Scaling Provider Ethernet

Paul Bottorff, Nortel
Panagiotis Saltsidis, Ericsson

ABSTRACT

We provide an overview of the IEEE provider backbone bridging standard and its associated traffic engineering standard. Furthermore, we relate these to the supporting work of the IEEE 802.1 committee.

INTRODUCTION

Ethernet and Ethernet bridging are rapidly becoming the technologies of choice for provider access and aggregation networks [1], and are viable as the provider core network of the future. This achievement stems from the intrinsic characteristics which are driven by the maturity of the technology and a high level of standardization: fast, simple, inexpensive, easy to manage, and backward compatible. In particular, the IEEE 802.1ad (provider bridges or PB) [2], the recently completed IEEE 802.1ah (provider backbone bridges or PBB) [3], and the IEEE 802.1ag (connectivity fault management or CFM) [4] standards have addressed the issues of creating provider services using Ethernet infrastructure, scaling Ethernet networks, and managing faults of Ethernet services. These Ethernet bridging technologies have provided the tools for building large provider Ethernet bridged networks while offering provider-customer independence and fault management. These standards, however, have not addressed the problem of providing traffic engineering capabilities and related rapid protection against failure. This is being addressed by the IEEE 802.1Qay (PBB traffic engineering or PBB-TE) [5] standard project, which is described in this article.

Provider bridge is the first bridging technology standardized by IEEE for provider applications. It solves two problems presented by the application of enterprise 802.1Q [6] bridges to provider networks. The first problem is separating the customer virtual LAN (VLAN) space from the provider VLAN space, and the second problem is allowing the provider the ability to multiplex many customer VLANs on a single provider VLAN. To address these problems a provider bridge supports the creation of service VLANs (S-VLANs) within an independent VLAN space controlled by the provider and used to carry customer VLANs (C-VLANs) [7, 8].

Provider backbone bridges were developed to address three problems with scaling of provider bridged networks (PBNs) [7]. The first problem is that a single PBN is only capable of supporting 4094 provider services. This limits the utility of isolated PBNs to small-scale deployments and access network applications, since they cannot provide enough services to support an entire metro or wide area backbone. PBB solved this problem by creating a provider backbone bridged network (PBBN) capable of connecting many PBNs and allowing identification of 16,776,959 provider service instances. The second problem is that PBNs do not separate the provider and customer address spaces. This problem limits the scaling of the PBN because the provider network must learn all the customer addresses that are associated with multipoint services in the provider's core bridges, resulting in the core state growing in proportion to the size of the customer networks rather than the size of the provider network. PBB solved this problem by encapsulating all service frames in a new frame that uses addresses internal to the PBBN. In this way only the edge bridges of the PBBN need to learn customer addresses. Many customer addresses are then carried using a much smaller number of backbone addresses. The third problem not addressed by PB is the separation of the service from the transport identifier. In the PB case the S-VLAN identifies both the service and the path used within the provider network to transport the service. The lack of separation between service and transport identifiers limits the scaling of the network, since it forces the network core to retain state for each service rather than limiting the core state by aggregating many services into a single transport channel. PBB solved this problem by having separate tags for service identification and VLAN identification. The service tag is only used at the edges of the PBBN to identify the service, while the core of the PBBN only uses a backbone VLAN to identify the path through the network. This allows many services to be multiplexed on a single backbone VLAN.

Connectivity fault management (CFM, IEEE 802.1ag) [4] was developed to address the issue of detecting, verifying, and isolating faults in bridged Ethernet networks. CFM also supports scaling of a network by allowing fault detection and isolation within hierarchies of multiple autonomies and technologies. The interconnection of these networks may result in faults that are not visible within the individual networks, but can be detected by running CFM over the entire network.

0163-6804/08/$25.00 © 2008 IEEE

Provider backbone bridge traffic engineering (PBB-TE, IEEE 802.1Qay) [5] further enhances the PBN and PBBN hierarchy with support for traffic engineered paths called Ethernet switched paths (ESPs) and support for 1:1 protection switching. PBB provided network scaling, however, retained the control methodology of 802.1Q bridges that offers only limited control over the routing of traffic by the manipulation of Multiple Spanning Tree Protocol (MSTP) cost parameters. In addition, the MSTP control plane supported by PB and PBB does not easily provide 50 ms protection, which is required in some provider applications. Furthermore, MSTP does not provide constant monitoring of alternate paths to provide assurance that, given a failure, it is possible to recover to a route that will provide an equivalent quality of service (QoS). The PBB-TE technology provides enhancements to PBB allowing traffic engineered point-to-point and point-to-multipoint paths within the PBB core. It also provides 1:1 protection switched paths.

