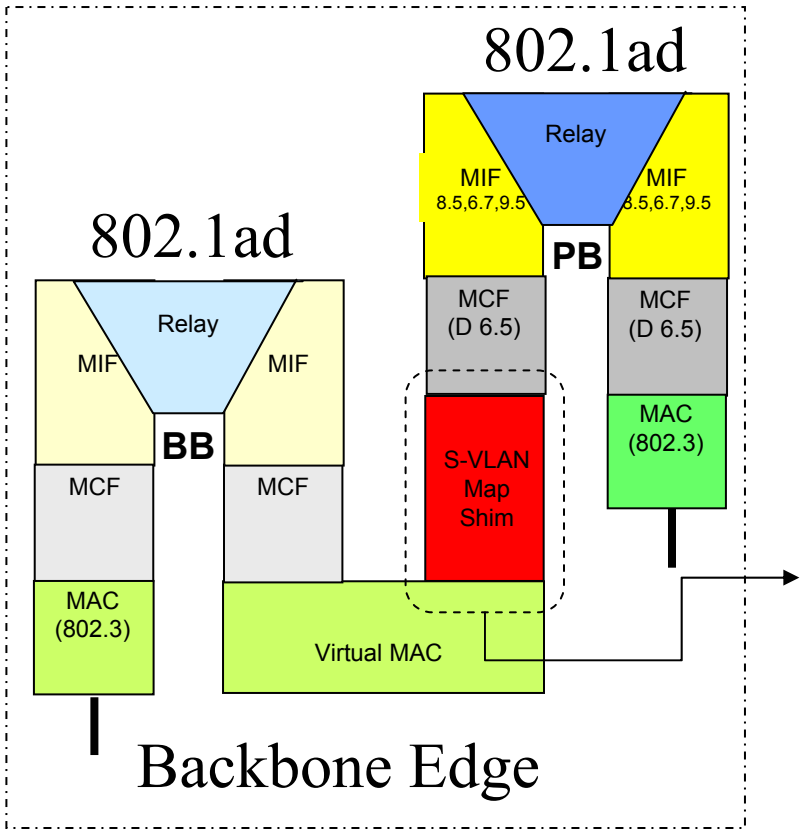

- Does encap/decap of 802.1ad frame
- Maps S-VID from 802.1ad into larger Extended Service VID (ES-VID)
- Learns and Correlates Backbone POP and Customer MAC addresses
- Filters L2 control packets sourced by core relays or by provider bridge relays (divides spanning trees)

S-VLANs Multiplex into B-VLANs

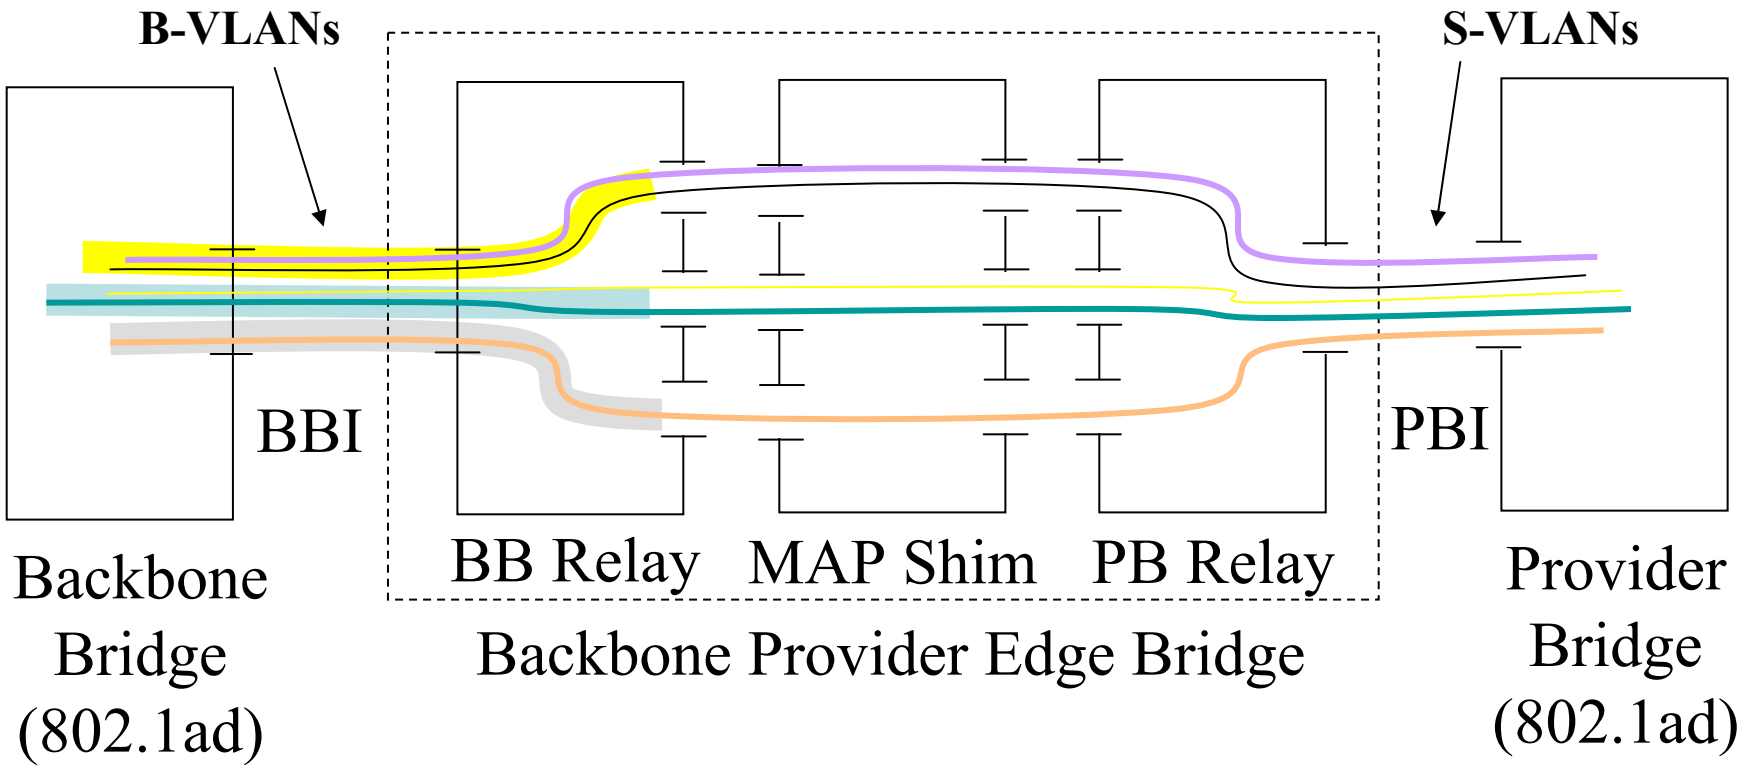

- MAP Shim performs encap/decap of frames to/from Provider Bridge Networks

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 (ES-VIDs)
  - PBN S-VIDs are local to the PBN
  - BBN ES-VIDs are local to the BBN

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 Backbone Provider Bridge Terminology
  - ES-TAG Extended Service VLAN TAG Field
  - ES-VID Extended Service VLAN ID
  - B-MAC Backbone MAC Address
  - B-VLAN Backbone VLAN (tunnel)
  - B-TAG Backbone TAG Field
  - B-VID Backbone VLAN ID (tunnel)

