THE POSSIBILITIES ARE INFINITE FUJITSU

Understanding PBB-TE for Carrier Ethernet

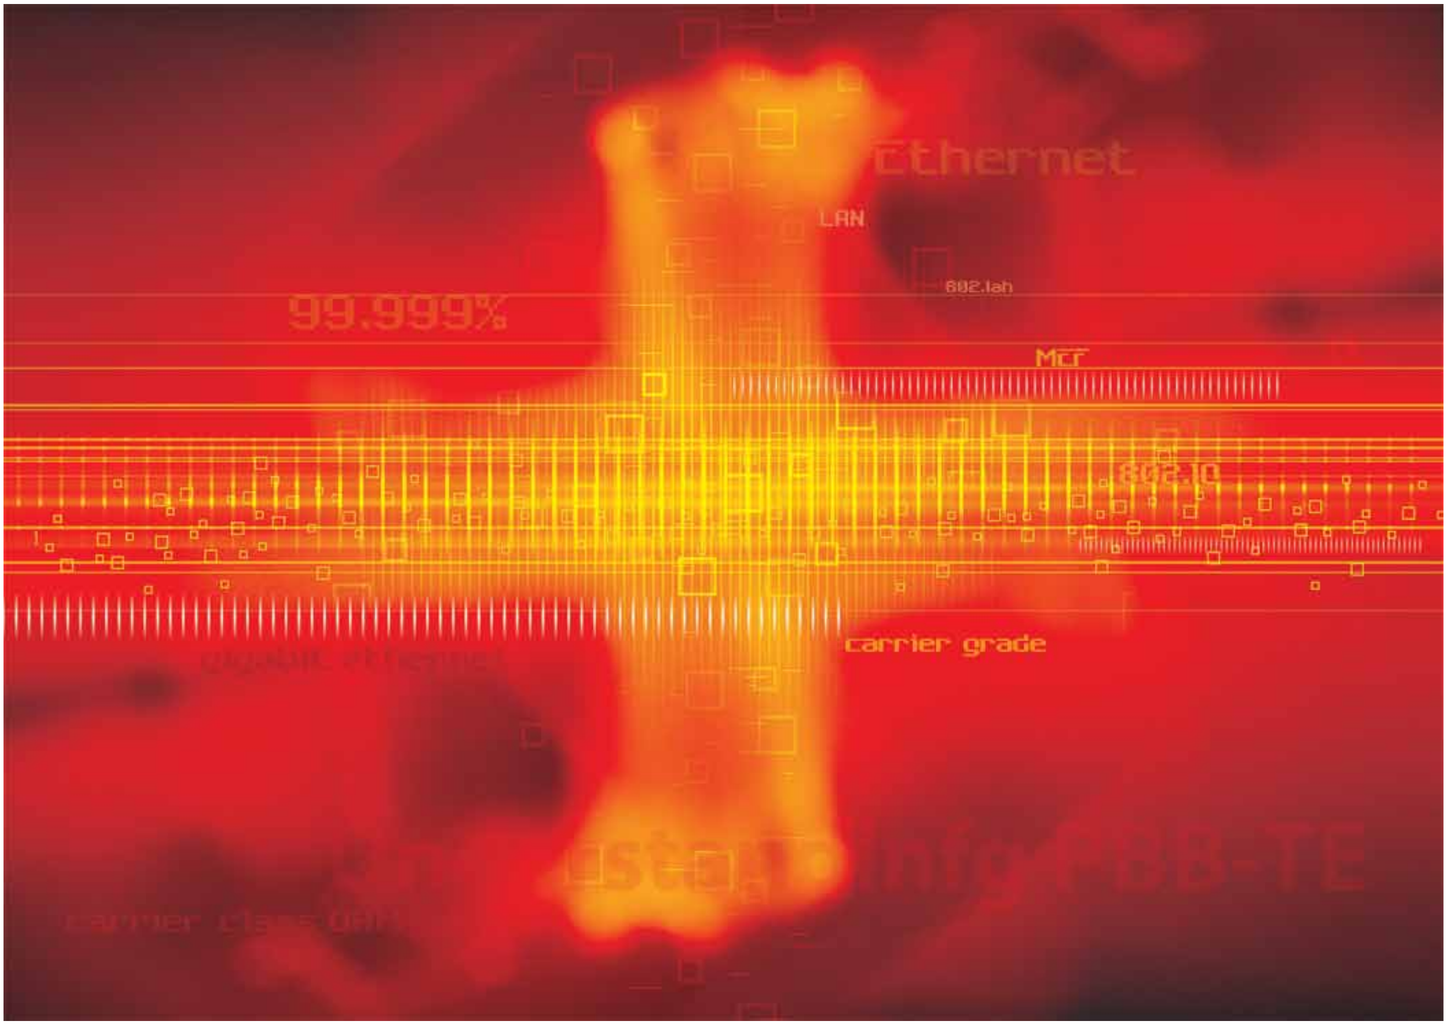

Introduction

Ethernet is evolving from an enterprise LAN technology to a much more robust, carrier-grade transport technology for metropolitan service networks. Ethernet is the dominant LAN technology in most offices, with ubiquitous 10/100 Ethernet and GigE interfaces used to interconnect high-speed corporate networks. Ethernet's popularity as a LAN technology is driven by the fact that it is simple and very inexpensive. However, some of the attributes that make Ethernet simple and inexpensive for office LANs impede its use for carrier-grade data services over larger metropolitan networks. Because of these limitations, mission critical corporate data networks have traditionally relied on Frame Relay, ATM or point-to-point private line (DS1, OC-n) connection-oriented services due to their high reliability and guaranteed service delivery.