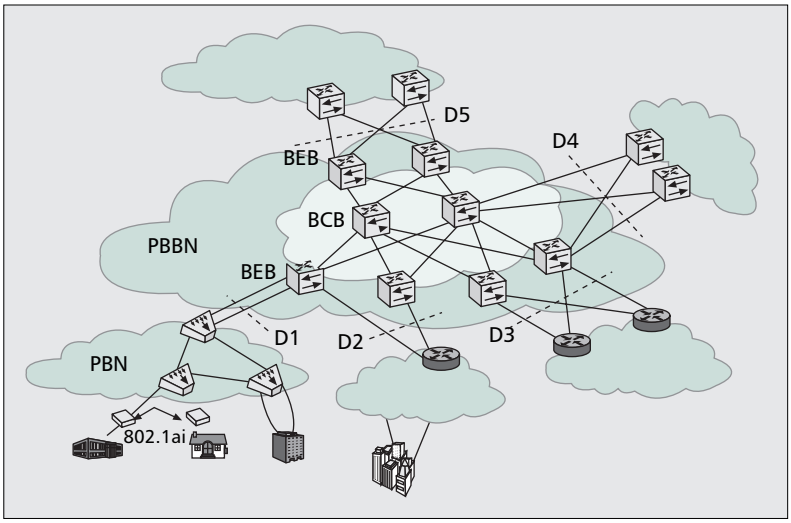

Figure 1. *Example PBBN.*

Provider Backbone Bridged Network Principles

Provider backbone bridged networks allow scaling of PBNs. Each PBBN may connect many PBNs, creating a hierarchy of provider networks with the PBNs serving as the access networks and the PBBN forming the core network. This bridge hierarchy may be further extended by using multidomain PBBNs. The scale of the resulting network is theoretically limited by the MAC address space, which limits the network to 2^{46} globally addressed nodes.

A PBBN provides many features to support carrier grade Ethernet, including:

- Support for 16,776,959 service instances per PBBN domain
- Extension to multiple domains using peer interfaces
- Extension to multiple domains using hierarchical interfaces
- Separation of carrier and customer address spaces
- Separation of service identification and network forwarding mechanisms
- Separation of service creation components and network transport components
- Support for extended CFM

A single PBBN is composed of two different bridge types called backbone edge bridges (BEBs) and backbone core bridges (BCBs). As can be seen from Fig. 1, the BEBs are located at the edge of the PBBNs. This ring of edge bridges performs encapsulation of service frames delivered to the PBBN from the attached networks, forwarding of these frames onto the core of BCBs, and de-encapsulation of service frames for delivery to the attached networks. Within the core of the PBBN, encapsulated service frames are carried on provisioned backbone service instances (BSIs) within B-VLANs. The BSIs, which are completely within the provider network, may be point-to-point, multipoint, or point-to-multipoint virtual networks.

A PBBN can carry traffic from a variety of attached network types. The standard supports direct attachment to PBNs, as illustrated in Fig. 1, interface D1. It is also possible for a PBBN to attach to an enterprise network by using a port-based interface or adding an IEEE 802.1ad adaptation layer between the PBBN and the enterprise bridged network. In addition, it is possible to use the PBBN to transport other multiprotocol services (Fig. 1, interfaces D2 and D3).

A PBBN may be extended by joining together multiple PBBNs to create a multidomain PBBN. Two types of interfaces are supported by 802.1ah for multidomain PBBNs. The first is called a hierarchal interface. This creates a hierarchy of PBBNs where each additional PBBN layer performs an additional encapsulation. Hierarchal encapsulation may be extended up to the maximum frame size allowed by the network. These successive encapsulations provide a multiplexing hierarchy. Each layer of this hierarchy multiplies the maximum number of provider services that could be carried by 16,776,959 since each of the PBBNs of the lower level is treated as a PBN by the next level, which introduces a completely new set of 16,776,959 service identifiers. A single PBBN may carry about 2^{24} services, two levels may carry about 2^{48}, and three levels allow about 2^{72} services (exceeding the IEEE medium access control [MAC] address limit of 2^{46} independent nodes).