Extended Service VLAN IDs In Backbone

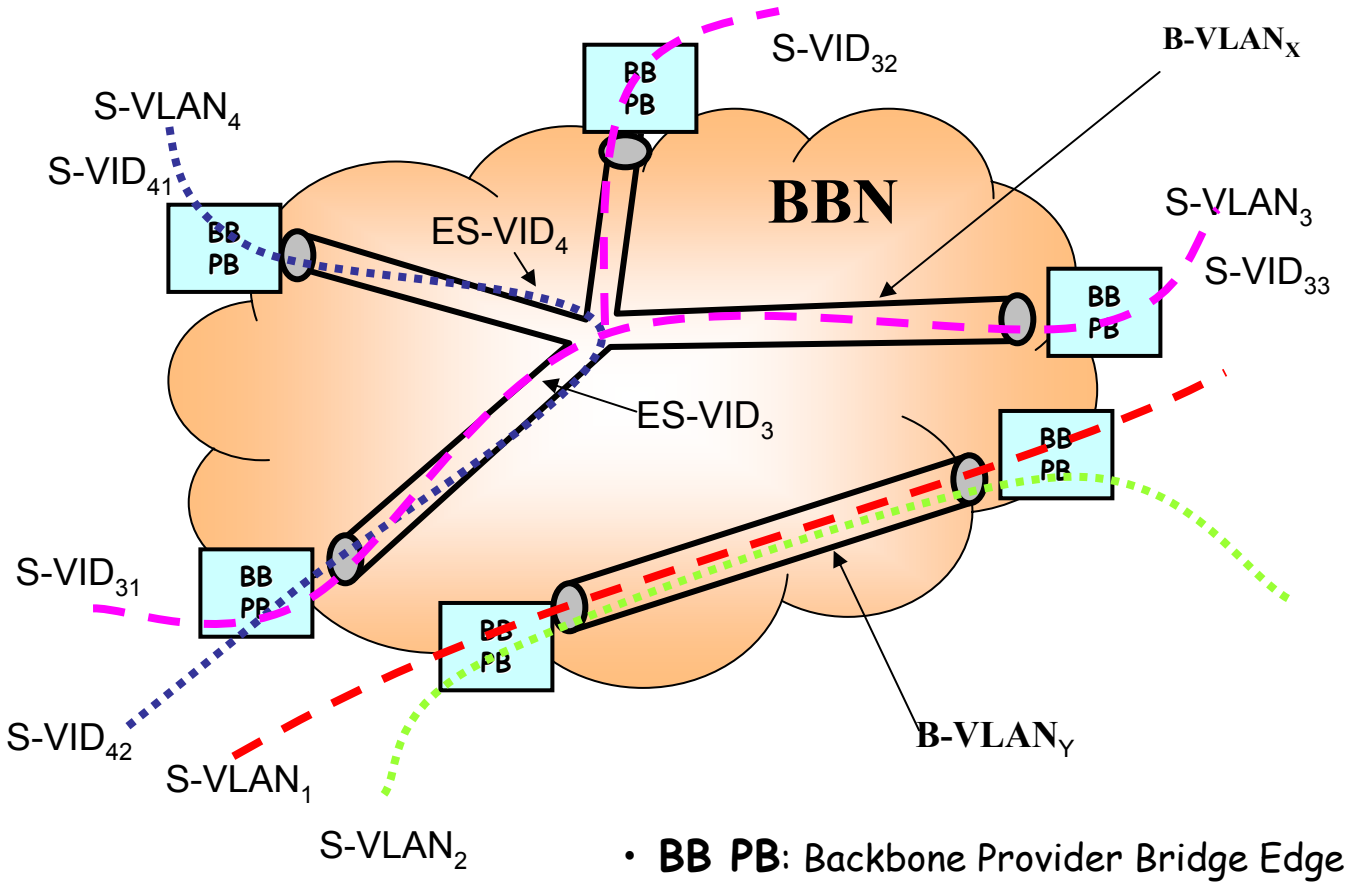

- **BB PB**: Backbone Provider Bridge Edge

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

Single ES-VID per S-VLAN

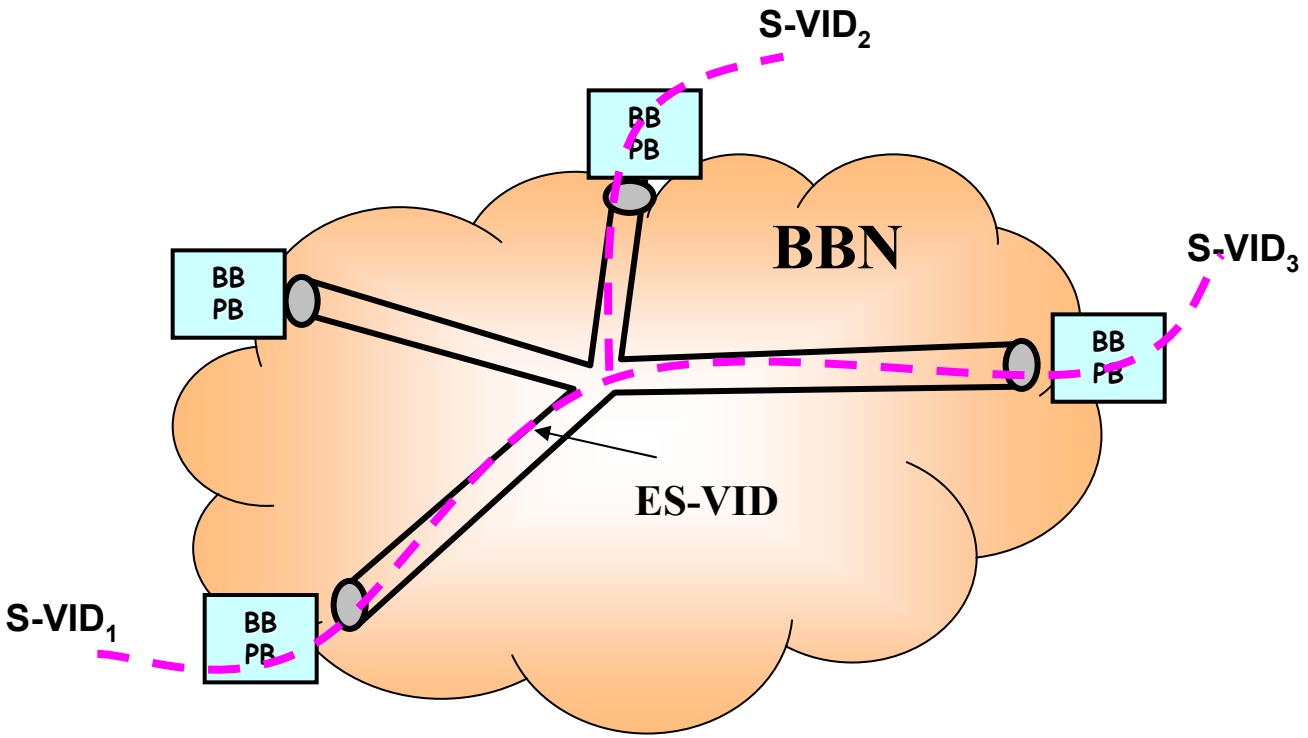

- Regardless of the ES-VID address size the map tables only have 4096 entries since only one ES-VID exists per S-VLAN and only 4096 S-VLANs exist per Provider Bridge.
- A different S-VID in each PBN maps to the ES-VID