In the last ten years, the amount of corporate, home, and internet data traffic has increased exponentially. Carriers are looking for new, less expensive ways to transport data services over metropolitan areas, while ensuring guaranteed end-to-end connections and SLAs. To fulfill this need for carrier-grade data connectivity over much larger areas, Ethernet has recently incorporated several enhancements. The latest, and most important, of these changes is a new Ethernet feature called Traffic Engineered Provider Backbone Bridges, or PBB-TE.

Issues with Legacy Ethernet

Ethernet was initially conceived as an enterprise LAN technology for interconnecting computers within a single office. Initial deployments within carrier networks simply relied on existing enterprise-quality legacy Ethernet switches, as shown in Figure 1.

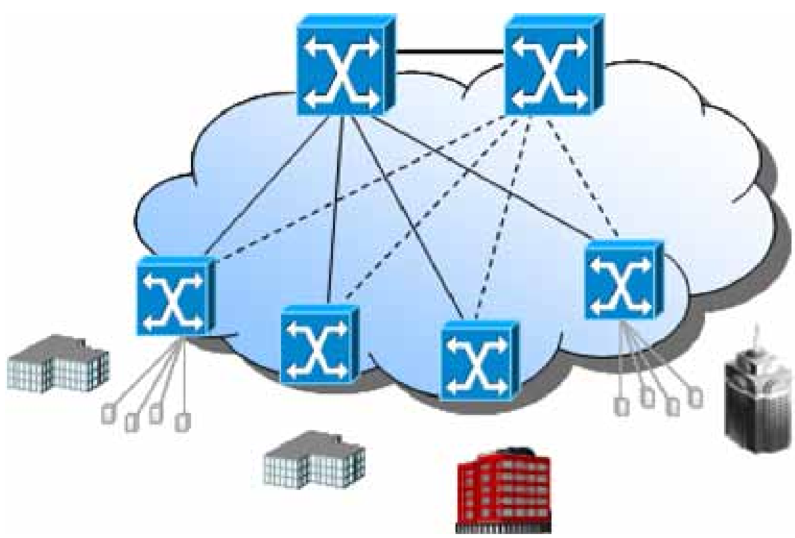

Figure 1: Legacy Switched Ethernet Metro Network

From a service and performance perspective, there are several issues with a network based on legacy enterprise-type Ethernet switches, including:

- Poor scalability
- Slow protection switching and restoration (2s to 120s)
- Limited QoS—hop-by-hop priority queuing
- Limited traffic engineering
- OAM worse than traditional transport networks

Even with these problems, switched Ethernet networks utilizing legacy enterprise LAN switches were deployed to support initial carrier Ethernet services. These initial switched Ethernet networks offered fine granularity data services (~1 Mbps), some level of protection switching and network restoration and support of both point-to-point and multipoint Ethernet data services.

Fundamental Ethernet Problems

While Ethernet is a simple and low-cost LAN technology, it suffers from three fundamental problems: scalability, flooding, and slow restoration times.

Ethernet is based on a flat MAC addressing scheme. In a flat MAC addressing scheme, Ethernet switches learn and store the MAC address of every device in the network. They learn the MAC addresses of all devices connected directly to their own switch ports, as well as all the devices connected to any other switch in their network (i.e., broadcast domain). The MAC addresses are stored in a MAC filter table and used to switch Ethernet frames from the input port to the correct output port. The MAC filter table is the heart of any Ethernet switch. As the size of a network grows, the MAC filter tables can grow extremely large, eventually limited by the amount of memory and processing power in the Ethernet switch. This isn't usually a problem for a few dozen or hundreds of devices in an office LAN, but it can quickly become a limiting factor in metro networks consisting of thousands to hundreds of thousands of devices.

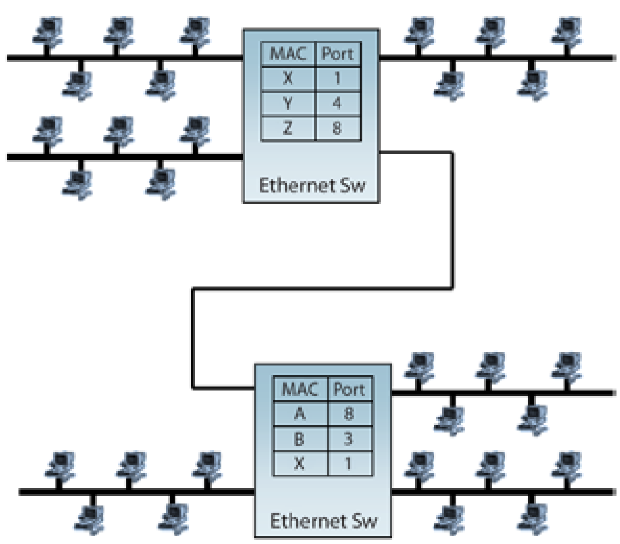

Figure 2: Flat MAC Address

Flooding is the second major issue with legacy Ethernet. When an Ethernet packet arrives at a switch which doesn't have a corresponding entry in its MAC filter table, it simply broadcasts, or "floods," the entire network with the unknown packet. Flooding is how an Ethernet switch learns which devices are on the network and where they are located—it's how they populate and update their MAC filter tables. As the size of a network increases, the constant flooding of unknown packets dramatically degrades overall network performance.

