Backbone Provider Bridging Networks

A Highly Scalable VLAN (Multicast) Architecture

Paul Bottorff

Version 1.0

July 12, 2004

NORTEL NETWORKS™

A Provider Bridge Scaling Solution "Backbone Provider Bridging"

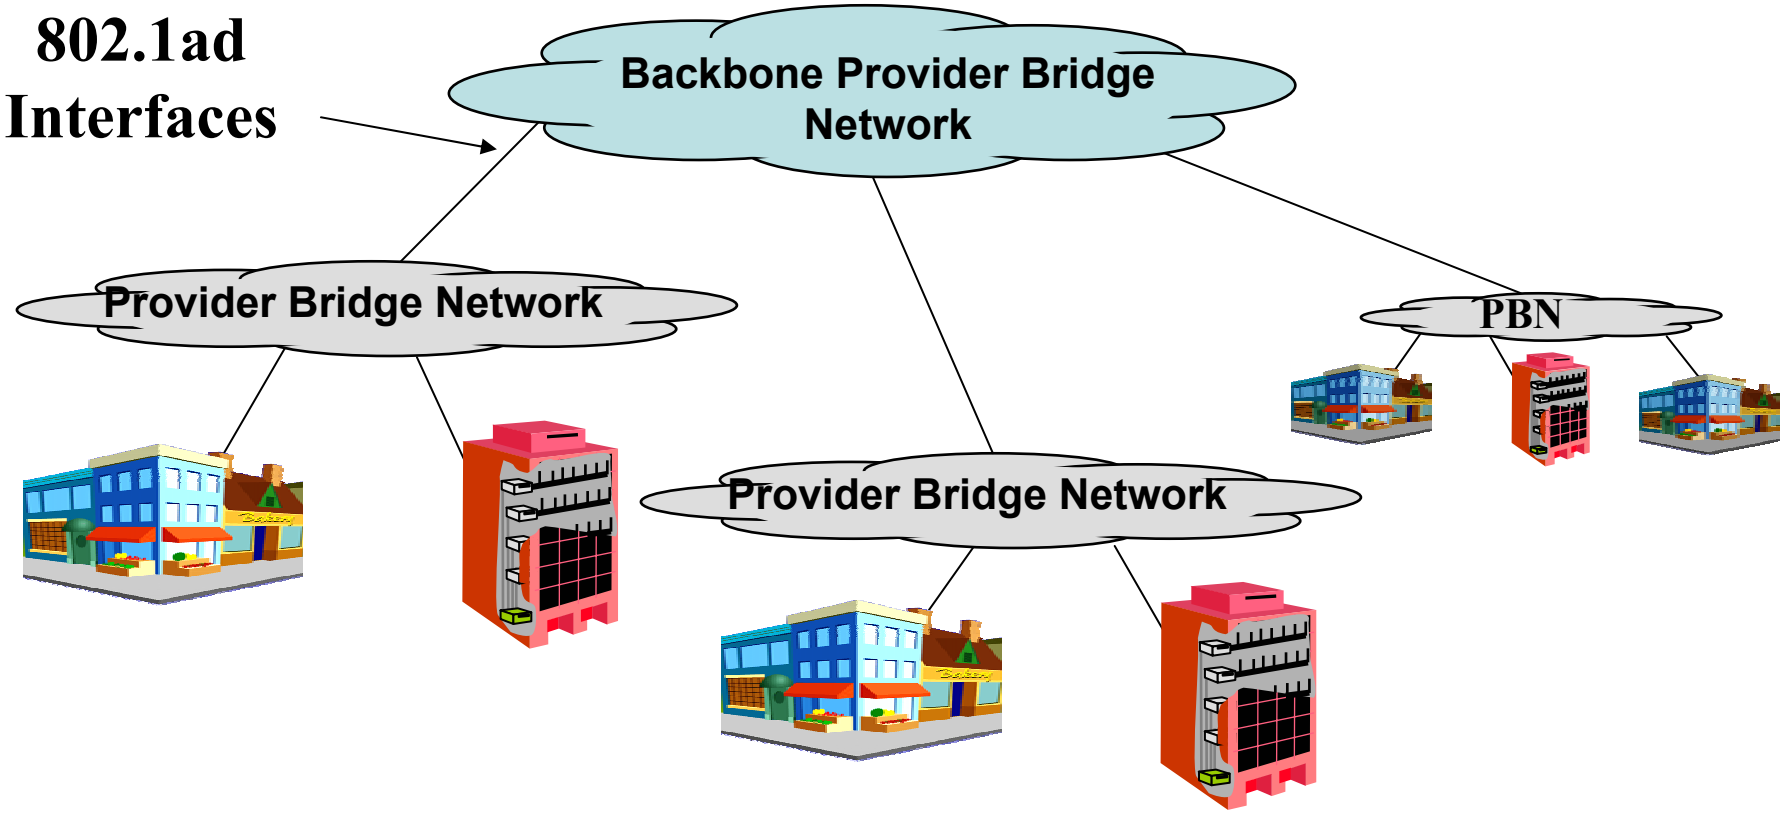

NORTEL NETWORKS

Ethernet Service Types

MEF Ethernet Virtual Connections(EVCs)

E-LINE
Router Mesh

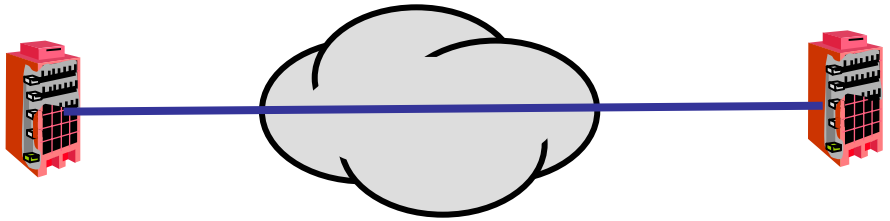

Pt-Pt, Like
Duplex Ethernet
Any-to-any

E-TREE
Hub & Spoke

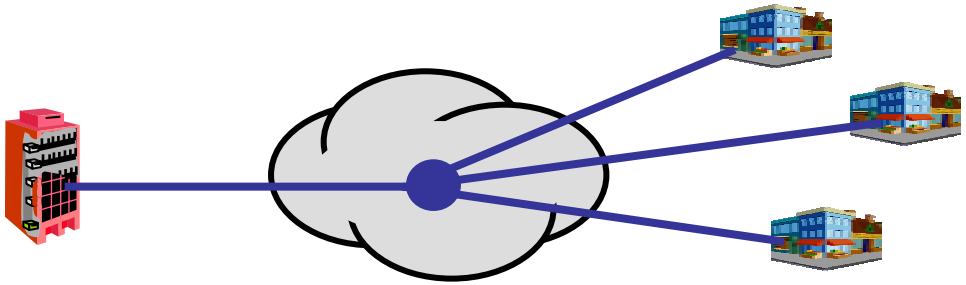

Pt-MPt, Like
EPON Ethernet,
Root-to-Leaf and
Leaf-to-Root

E-LAN
Multi-Site

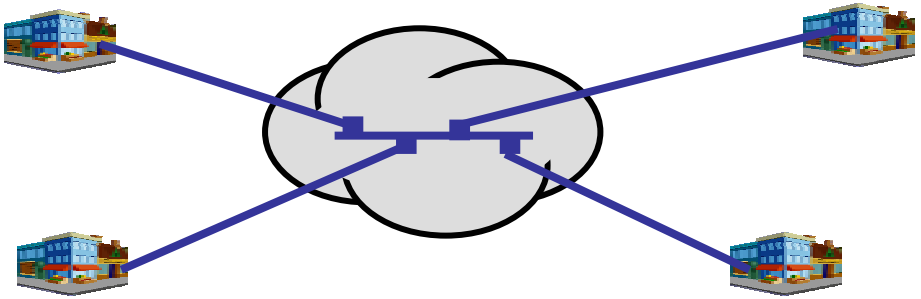

MPt, Like VLAN,
Any-to-any

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 Promis
  - 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