Site Connectivity B-VLAN ID

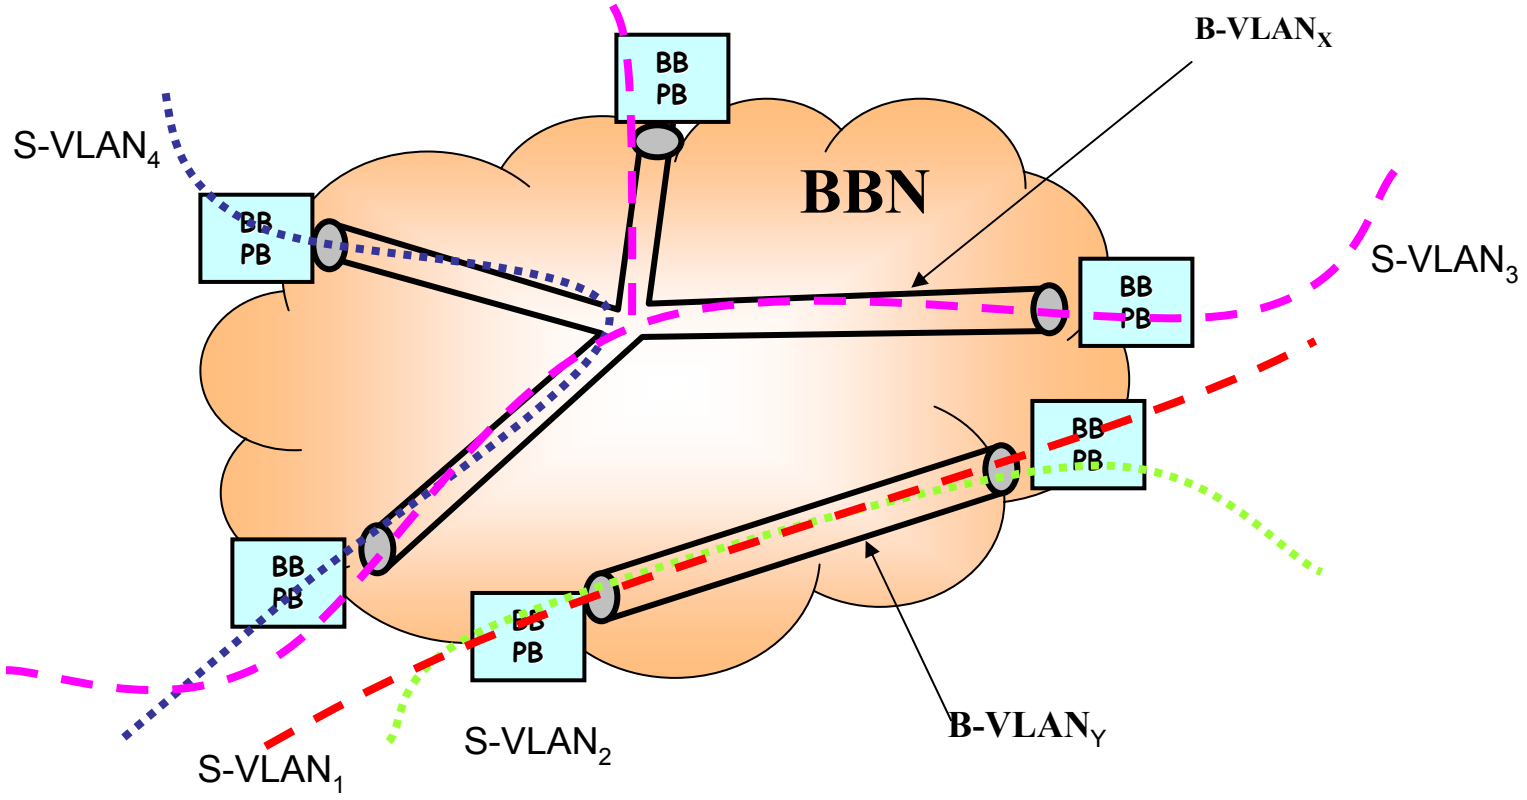

- B-VLANs are addressed like regular VLANs with a 12 bit B-VID
- B-VID and ES-VID need to be separate ID spaces to allow many S-VLANs to be carried in a single B-VLAN

Backbone POP MAC Address

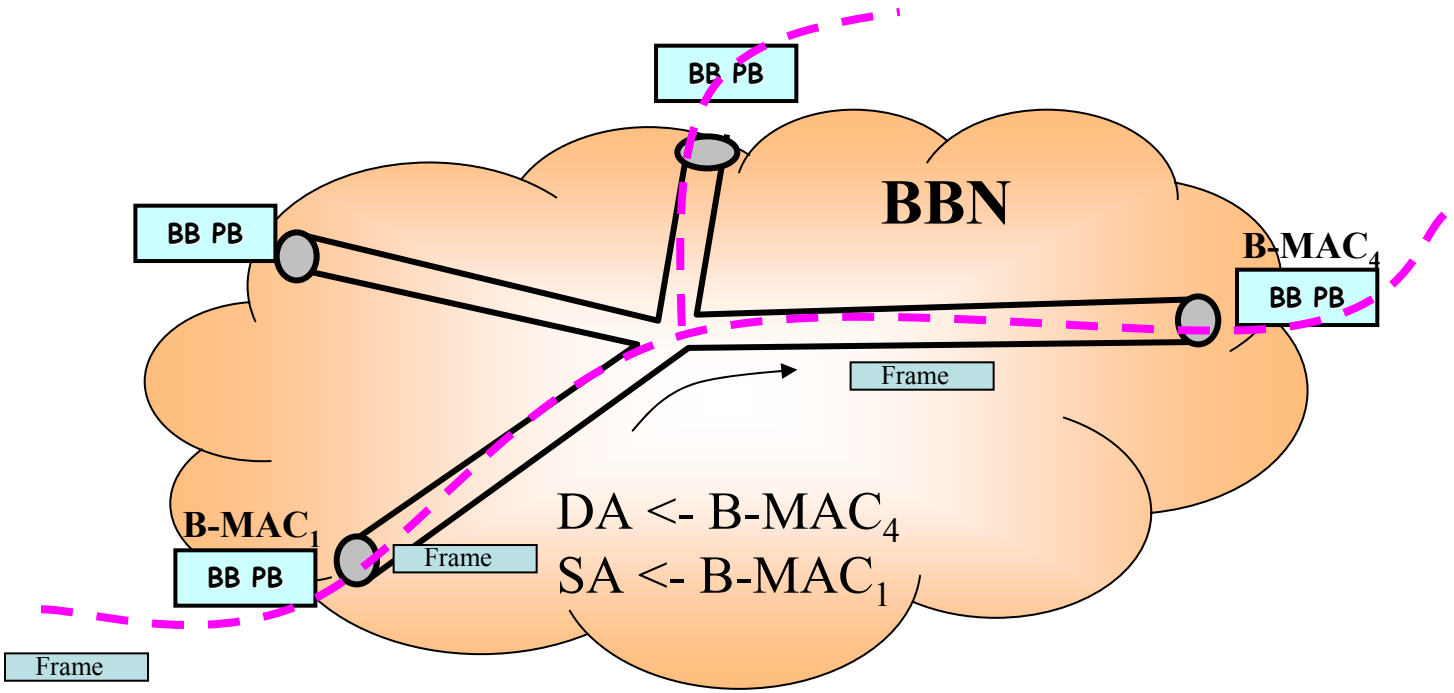

- B-MAC Addresses identify the Edge Backbone Provider Bridges (BB PB)
- B-MAC Addresses are learned by other Edge Backbone Edge Bridges
- The backbone edge MAC address determines which edge on the B-VLAN will receive the frame.
- Frames may be flooded by sending with broadcast or multicasts DA B-MACs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