The only thing worse than flooding is any form of Spanning Tree Protocol (STP, RSTP, MSTP). STP is the method Ethernet uses to identify redundant paths in a network and use those redundant paths for protection switching when a link fails. The problem with STP is that it is very slow, typically requiring 10s–100s for full restoration. Rapid STP improves protection switch times to the 1s–10s range, but this is an order of magnitude slower than the carrier-grade standard of <50 msec.

Ethernet Evolving

To evolve Ethernet to a true carrier-grade service and transport technology, several enhancements have been added to the standard over the last couple of years. In addition, the MEF has been active in defining five attributes of Carrier Ethernet services.

- Standardized Services
  - Defined set of Ethernet services (E-Line, E-LAN)
  - Seamless integration of TDM services
  - Interoperable hand-offs between carriers
- Scalability
  - Ability to support thousands to millions of Ethernet services
- Reliability/Protection
  - 99.999% availability
  - <50 msec protection switching
- Service Management
  - Carrier-class OAM
- Quality of Service
  - Guaranteed end-to-end SLAs

PBB-TE achieves these five objectives for carrier-grade Ethernet services and represents the latest enhancement to the Ethernet standard. PBB-TE is widely expected to become the dominant technology for guaranteed, connection-oriented data transport in the aggregation and transport portion of metropolitan networks. PBB-TE is based upon the 802.1Q VLAN standard, which provides a starting point for understanding the incremental enhancements that led up to PBB-TE.

802.1Q Virtual LANs – The First Step

Virtual LANs, also known as 802.1Q, provides a method for assigning a "Virtual Identifier" or VID tag to each Ethernet frame. Within an office, VIDs allowed different organizational groups to be connected together in virtual LANs, improving overall network efficiency and performance (i.e., engineering VID=100, marketing VID=200, sales VID=300). For a carrier, the VID became a useful tag for identifying services from a specific customer, as shown in Figure 3.

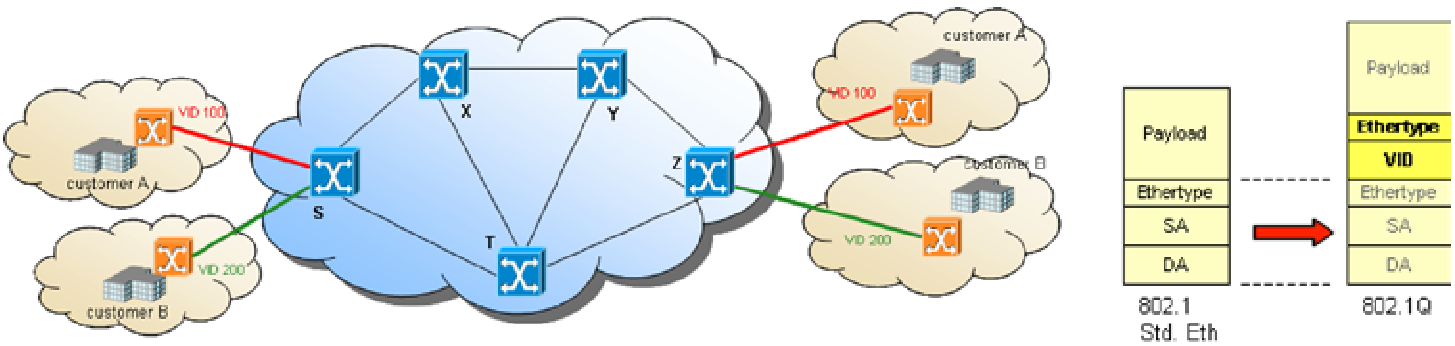

Fig 3: Virtual LANs

The VLAN ID works well for enterprise applications, but there are several problems when using 802.1Q in a carrier application. The first issue is that the VIDs had to be globally unique across an entire network; no two Ethernet switches could use the same VIDs. Ensuring globally unique VIDs became an administrative nightmare for carriers, since they were forced to coordinate VID assignments across all customer switches.

The second major issue is the lack of network scalability. VLAN IDs are limited to 12 bits in length, which only supports 4,094 service instances. While this number of VLAN IDs is fine for internal enterprise applications, it clearly doesn't scale to the tens of thousands or hundreds of thousands of Ethernet services required in a major metropolitan carrier network.

802.1ad Provider Bridges—Solving the Administrative Issue

In order to solve the administrative nightmare of coordinating VLAN IDs across both carrier and customer networks, a new Ethernet enhancement was added that provided separation of C-VIDs and S-VIDs. This new standard is called Ethernet Provider Bridging (802.1ad), but is also commonly referred to as "Stacked VLANs" or "Q in Q." With Provider Bridging, service providers have their own set of S-VIDs that are completely separate and independent of customer provisioned C-VIDs.

The Ethernet frame was modified slightly so that C-VIDs are carried transparently across the network. In effect, C-VIDs simply become part of the payload. Figure 4 shows the evolution of the Ethernet frame structure and why 802.1ad is commonly referred to as "stacked VLANs."