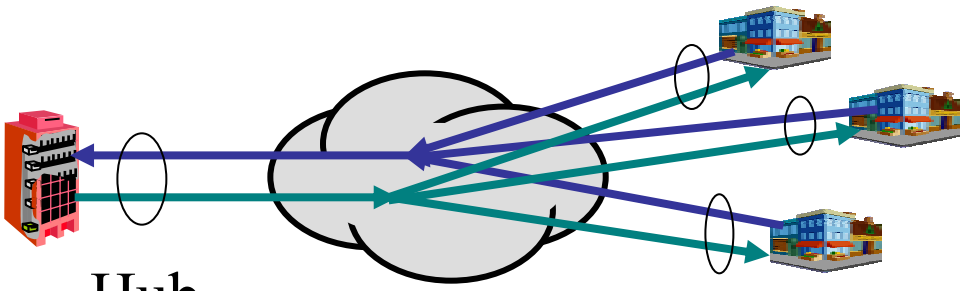

Hub
Port

Spoke
Ports

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

NORTEL NETWORKS™

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

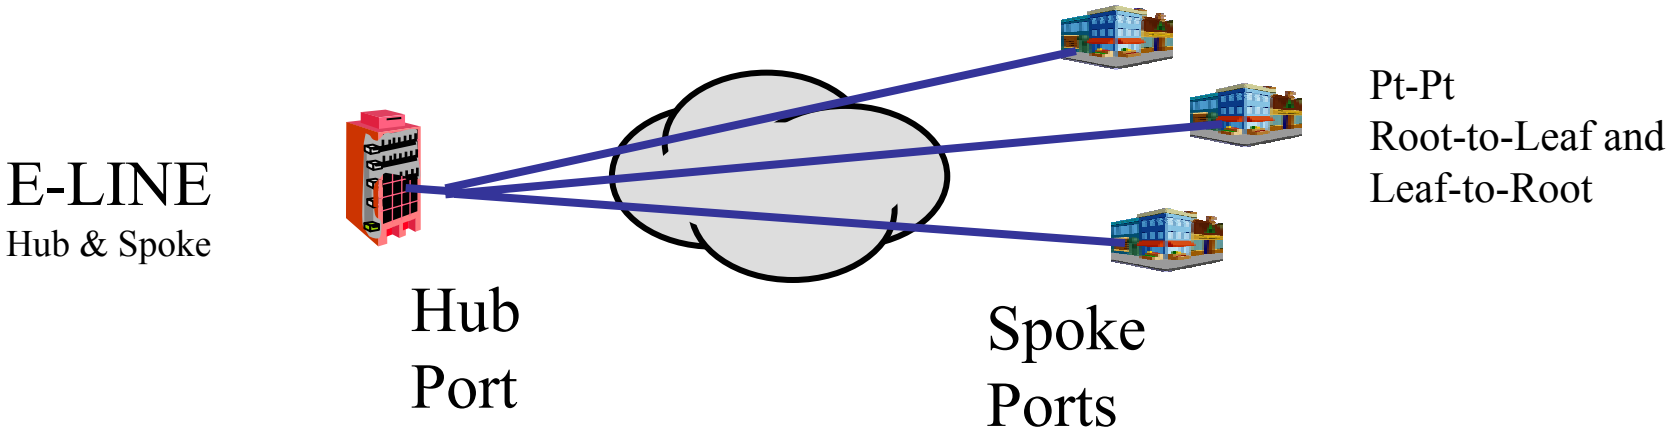

- Hub port would usually be multipexed to allow the multiple Pt-Pt attachments.
- Each E-LINE is a seperate managed S-VLAN
- This arrangement allows use of E-LINE management, billing, and QoS
- Many more S-VLANs are required

NORTEL NETWORKS

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

NORTEL NETWORKS™

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	>10,000 **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