Customer/Provider Addresses

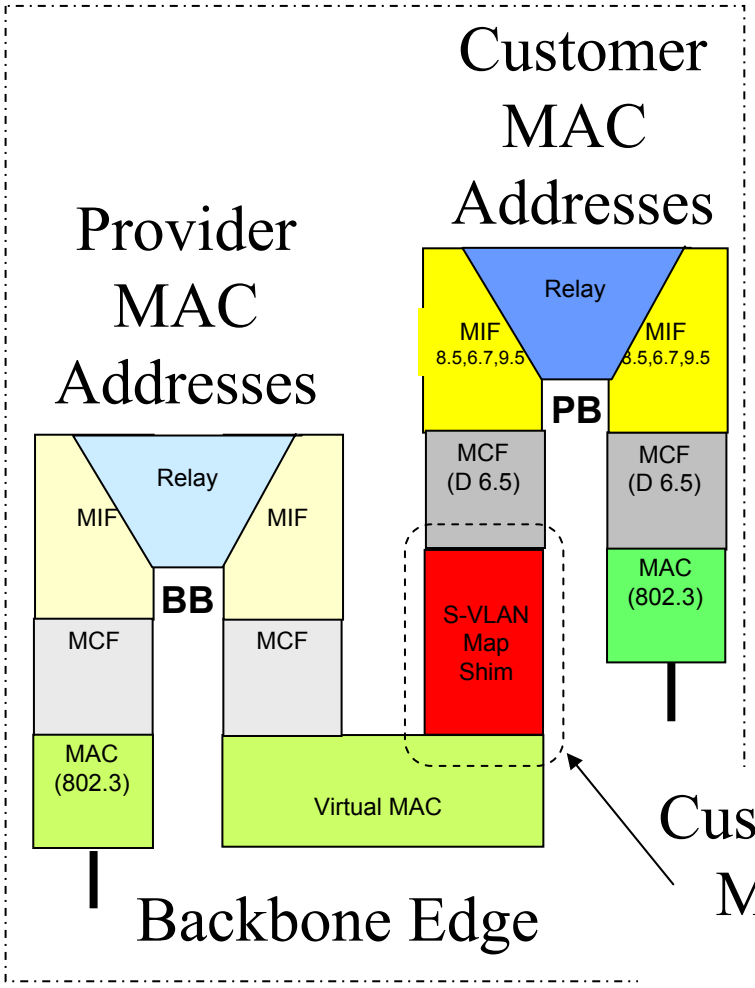

- PB Relay Learns Customer Address Per S-VLAN
- BB Relay Learns Provider Addresses Per B-VLAN
- MAP Shim Learns Correlated Customer and Provider MAC Addresses per S-VLAN