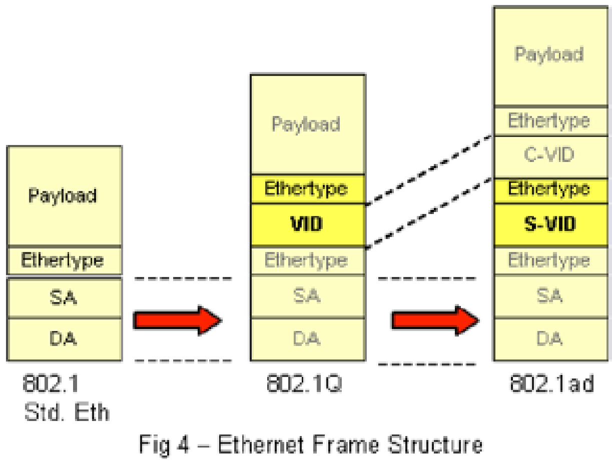

Fig 4: Ethernet Frame Structure

Figure 5 shows a typical network implementation with the S-VID assigned as the Ethernet frame enters the network and removed by the last switch, or egress point, in the carrier network.

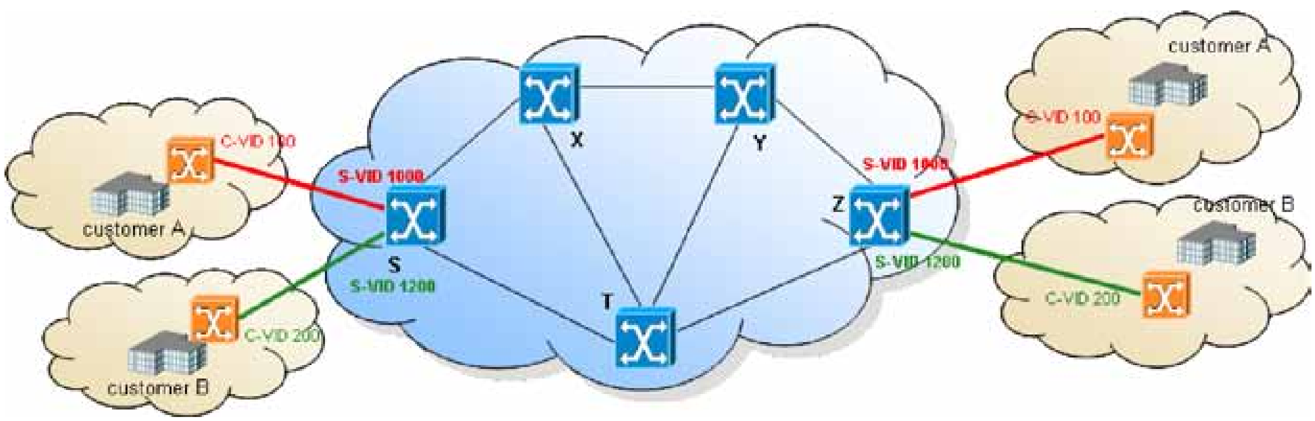

Fig 5: Provider Bridging

Control Plane Problem

While 802.1ad Provider Bridging solved a number of issues, specifically the separation and administration of VLAN IDs, a number of problems still existed between customer and carrier Ethernet networks. Ethernet relies on a single, flat MAC address space, which operates across both customer and service provider networks. With a single flat MAC addressing scheme, MAC filter table updates, flooding and spanning tree protocols operate across the customer-carrier boundary. In a nutshell, there is no separation of the customer network from the service provider network—it simply looks and acts like one big Ethernet switch domain, as shown in Figure 6. As mentioned previously, flat MAC addressing leads to large MAC filter tables and scalability problems. The network is exposed to flooding of unknown packets across the network (customer & carrier). In addition, the network still relies on relatively slow STP/RSTP/MSTP for restoration. To provide carrier-grade Ethernet services and effectively manage their network, service providers need clear separation of their Ethernet infrastructure from their customers' Ethernet domains.