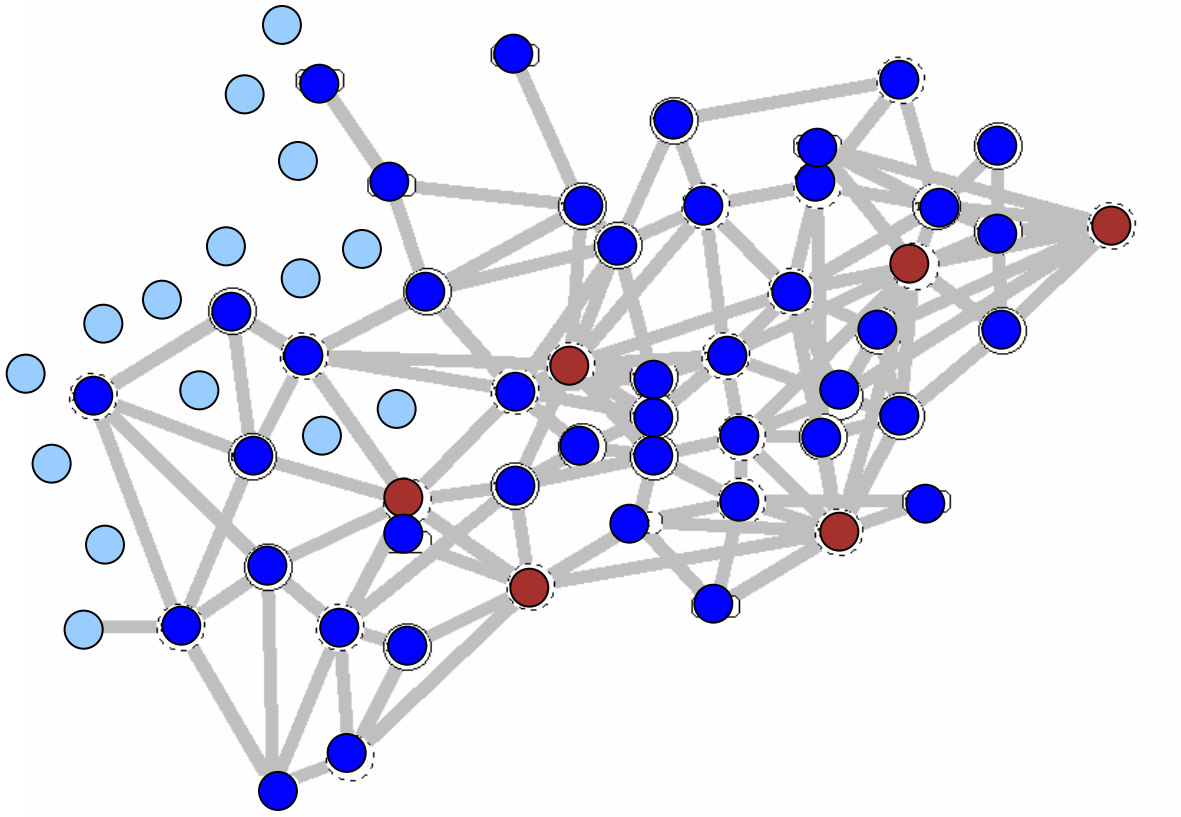

NORTEL NETWORKS

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

NORTEL NETWORKS™

Large Service Address Space Needed

- Carriers need to separate the service address space to allow administration of networks
  - Allocation 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