Backbone Frame Format

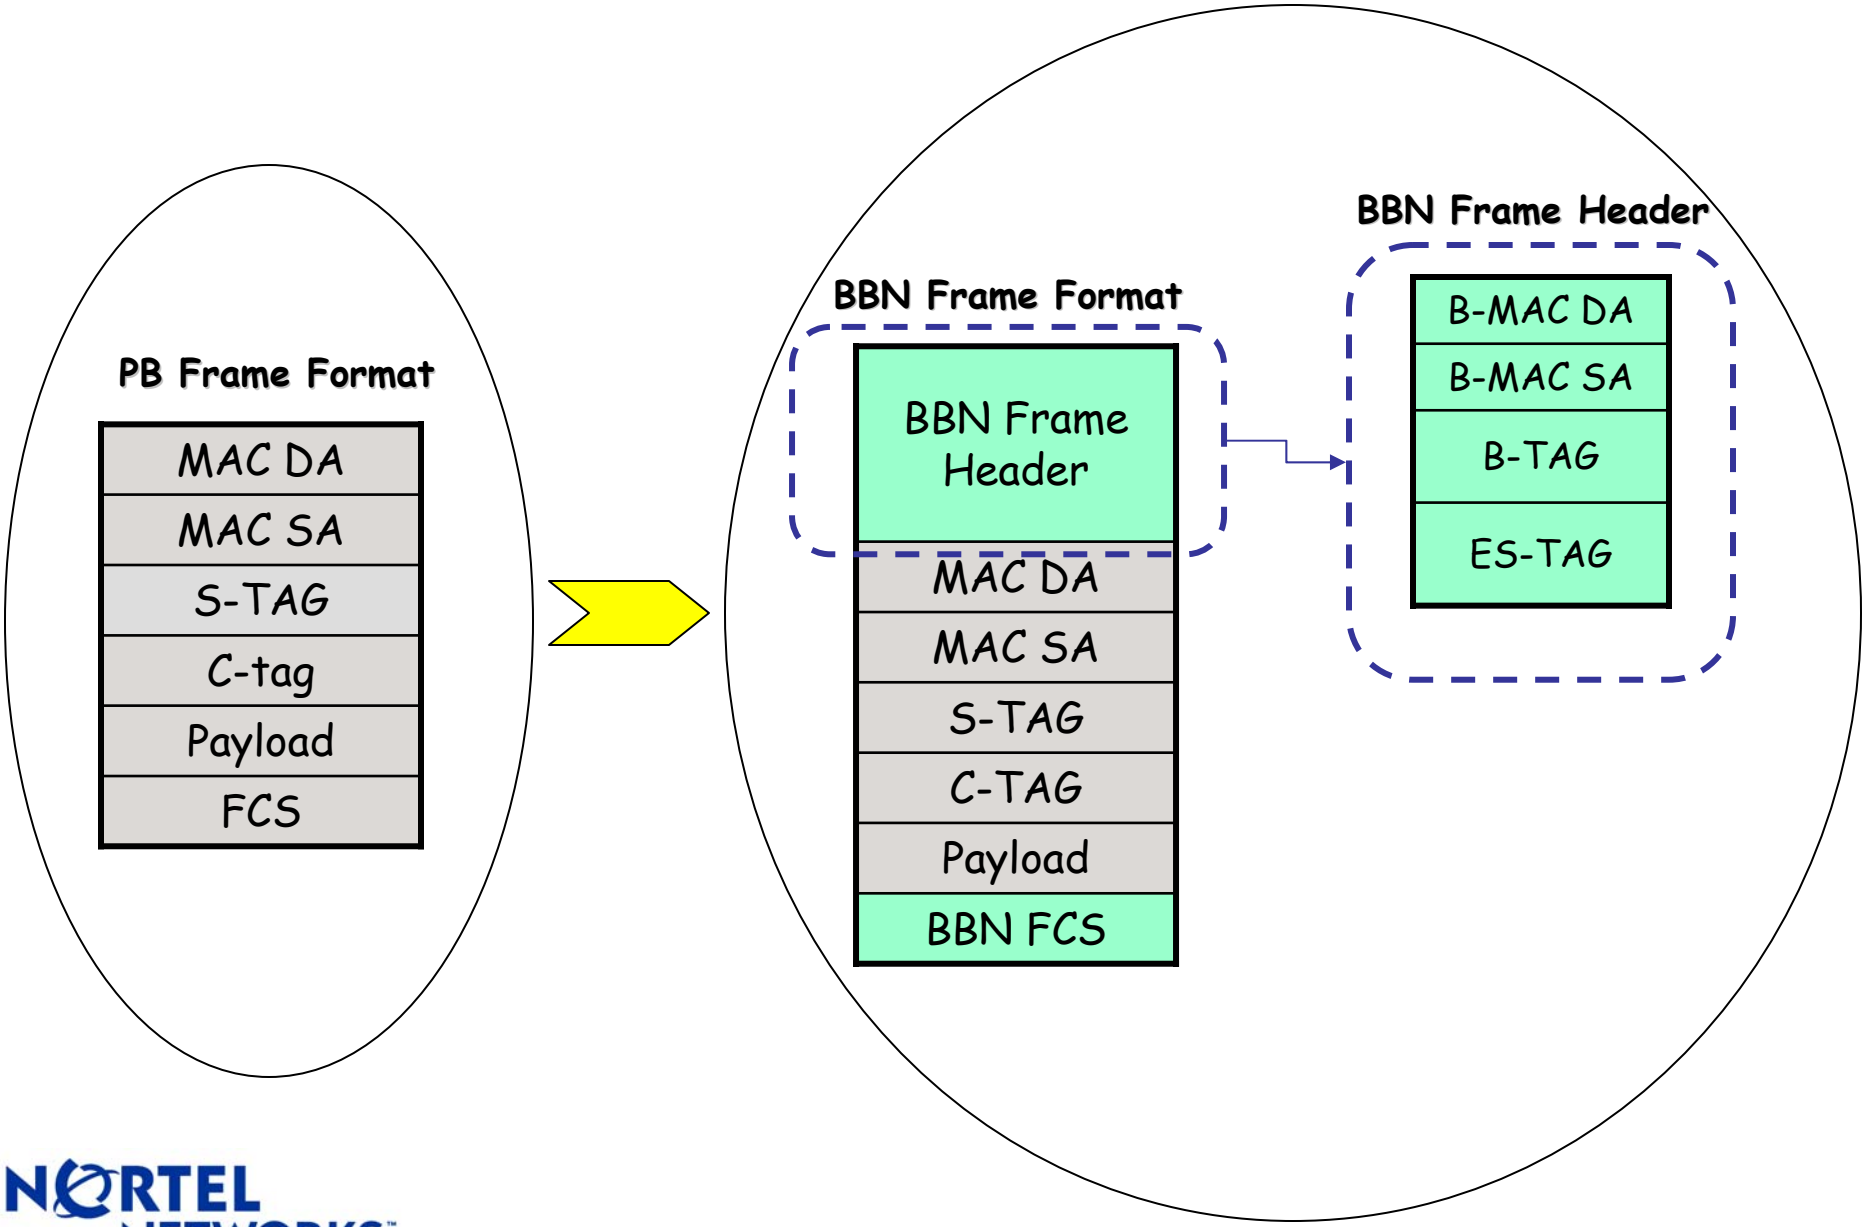

MAP Shim Correlation Table

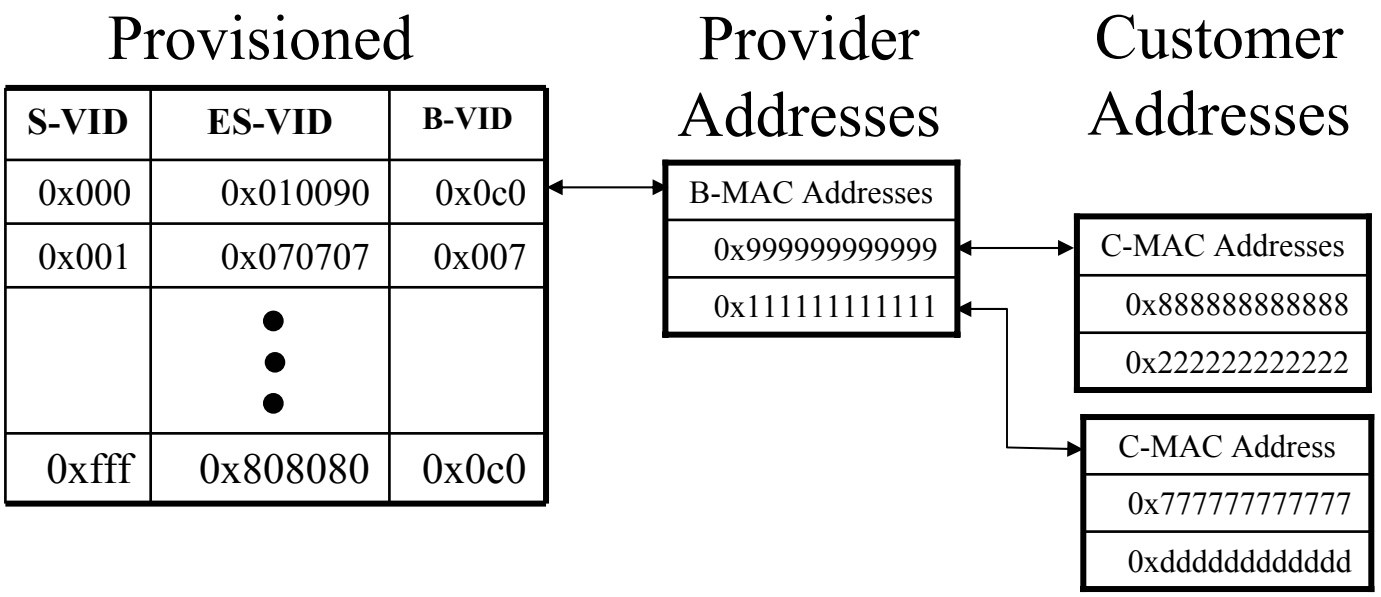

- In the beginning the MAP Shim is provisioned with the correlation between the S-VID, ES-VID, and B-VID
- During operation the MAP Shim learns both B-MAC addresses and C-MAC addresses
- The MAP Shim keeps track of which C-MAC addresses are behind which B-MAC
- The correlation data is used to encapsulate 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 ES-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
  - ES-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 Backbone Provider 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 Backbone Provider Bridges

Backup Material

PAR Slides

Backbone Provider Bridge PAR

- Allows scaling Provider Bridge networks to support a large population of users
- May use reuse much of 802.1Q bridge technology
- Recommend a new standard to allow removing bridge functions which are not important to Backbone Provider Bridge

PAR Title

- Standard for Local and Metropolitan Area Networks – Virtual Bridged Metropolitan Backbone Provider Networks

PAR Scope

- To develop bridge protocols, architecture, and management compatible and interoperable with Provider Bridged(1) Network protocols and equipment allowing interconnection of multiple Provider Bridged Networks in such a way as to extend the Customer MAC Service Instances provided by these networks between the multiple Provider Bridged Networks, to allow scaling to at least 2^20 Service Virtual LANs, and to support management of the Customer MAC Service Instances.