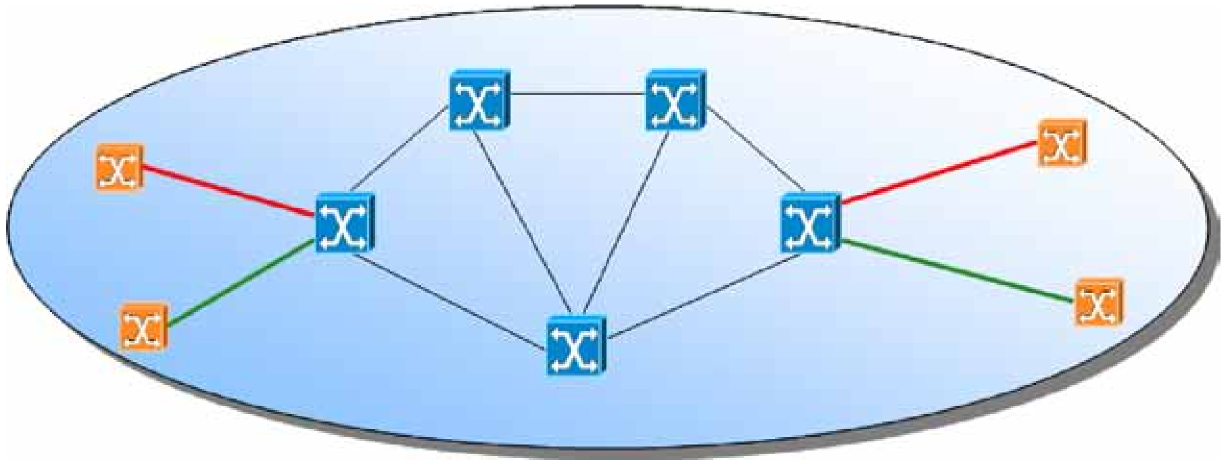

Fig 6: Control Plane Perspective—One Big Network

802.1ah Provider Backbone Bridges—Control Plane Separation

The lack of separation of customer networks from the carrier network (control plane), lack of scalability (limit of 4094 VIDs), and lack of end-to-end QoS need to be resolved to achieve connection-oriented, carrier-grade Ethernet. Solving these remaining problems resulted in "tweaking" the Ethernet structure yet again, with the new standard being known as "Provider Backbone Bridges" or 802.1ah.

The previous section describes how stacking VLAN IDs (802.1ad) allowed separation of the C-VIDs from the S-VIDs. Since the control plane operates at the MAC layer and the goal is to provide separation of the customer control plane from the service provider control plane, one suggestion was to simply "stack" MAC addresses in a similar manner. This is exactly what is defined in the 802.1ah Provider Backbone Bridging standard, as shown in Figure 7, which is also referred to as MAC in MAC.

The service provider Ethernet switches (also called Backbone Bridges) use their own backbone MAC addresses (B-SA, B-DA) for implementing and maintaining their own control plane (MAC filter tables, flooding, protection). Because the carrier network operates on its own set of backbone MAC addresses, the carrier control plane does not interact with customer Ethernet switches. At the edge of the service provider network, customer MAC addresses (SA, DA) are moved into the Ethernet payload and transported transparently across the network, similar to what was done with customer VLAN IDs.

FUJITSU NETWORK COMMUNICATIONS INC.
2801 Telecom Parkway, Richardson, Texas 75082-3515
Telephone: (972) 690-6000
(800) 777-FAST (U.S.)
us.fujitsu.com/telecom

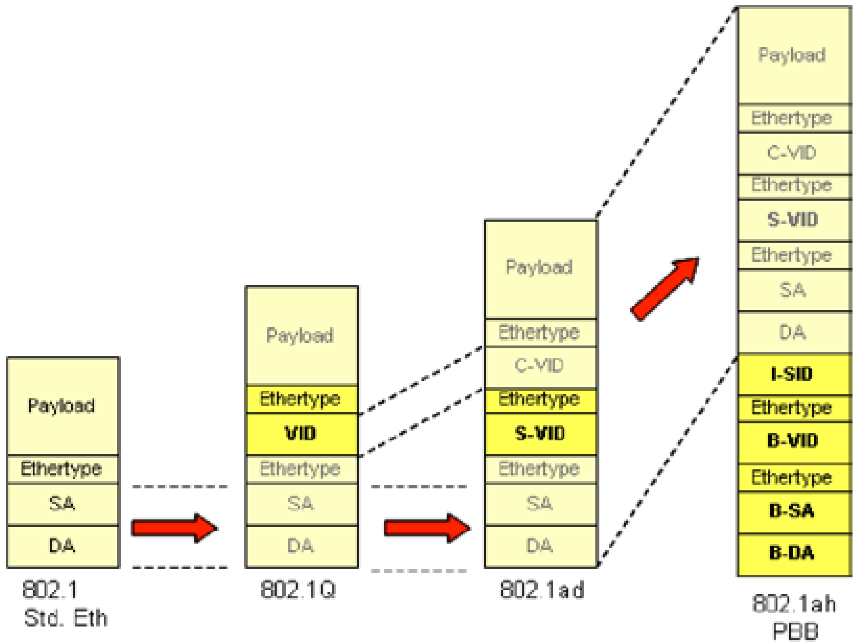

Fig 7: 802.1ad "MAC in MAC"

In addition to backbone MAC addresses, the 802.1ah standard incorporates a new service identifier to solve the VLAN ID scalability issue. Scalability has always been a problem due to the fact that VLAN IDs are limited to 12 bits, which only provides 4094 service instances. The 802.1ah standard defines a new tag called the I-SID, which can be used in place of S-VIDs. The I-SID is 24 bits long, supporting over 16 million services within a single metropolitan or regional area network.

Finally, the Provider Backbone Bridging standard incorporates a new feature called Backbone VLANs. In most metropolitan networks, data traffic flows to a few major endpoints, such as a data center, ISP gateway, or long-haul POP. Backbone VLANs allow a carrier to define a few "super highways" between major network ingress and egress points. These "super highways" are identified by a B-VID. As new customers and services are added to the network, they can simply be assigned to existing B-VID tunnels. Using B-VIDs eliminates the need to manually provision each and every service flow across intermediary nodes on the network.

Backbone VLANs are end-to-end provisioned "pipes" across the network. Provider Backbone Bridging has the added benefit that both primary and backup VLANs can be defined across the network. Hmm—these things are starting to look a lot like STS circuits across a SONET network and that's a good thing! We're starting to make Ethernet perform like a guaranteed, connection-oriented network with end-to-end QOS and to allow provisioning in a similar manner to existing methods and procedures.