New Project for Backbone Provider Bridge

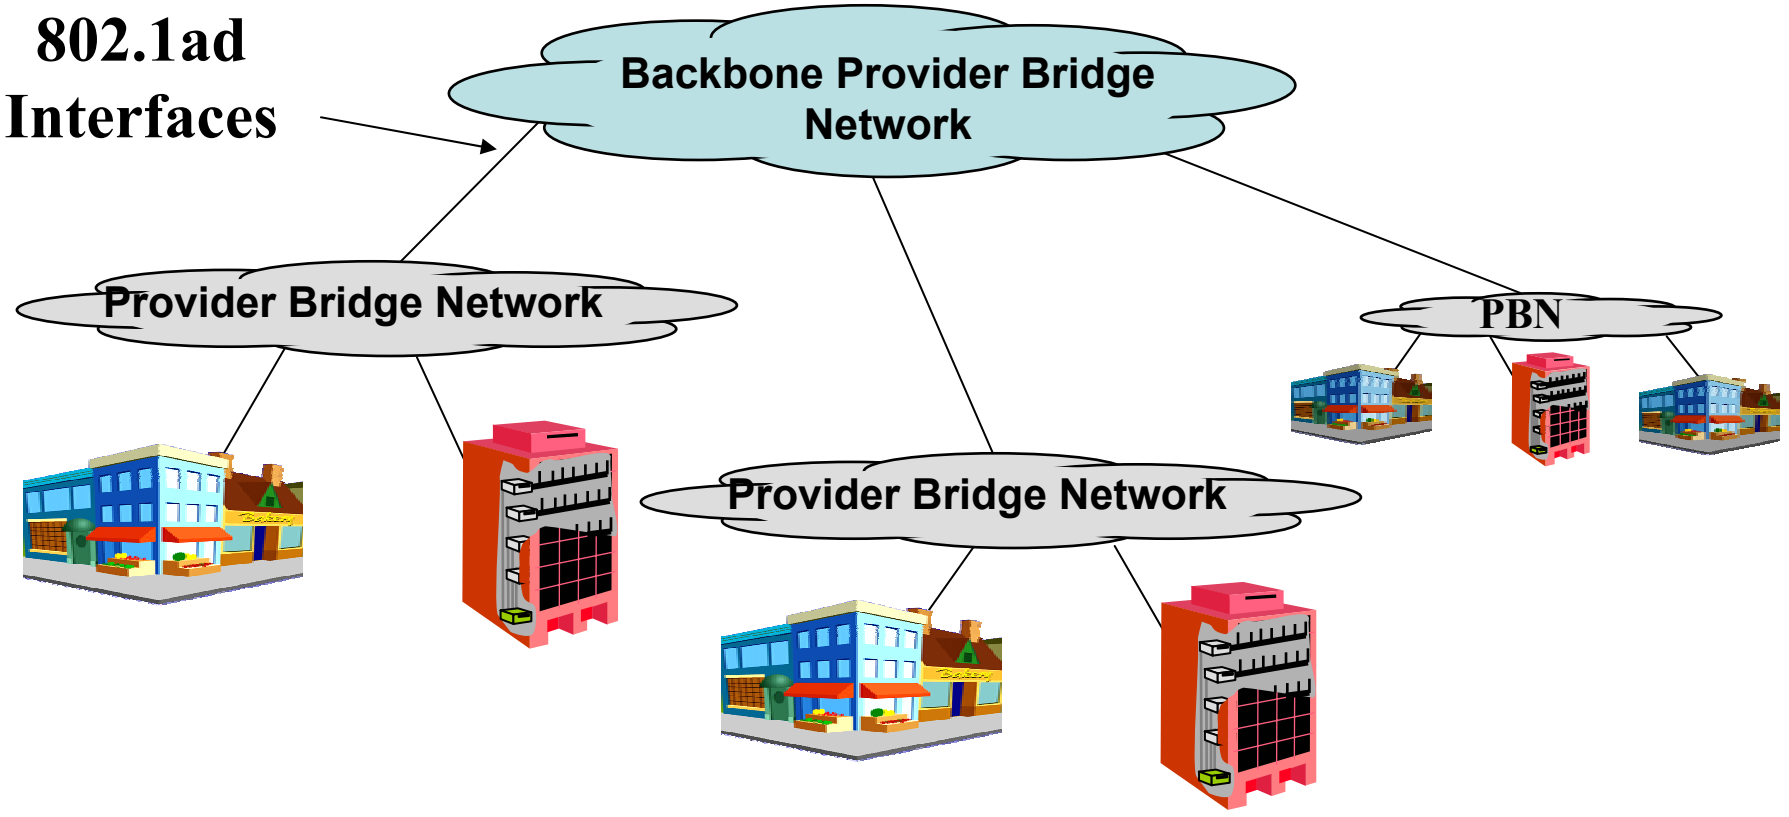

Possible Directions

Backbone Provider Bridge Model

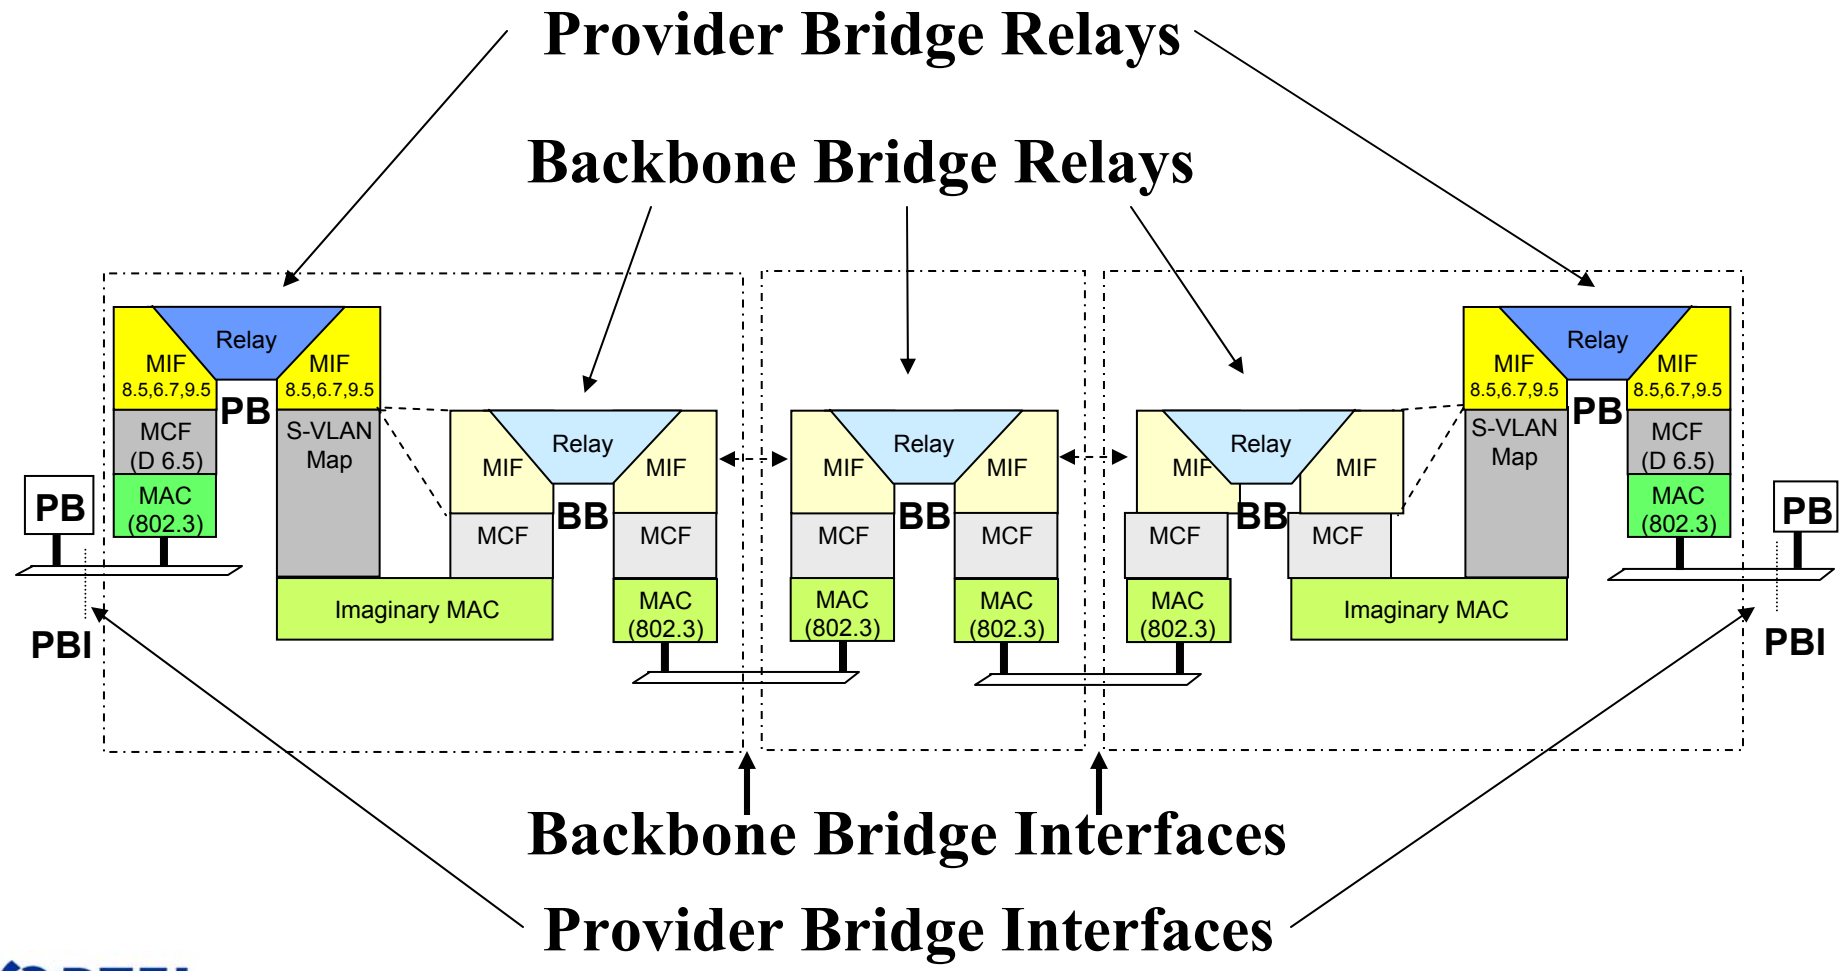

Backbone Provider Bridge Network

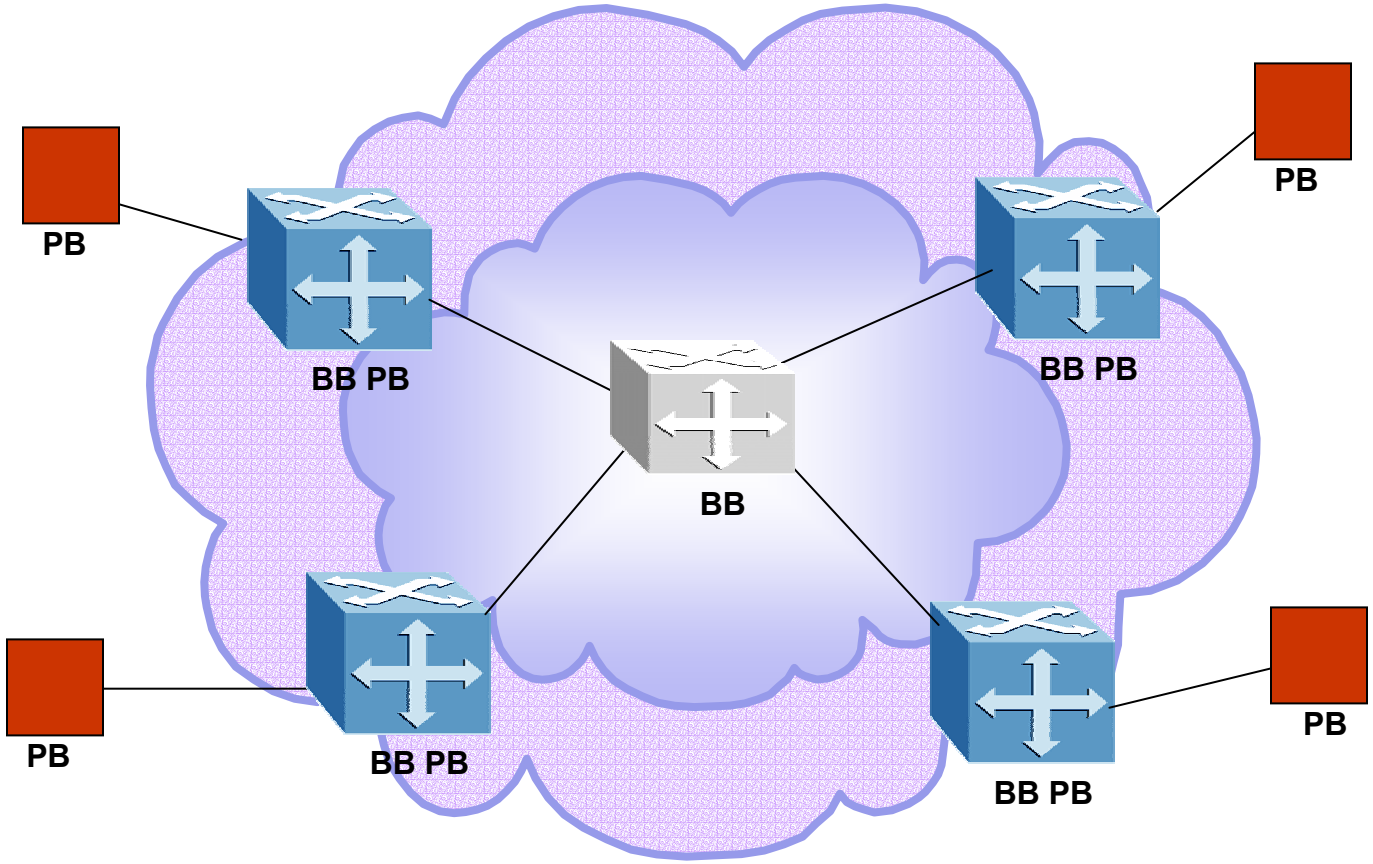

- **PB**: Provider Bridge (as defined by 802.1ad)
- **BB PB**: Backbone Provider Bridge Edge
- **BB**: Backbone Provider Bridge

NORTEL NETWORKS

Backbone Provider Bridging Principles

- BBN encapsulates PBN frame with BBN header
- BBN header consists of
  a) Service Instance identifier
     - Identifies the BBN service instance of the PBN flow