-1-IEEE Std. 802.1ad

NORTEL NETWORKS

Purpose

- This standard will enable a Service Provider to scale the number of Service Virtual LANs in a Provider Network by interconnecting many independent Provider Bridged (IEEE Std. 802.1ad) Networks while interconnecting the Service Virtual LANs provided by these Provider Bridged Networks, and provide for interoperability and consistent management.

Multicast Relay Scaling

Backbone Provider Bridging

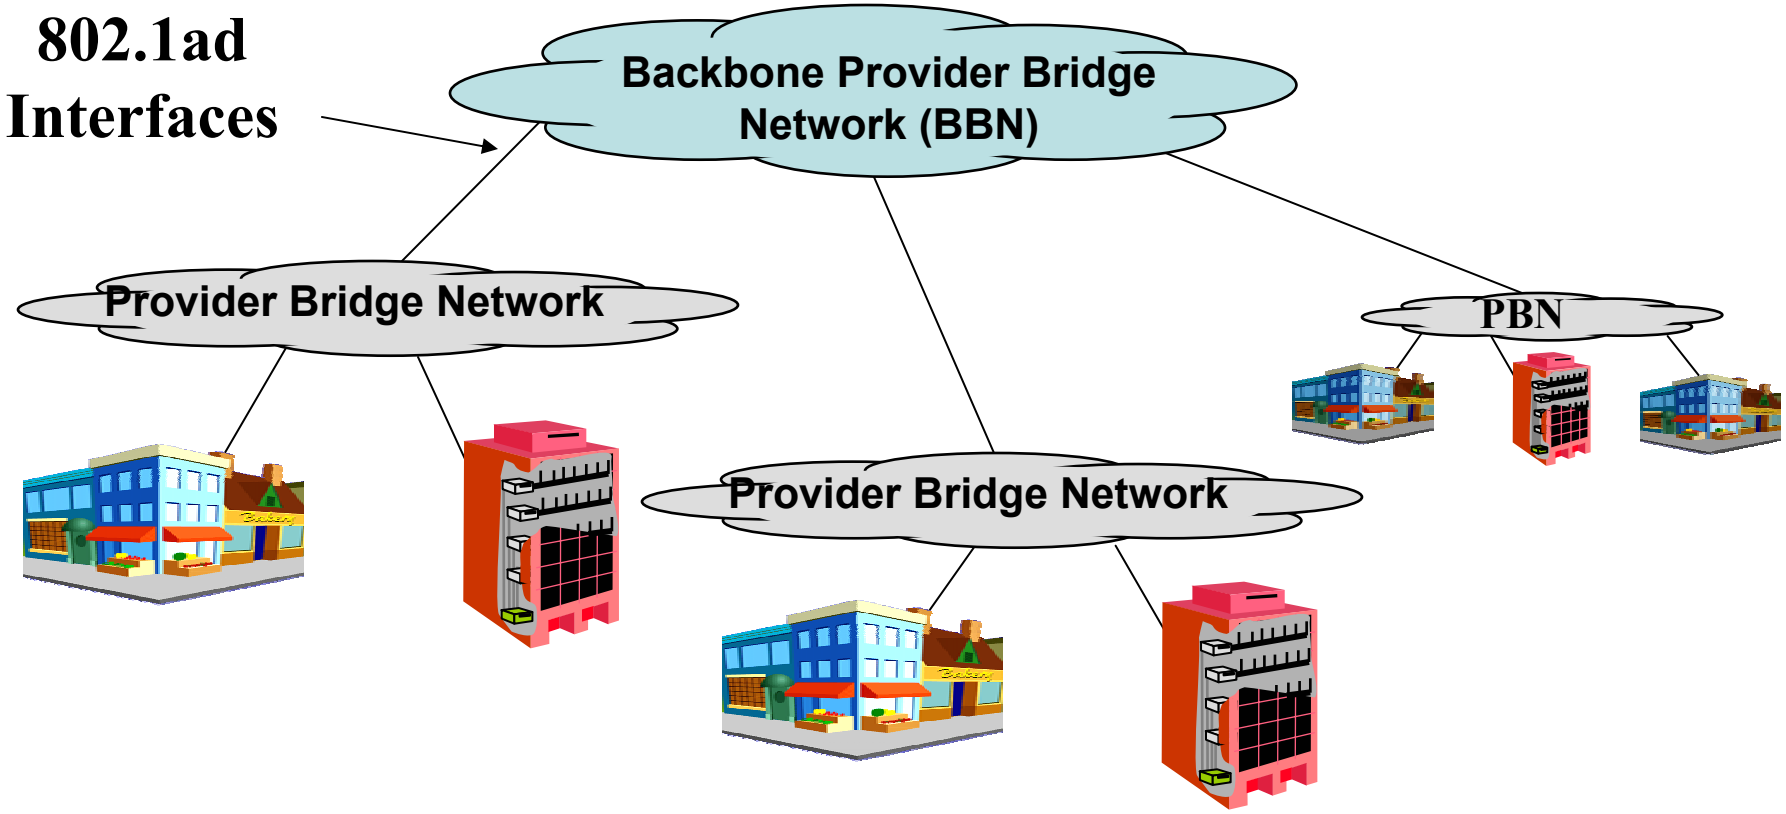

Terminology

- 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 Backbone Provider Bridge Terminology
  - ES-VID Extended Service VLAN ID
  - B-VLAN Backbone VLAN (tunnel)
  - B-VID Backbone VLAN ID (tunnel)
  - B-MCD Backbone Multicast Domain
  - B-TAG Backbone TAG Field
  - B-MAC Backbone MAC Address