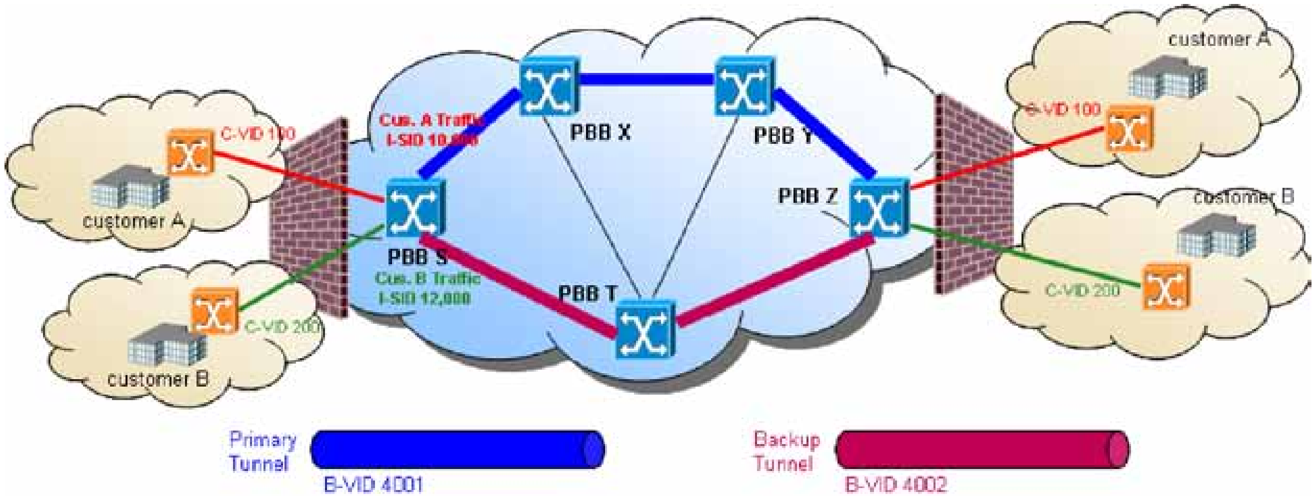

Fig 8: 802.1ah Provider Backbone Bridging

Carrier-Grade Ethernet

Provider Backbone Bridging (802.1ah) resolved many of the issues with legacy Ethernet LANs. The scalability problems were resolved by the adoption of the I-SID, which allows over 16 million individual uniquely identifiable services. The new "MAC in MAC" frame structure allows clear separation of customer and carrier network control planes (MAC filter tables, Flooding, RSTP). Finally, PBB VLAN IDs provide a means to establish end-to-end paths across a network.

Despite all of the progress a few problem areas remained, specifically within the service provider network. Within the carrier cloud shown in Figure 9, many of the old, problematic Ethernet mechanisms are still utilized including:

- MAC filtering tables & updates
- Flooding of unknown packets
- Long protection switching/restoration times (STP/RSTP)

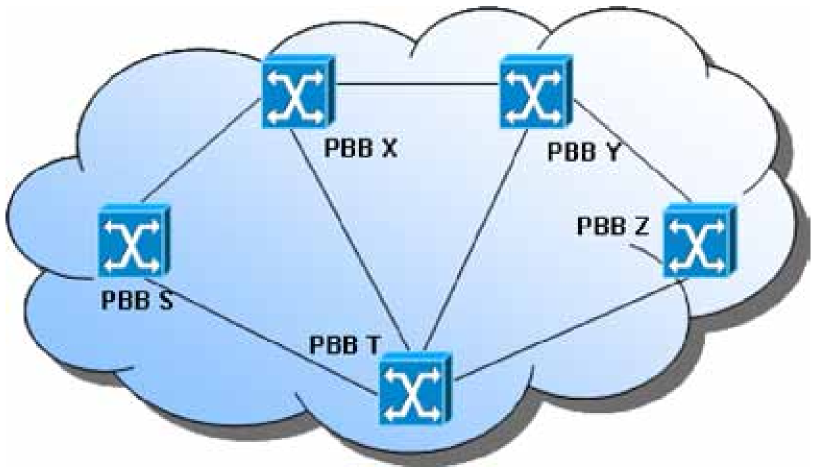

Fig 9: Service Provider Ethernet Network

802.1Qay Provider Backbone Bridges Traffic Engineered (PBB-TE)

Provider Backbone Bridging solved most of the carrier Ethernet problems, but left a few undesirable issues in the core. If MAC filter table updates, unknown packet flooding, and long protection times due to RSTP are so problematic, why not simply turn these mechanisms off? This is exactly what is done with the new PBB-TE (802.1Qay) standard. PBB-TE is that simple, just a matter of turning off these problematic control plane functions.

By turning off these legacy Ethernet functions, there needs to be an alternative method to provision the MAC Filter tables, backbone tunnels, and Ethernet service parameters. PBB-TE simply states that these functions can be provisioned by an external management system. For Fujitsu networks, this would be our NETSMART® 1500 network management platform.

Through these incremental Ethernet changes, VLANs (802.1Q), Provider Bridges (802.1ad) Provider Backbone Bridging (802.1ah), and PBB-TE (802.1Qay), we've taken Ethernet from a relatively low-quality, best-effort network to a connection-oriented, carrier-grade network that behaves similarly to proven, rock-solid, SONET networks. Even the provisioning is similar to current methods, such as A-Z path provisioning feature of the NETSMART 1500 Network Management System.