     - 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
     Addresses POP within Site Connectivity
- PBN service instances (S-VLANs) map to BBN service instances
  - PBN service instances are local to the PBN
  - BBN service instances are local to the BBN

NORTEL NETWORKS

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-MCB Backbone Multicast Domain
  - B-TAG Backbone TAG Field
  - B-MAC Backbone MAC Address

BBN Provides Multi-Point B-VLANs Between PBNs

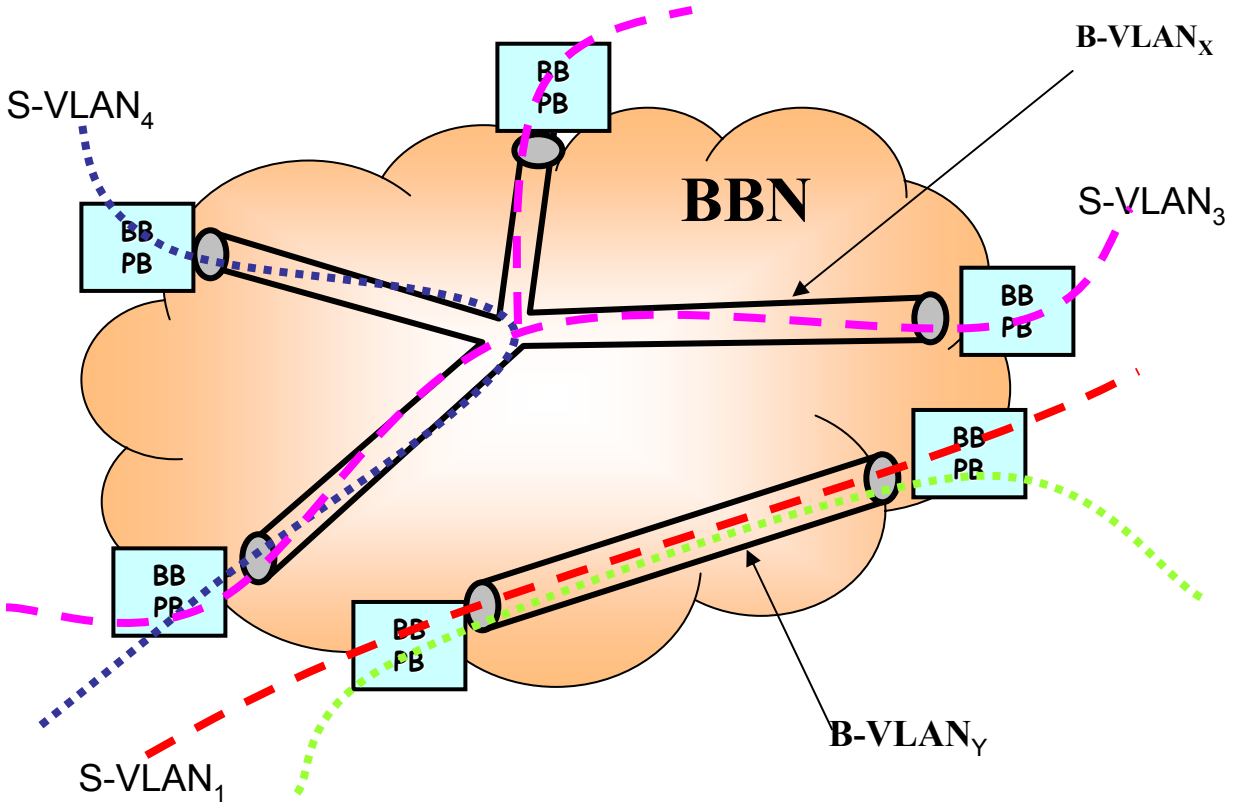

- **BB PB**: Backbone Provider Bridge Edge

- 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.