A Simple Backbone Provider Bridge Network

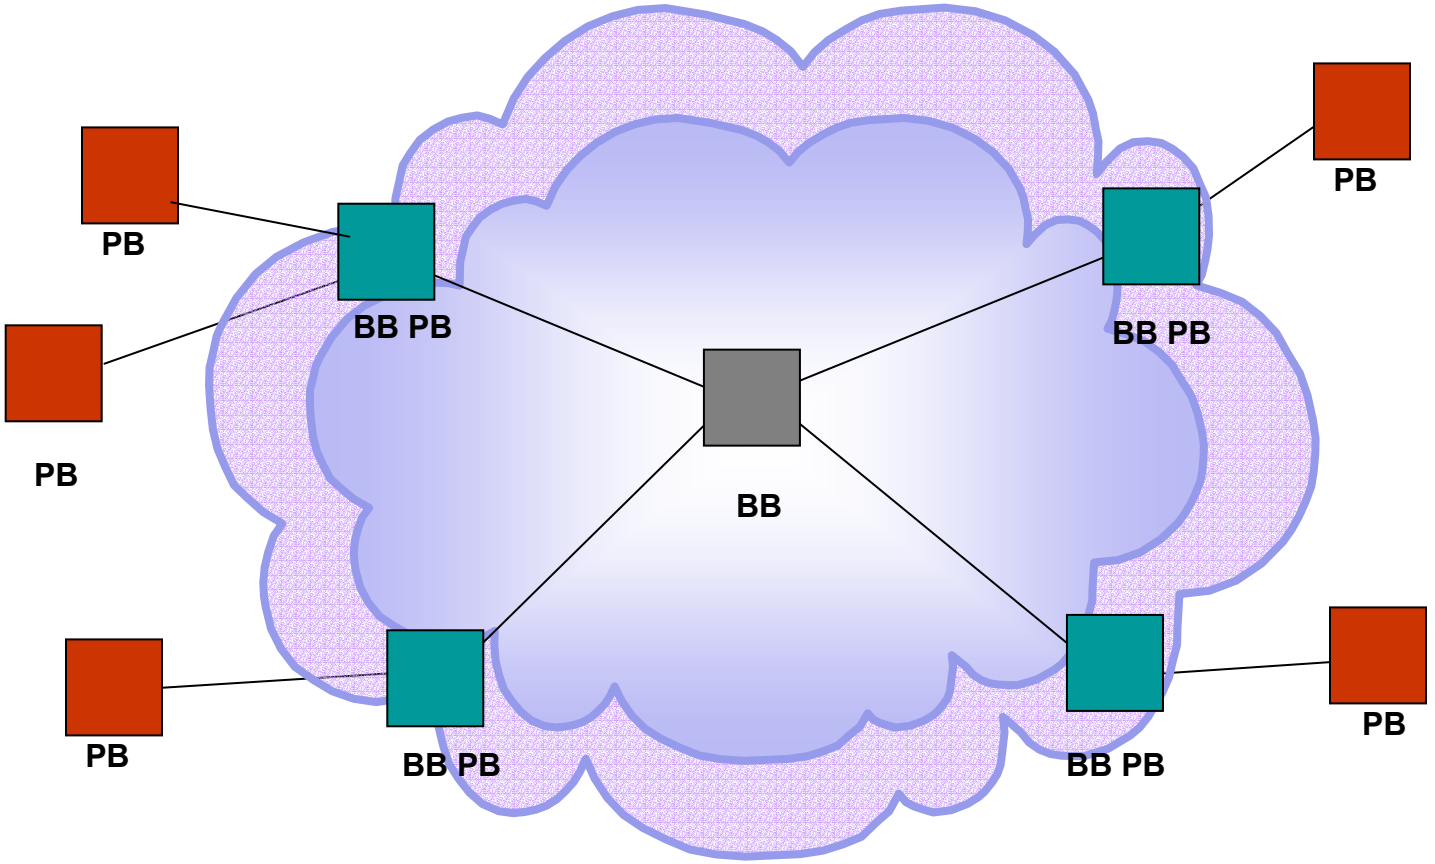

- PB: Provider Bridge (as defined by 802.1ad)
- BB PB: Backbone Provider Bridge Edge
- BB: Backbone Provider Bridge

Simple BBN Principles

- BB PB edge encapsulates received PBN frames with BBN header
- BBN header includes
  - Extended Service VLAN Identifier (ES-VID)
    - Identifies the S-VLAN associated with the PBN S-VIDs on the BBN
    - Must be large enough to support millions of S-VLANs
  - Backbone VLAN Identifier (B-VID)
    - Identifies a backbone VLAN (B-VLAN or tunnel) that is used to transport the S-VLANs over the BBN
    - A B-VLAN(tunnel) can be point-to-point or multi-point in nature
    - The B-VID must have a large enough address space to support all available multi-point tunnels among BB bridges
  - Backbone POP Address
    - Addresses POP within Site Connectivity

Use Hierarchical Architecture to Scale BBN size

- Support of B-VLANs (i.e., multicast) with large number of bridges is challenging
  - frame replicators to large number of points limit performance
- Hierarchy of BB bridges creates small multicast domains
  - each domain has a small number of bridges, which limits number of multi-point tunnels and number of replications