A PBBN may also be extended to a multidomain network using the peer interface as specified by 802.1ah. The peer interface allows the connection of two PBBNs without increasing the encapsulation level. When two PBBNs are connected by a peer interface, the service instances that are extended between the PBBNs are mapped by the interface into the other network. The scaling of peer connected networks comes from locality of some service instances. In a peer connected network with two PBBNs, the total number of services is $2 \times 16{,}776{,}959$ – the number of services that extend over both networks. So, for instance, if only one service crosses the peer connection, the total number of services possible is 33,553,917.

Basic PBBN Operation

A service frame passing from a PBN through a PBBN and back to another PBN is carried by an S-VLAN, which the PBBN extends between the

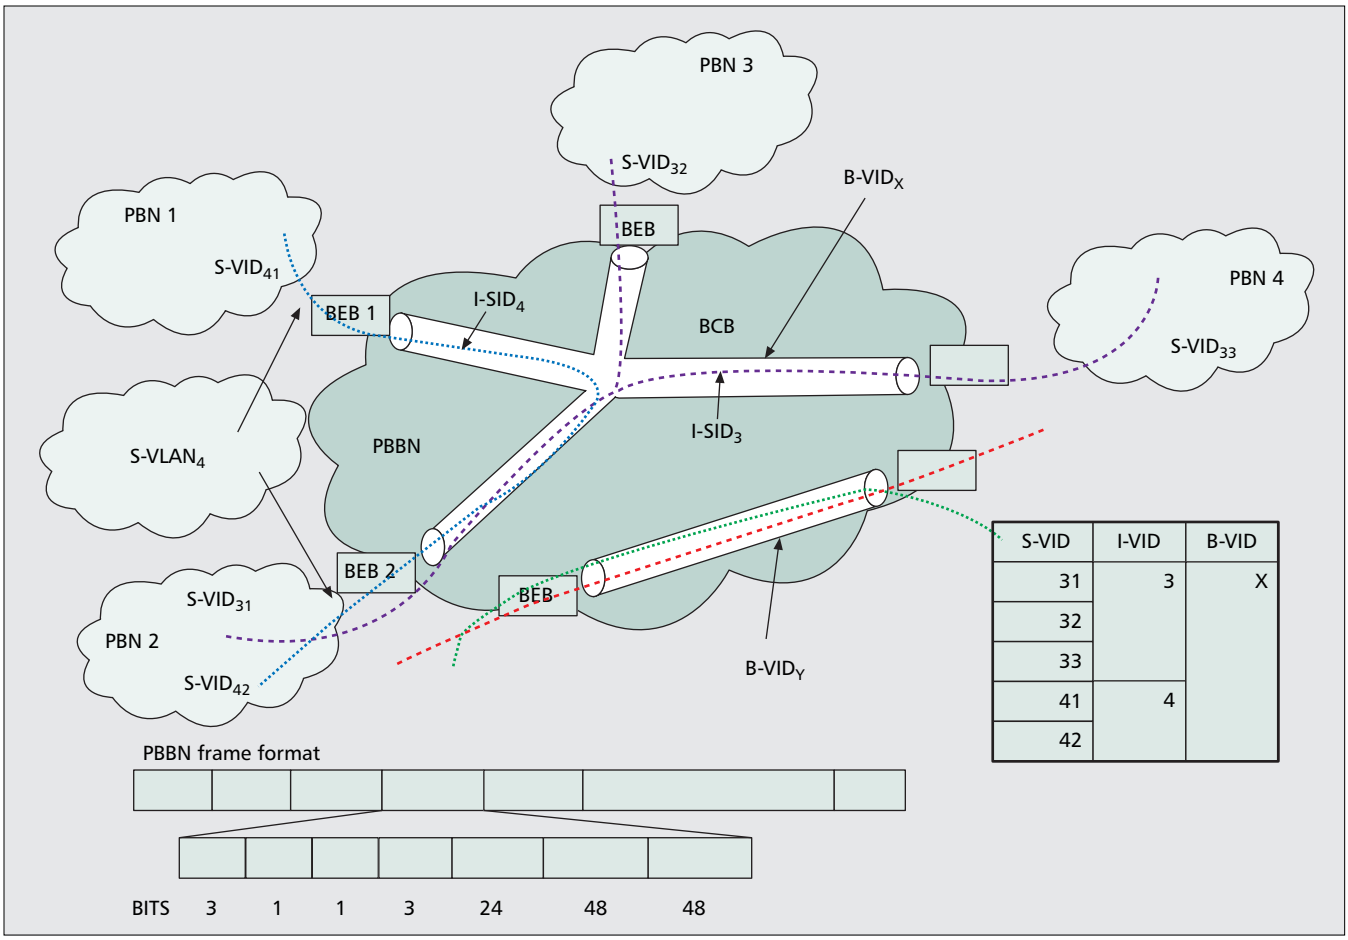

Figure 2. *PBBN functions.*

two PBNs using a BSI (Fig. 2). For instance, S-VLAN$_4$ is extended from PBN 1 to PBN 2 by the PBBN. In PBN 1 S-VLAN$_4$ is identified by S-VID$_{41}$, while it is identified by S-VID$_{42}$ in PBN 2. As a service frame on S-VLAN$_4$ enters BEB 1, the frame is encapsulated. The encapsulation adds new source and destination MAC addresses (the backbone destination address [B-DA] and backbone source address [B-SA]), a service instance tag called an I-TAG, and a new S-VLAN Tag called a B-TAG, which is independent from and in addition to any S-VLAN tag used by the PBN networks.

The B-DA and B-SA added by the BEB are MAC addresses within the PBBN that identify the source and destination within the PBBN. For instance, the service frame entering BEB 1 on S-VLAN$_4$ will be encapsulated with a B-SA identifying a port on BEB 1 where the frame was encapsulated. The B-DA for the same frame will identify a port in BEB 2. The B-DA address is determined in BEB 1 by a learned association between the service frame's DA with a B-DA.