NORTEL NETWORKS

Backbone Frame Format

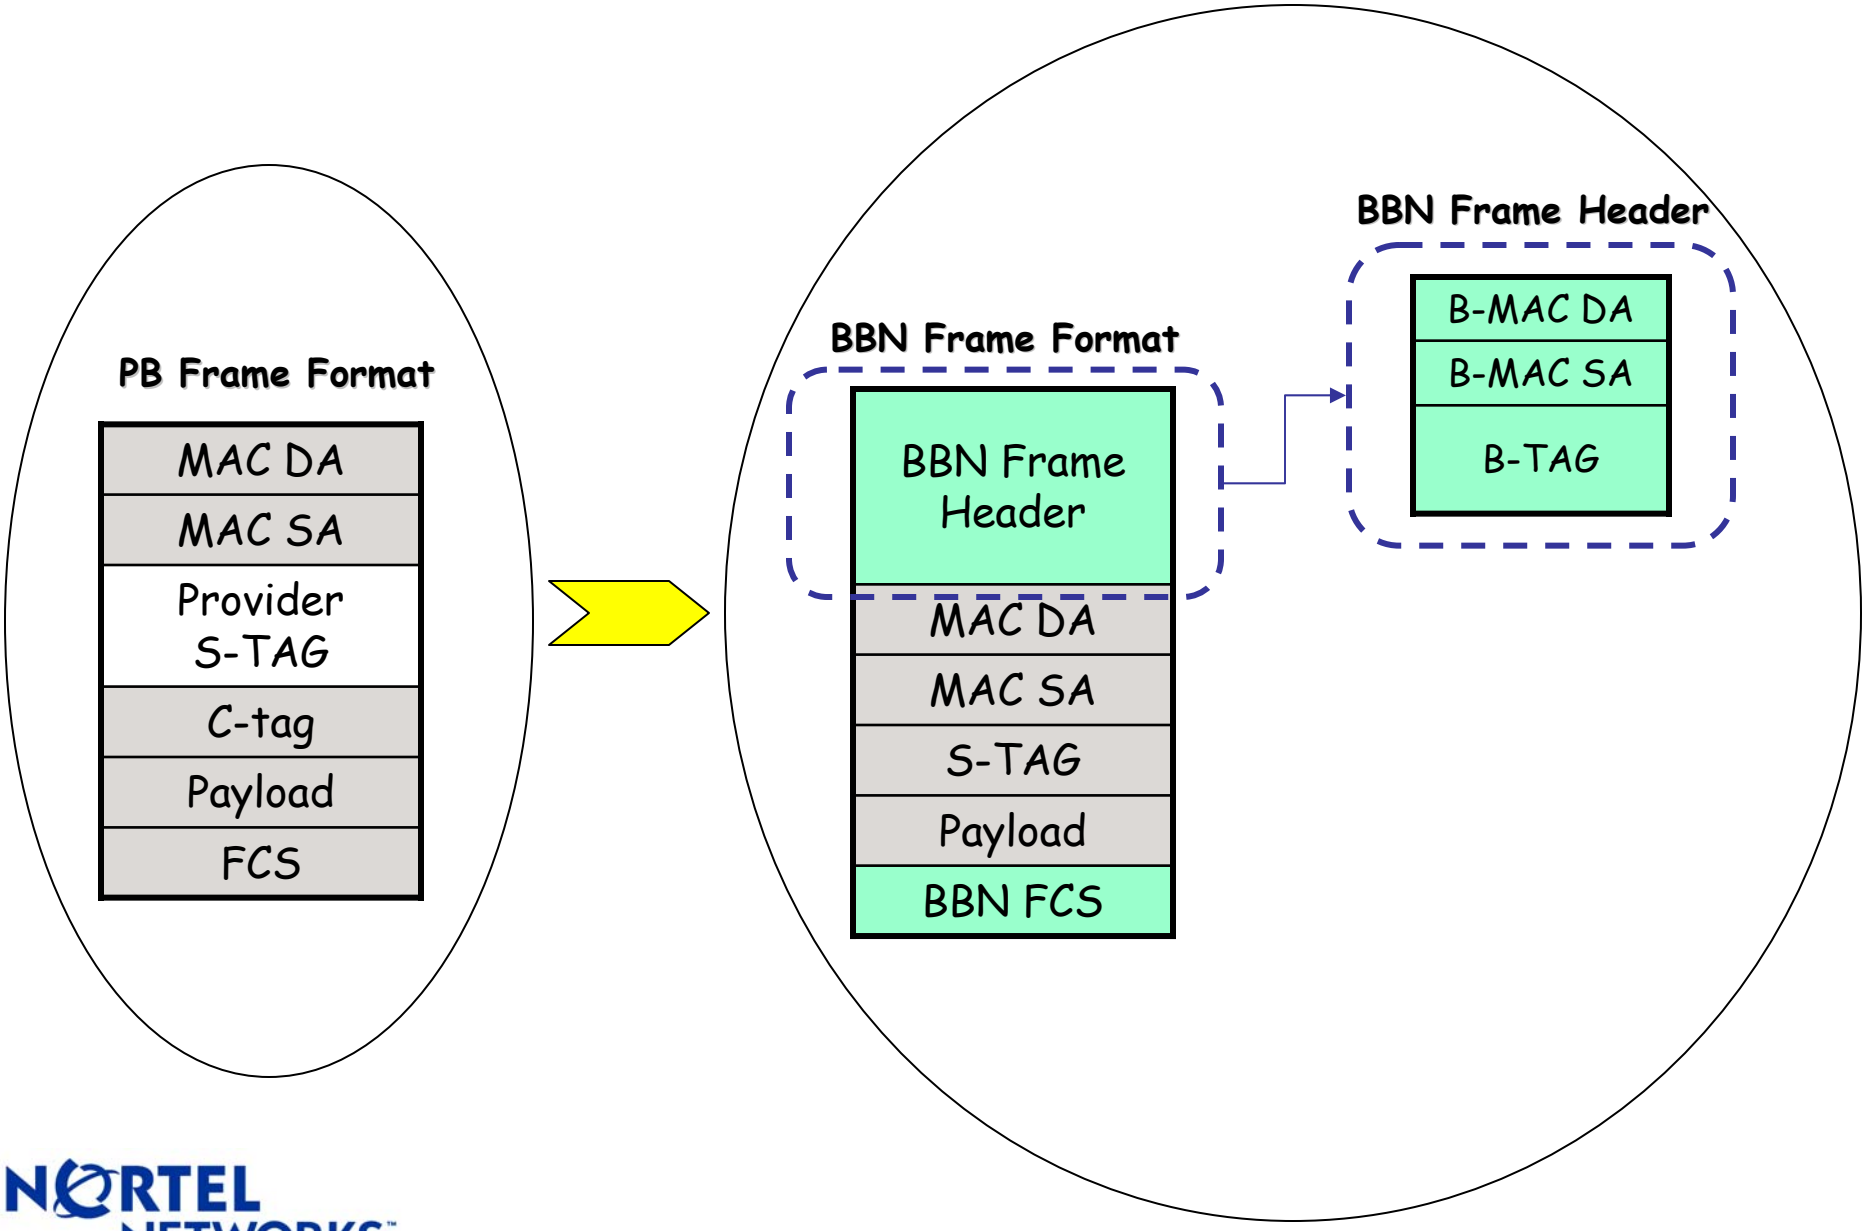

Basic Backbone Forwarding

Basic Back Backbone Operation

- Frames from PBN are encapsuled with BBN header
  - S-VID maps from/to provisioned ES-VID
  - DA B-MAC from provisioned ES-VID table
- Forwarding through BBN may occur as in 802.1D Bridge
- At destination BBN header is de-encapsuled

But E-LAN/E-TREE Are Flooded

- These are filtered at BB POP based on ES-VID lookup, but still
- Add to the basic picture the B-VID tunnel
  - This tunnel may have many service instances multiplexed since the ES-VID separates the service instances
- Change mapping for E-LAN and E-LINE to use a B-VID (for a B-VLAN) with the correct backbone connectivity
- Solution allows E-LINE scaling to Millions, but E-LAN and E-TREE are limited to 4K instances
- Not bad, but we could go further