A Two Layer Hierarchical BBN

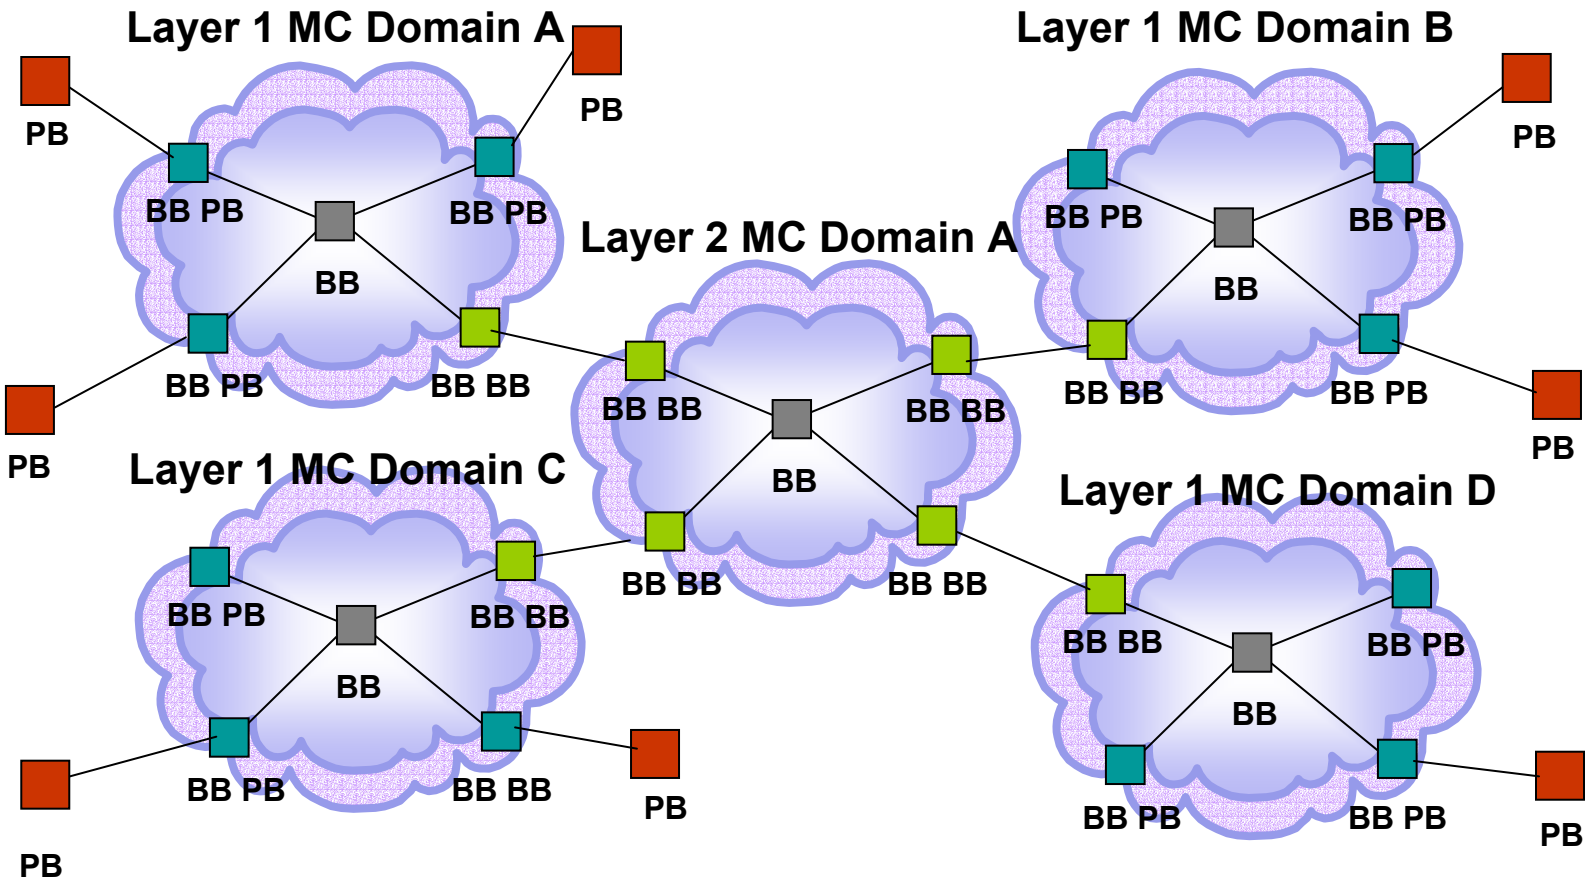

- **PB**: Provider Bridge (as defined by 802.1ad)
- **BB PB**: Backbone Provider Bridge Edge
- **BB**: Backbone Provider Bridge
- **BB BB**: Backbone Provider Bridge Layer Edge

NORTEL NETWORKS

Hierarchical BBN Principles

- **BB PB edge**
  - Encapsulates received PBN frame with BBN header
    - swaps S-VID to/from a much large ES-VID
    - Creates a B-VID from the ES-VID
  - De-encapsulates frames to be transmitted to the PBN by stripping the BBN header
    - Swaps the ES-VID to a S-VID for the PBN
    - Removes the final B-VID
  - Both S-VID and ES-VID identify the S-VLAN carried through the PBNs and BBN

- **BB BB layer edge bridge**
  - swaps the B-VID to a new B-VID based on the ES-VID
    - The new B-VID allow transport over the current BBN multicast domain (MC-DOM)
  - source route addressing (with Backbone Connectivity identifier stacking) can also be used to avoid the need for table lookups and B-VID swapping at layer boundaries
  - The ES-VID is the same throughout the BBN
  - The ES-VID is swapped with the S-VID at the BB-PB edge
  - B-VID must be large enough to address all possible multi-point tunnels within a given layer domain (e.g., 12 bits is enough to support 12 BB bridges in a layer. More bits are required for more bridges).

- **Scalability**
  - Hierarchical BBN can have as many layers as required

Customer, PB, BB Spanning Trees

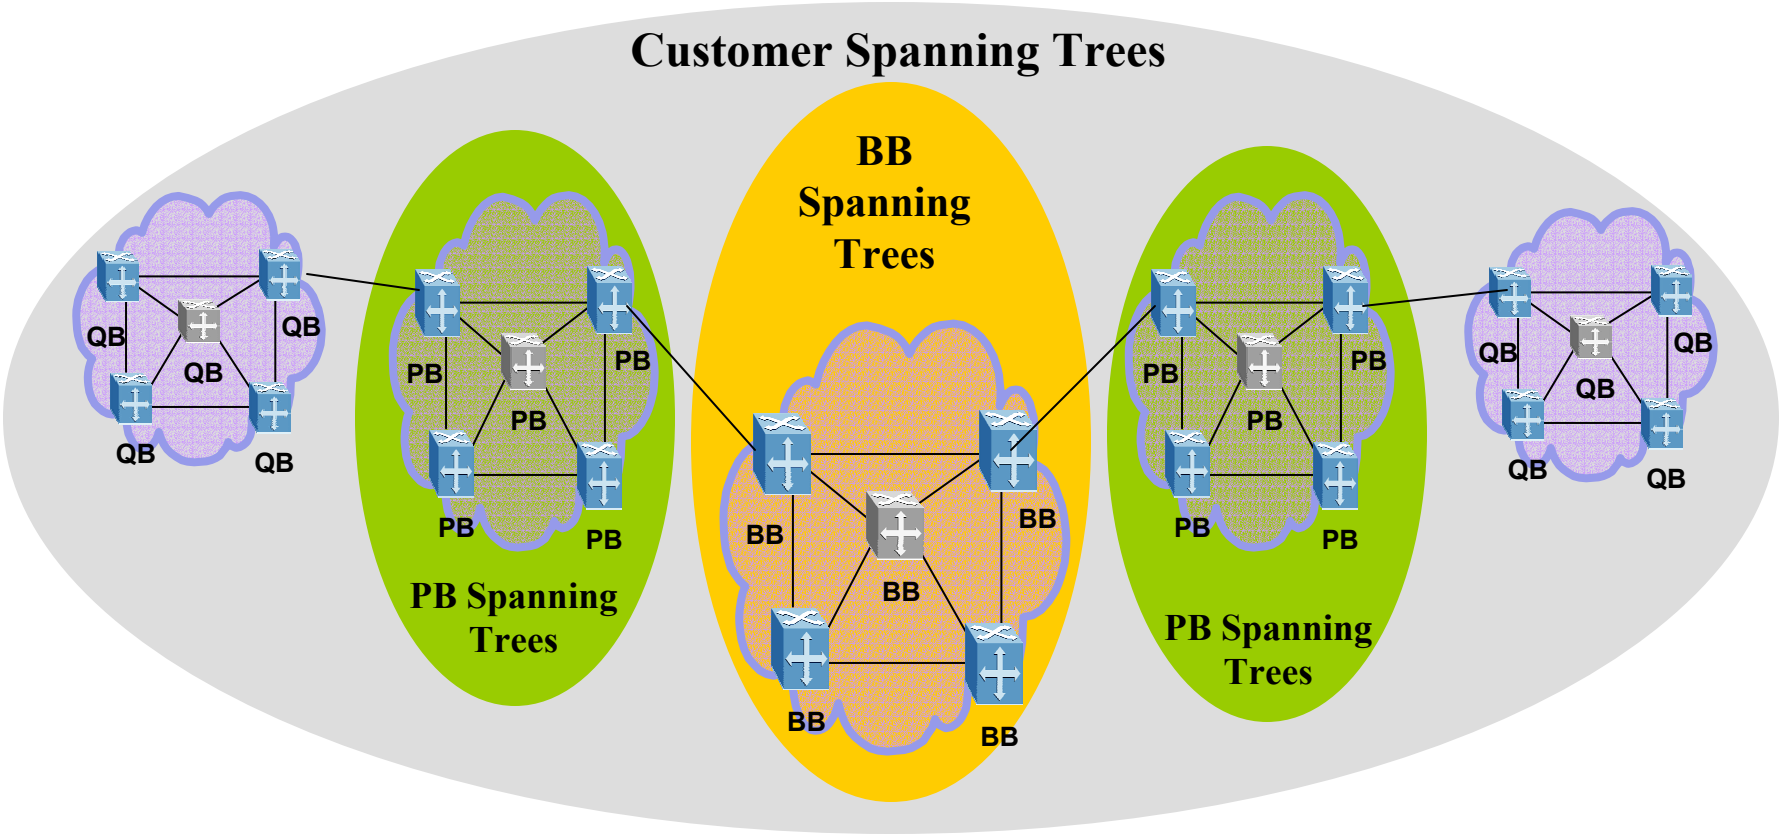

- Customer spanning trees may extend over Provider Network
- PB Network and BB Network spanning trees must be decoupled to scale the provider network