PBB-TE achieves the objective of providing carrier-grade, connection-oriented Ethernet across metro and regional areas. The use of an external management control aligns very closely with existing carrier methods and procedures for provisioning legacy services, as well as aligning their organizational and staff structures.

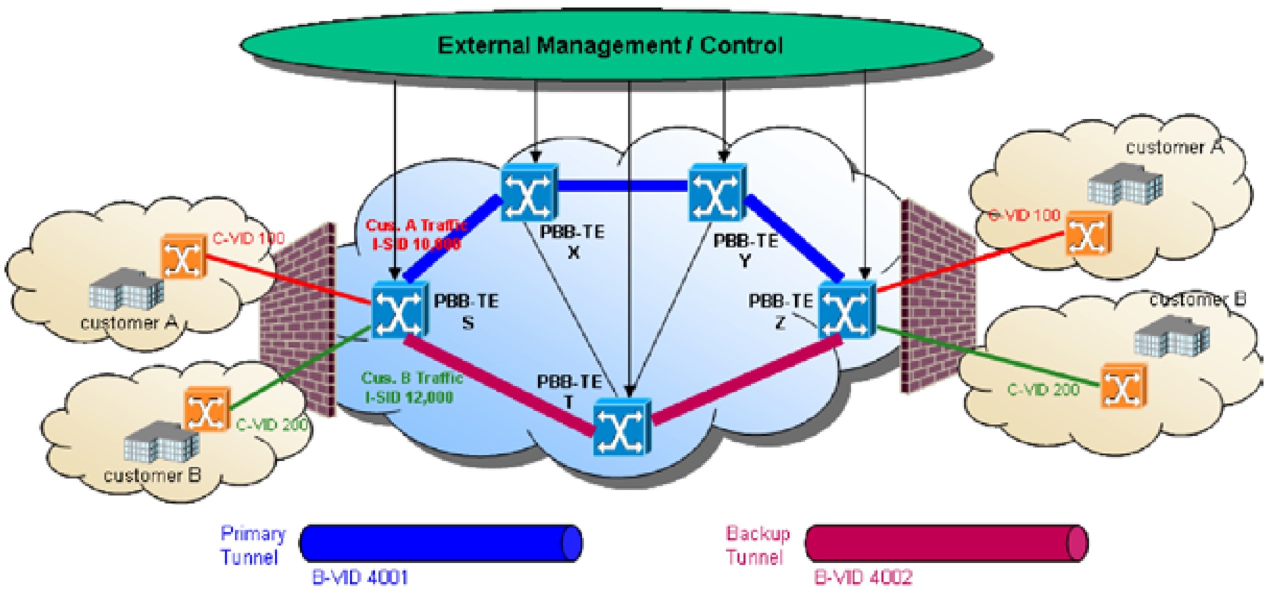

Fig 10: PBB-TE

PBB-TE vs IP/MPLS

There has been a great deal of discussion in the industry over which Ethernet tunneling technology, PBB-TE or IP/MPLS, is better suited to providing carrier-grade, connection-oriented data services. Both of these technologies have a fit and function in carrier networks, but in different areas. By deploying both technologies, a carrier optimizes their network efficiency, minimizes their network costs, and aligns their network to their internal staffing, training, and organizational structure.

In most metro networks, there is a natural division between the core routing/service elements, located at a handful of major COs, and the aggregation and transport parts of the network which operate over hundreds to thousands of nodes. Due to the large number of nodes deployed in the aggregation and transport part of the network, these elements need to be relatively low-cost and easily managed and maintained by large telco craft organizations. In contrast, there are only a handful of core routing/service elements deployed in a metro area. These nodes tend to be much more expensive and require a much higher level of technical expertise for provisioning, operations and maintenance support than the transport nodes.

This segmentation of the network into aggregation and core improves overall system scalability, transport efficiency, and service element utilization, while reducing provisioning and management complexity.