Multicast Relay Scaling

Backbone Provider Bridging

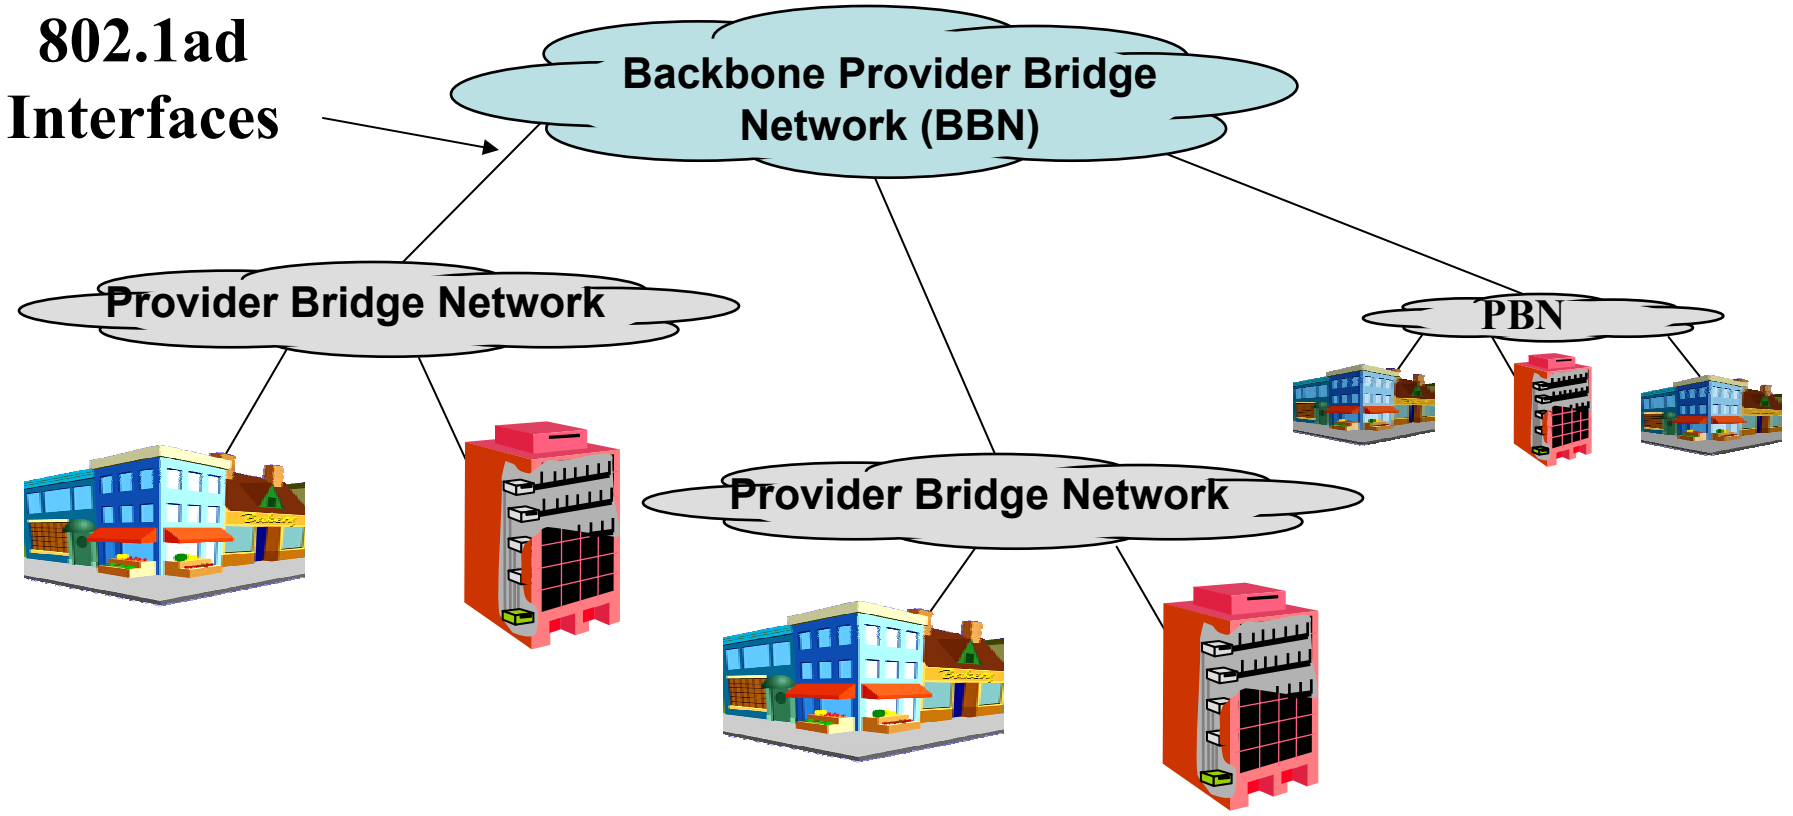

Terminology

- 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

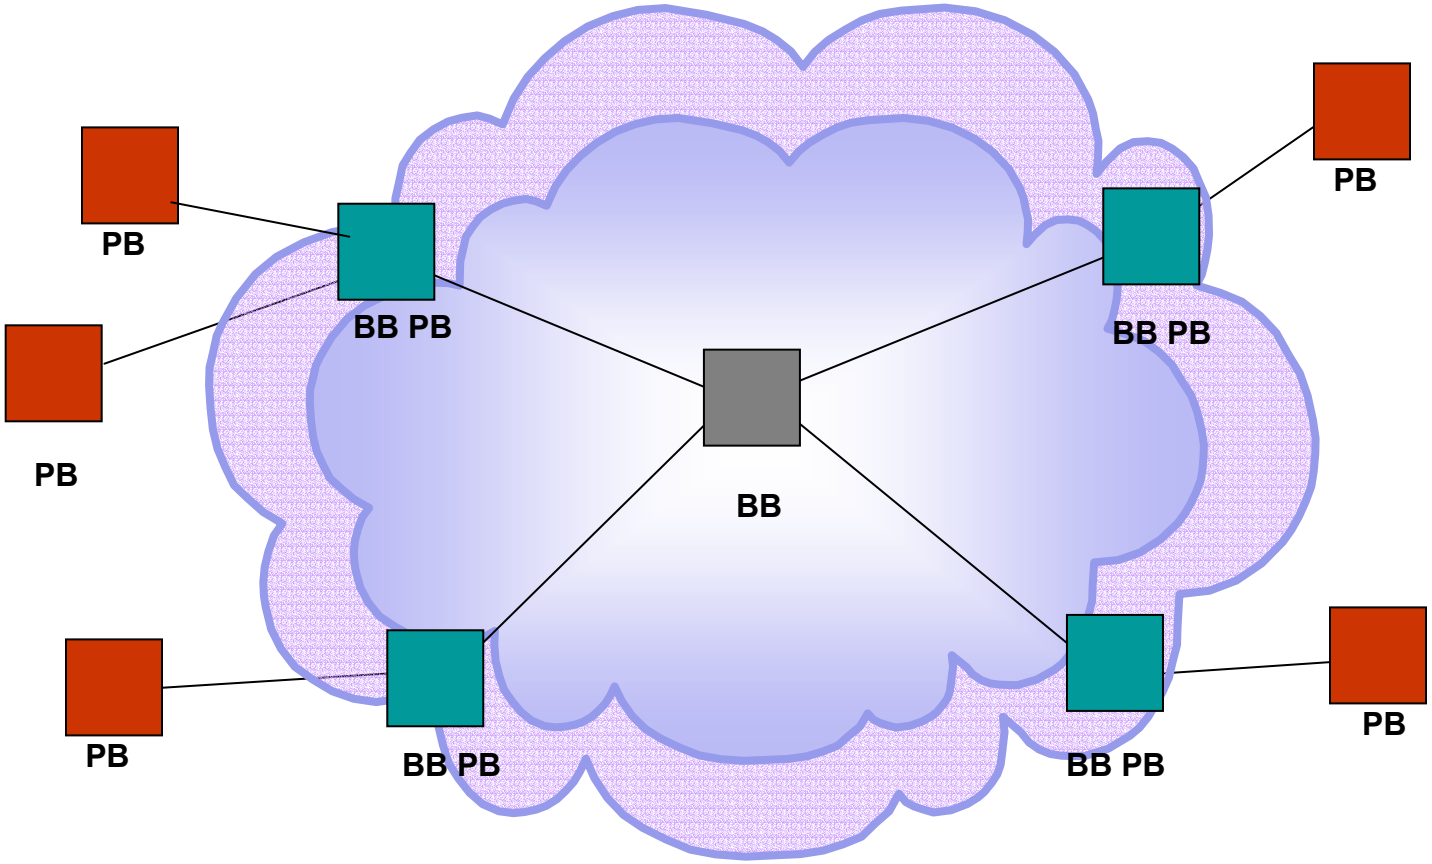

- PB: Provider Bridge (as defined by 802.1ad)
- BB PB: Backbone Provider Bridge Edge
- BB: Backbone Provider Bridge

NORTEL NETWORKS

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

NORTEL NETWORKS

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

NORTEL NETWORKS

A Two Layer Hierarchical BBN

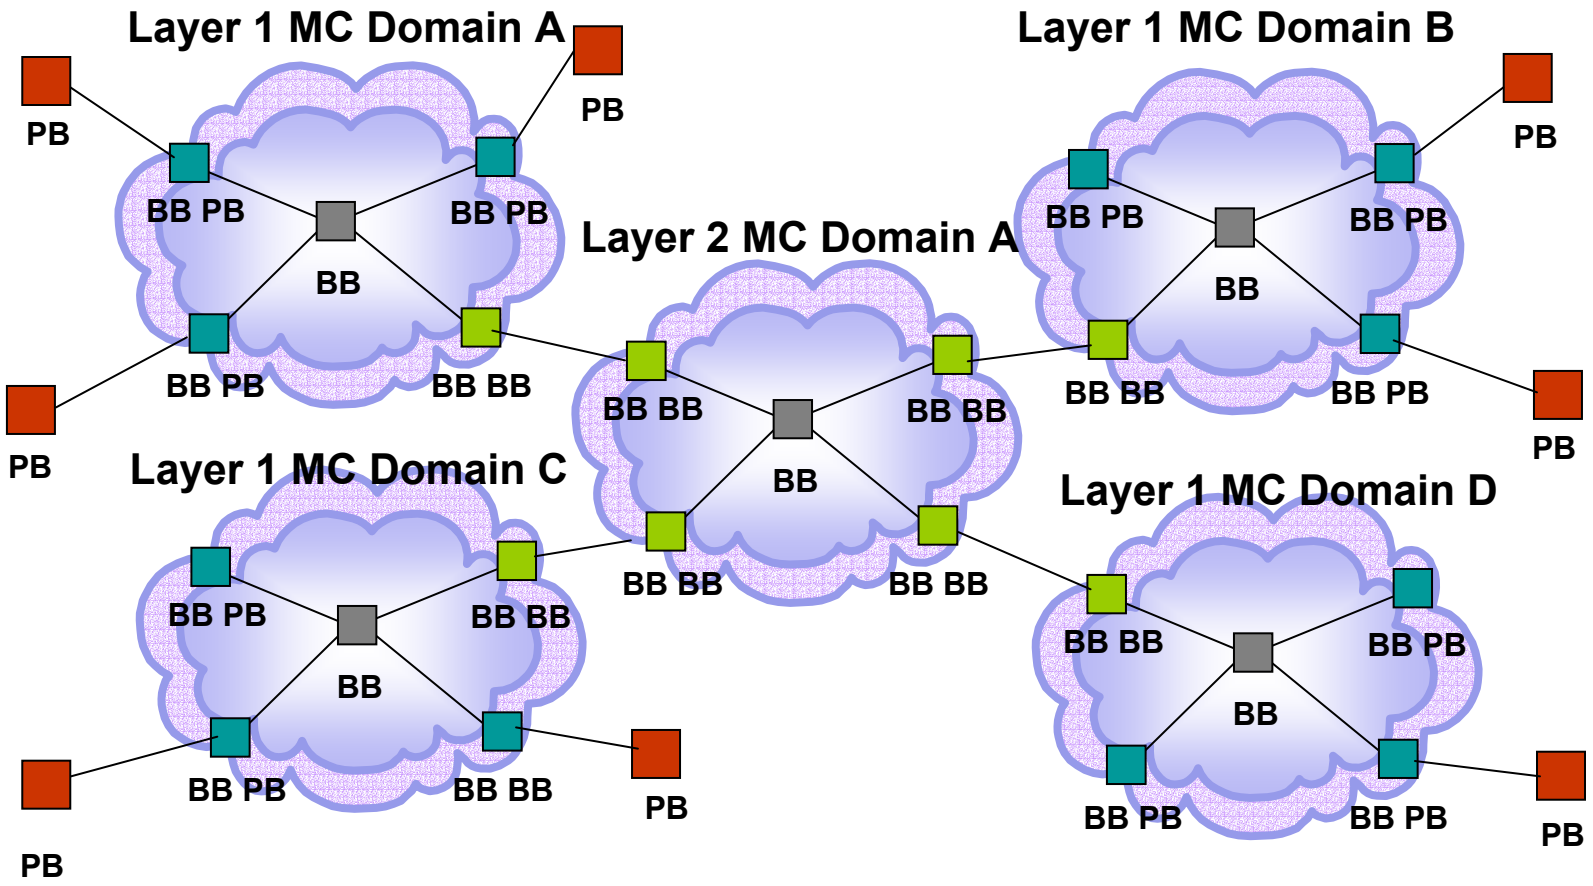

- **PB**: Provider Bridge (as defined by 802.1ad)
- **BB PB**: Backbone Provider Bridge Edge
- **BB**: Backbone Provider Bridge
- **BB BB**: Backbone Provider Bridge Layer Edge

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

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

NORTEL NETWORKS

PAR Title

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

PAR Scope

- To develop bridge protocols and architecture, 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.

NORTEL NETWORKS™

Objectives

- interconnect Provider Bridge Networks in a manner that allow 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
- provide rapid distribution of S-VLAN tags
- provide rapid healing of network failures without interruption of service to unaffected S-VLANs