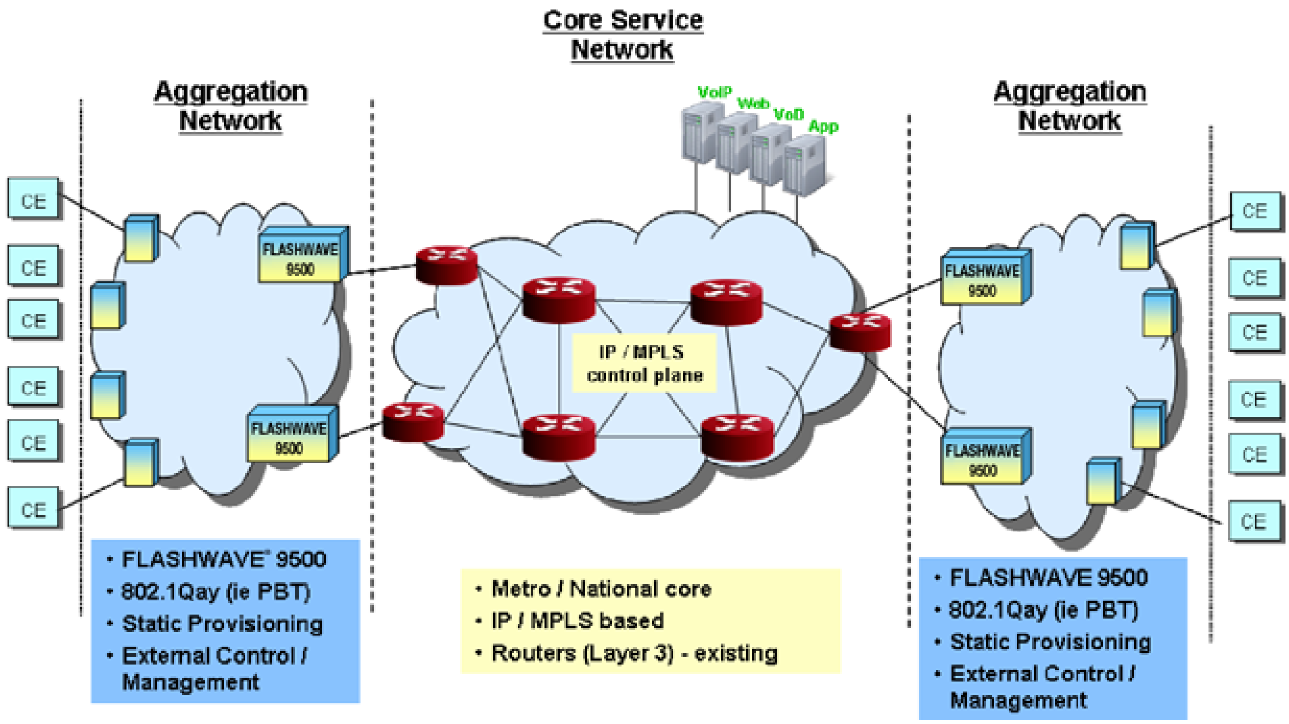

Fig 11: Aggregation/Core Network Separation

This separation of the network into aggregation/transport and core/services also provides a natural alignment for the Ethernet tunneling technology (PBB-TE and IP/MPLS). For the aggregation/transport part of the network, PBB-TE has many advantages and benefits over IP/MPLS. Likewise, IP/MPLS is the default choice for most Layer 3 routers in the core/services portion of the network.

PBB-TE	IP/MPLS
Simple aggregation network • Static provisioning of tunnels (B-VIDS) – Similar to current transport provisioning • Simple training requirements for staff • Simple protection scheme	**Complex core network** • Requires multiple L3 protocols – ISIS-TE, OSPF-TE, RSVP-TE, BGPv4, LDP, BFD, VCCV • Complex training requirements – Detailed Layer 3 routing expertise required • Complex protection algorithm (FRR)
Separation of responsibilities – Aligns with carrier methods & organizational structure	**Integration across network layers** • Substantial coordination required between access/transport groups and Services groups
Improved network scaling	**Scaling Issues** • Scaling routed networks ?? • Convergence times ?

Conclusion

Ethernet has become the dominant and ubiquitous enterprise LAN technology, due to its low cost, simplicity and high speed. As the need for data networking services has grown, there has been interest in using Ethernet as a transport technology for metropolitan and regional area networks, essentially replacing or supplementing existing EoS, Frame Relay, ATM and private line connections. The Ethernet that works so well in the office is far removed from a carrier-grade, connection-oriented technology that service providers need to support guaranteed data services. The primary issues with legacy Ethernet are the flat MAC addressing scheme, flooding of unknown packets and poor restoration times. Several enhancements have been amended to Ethernet over the last few years to improve its ability to function as a carrier transport technology, including VLANs (802.1Q), Provider Bridging (802.1ad), Provider Backbone Bridging (802.1ah). PBB-TE (802.1Qay) is the latest of these enhancements and provides full support for the deployment of carrier grade Ethernet services. PBB-TE provides lower operating costs, lower equipment costs, lower training cost and lower implementation cost than alternative options, such as IP/MPLS. Fujitsu has adopted PBB-TE as the connection-oriented Ethernet technology for our next generation Packet ONP.

Acronym	Descriptor
ATM	Asynchronous Transfer Mode
B-DA	Backbone Destination MAC Address
B-SA	Backbone Source MAC Address
B-VID	Backbone VID
CO	Central Office
C-VID	Customer VID
DA	Destination MAC Address
EoS	Ethernet over SONET
GigE	Gigabit Ethernet
I-SID	Individual Service Identifier
LAN	Local Area Network
MAC	Media Address Control
MEF	Metro Ethernet Forum
MPLS	Multi-protocol Label Switching
MSTP	Multiple Spanning Tree Protocol
OAM	Operations, Administration and Maintenance
PBB	Provider Backbone Bridge
PBB-TE	Provider Backbone Bridges with Traffic Engineering
Packet ONP	Packet Optical Networking Platform
QOS	Quality of Service
RSTP	Rapid Spanning Tree Protocol
SA	Source MAC Address
SLA	Service Level Agreement
SONET	Synchronous Optical Network
STP	Spanning Tree Protocol
S-VID	Service Provider VID
TDM	Time Division Multiplexing
VID	Virtual Identifier
VLAN	Virtual Local Area Network

FUJITSU NETWORK COMMUNICATIONS INC.
2801 Telecom Parkway, Richardson, Texas 75082-3515
Telephone: (972) 690-6000
(800) 777-FAST (U.S.)
us.fujitsu.com/telecom