The BEB also adds a tag used to identify the service at the edges of the PBBN. This tag is the I-TAG. The I-TAG contains a 24-bit service ID (I-SID), a 3-bit priority code point (I-PCP) field, 1 bit of drop eligible indicator (I-DEI), and a customer SA (C-SA) and customer DA (C-DA) field. The I-SID is a provisioned value identifying the S-VLAN within the PBBN. The C-DA and C-SA carry the DA and SA fields of the service frame. The PCP and DEI fields are taken from the service frame, and carry the PCP and DEI of the service frame over the PBBN.

The B-TAG added by the BEB identifies a backbone VLAN (B-VLAN) used to carry the frame in the core of the PBBN. B-VLANs are normal S-VLANs identified by a 4094 B-VID.

All forwarding of frames over the PBBN is based on the B-VLAN and B-DA. The I-TAG is used only within the BEBs at the edge of the network to create the mapping between the attached S-VLANs and the PBBN.

To create a standard for PBBNs the IEEE 802.1 committee extended the 802.1Q [1] architectural reference model, abstracting the bridge relay model (commonly called the "baggy pants" model) into a logical bridge component. Bridge components allowed the IEEE to define a set of building blocks from which the new provider bridges could easily be specified. For a description of the 802.1ah architecture and discussion of the S-VLAN component, B-component, and I-component types used as building blocks for 802.1ah, see [9].

Traffic Engineering Principles for PBBNs

The same network that provides traditional LAN services can be enhanced to provide traffic engineered services. The first step in adding TE

capabilities to a PBBN is to remove a range of VIDs from the control of MSTP and assign them to PBB-TE. A PBB-TE region is formed by the set of PBBs that allow an external agent control over the common subset of VIDs assigned to PBB-TE. A PBB-TE service is then provided at special interfaces that encapsulate customer frames onto ESPs, thus allowing the PBBN operator to offer TE services.

To create a PBB-TE region, the operator configures a set of VIDs used throughout the region. A VID belonging to this space is called an ESP-VID and is associated with a special value of the multiple spanning tree identifier, indicating that this specific VID is not under the control of a spanning tree protocol.

The allocation of customer frames to BSIs that are transported by TE service instances is done at the IB-BEBs (a BEB that contains one B-component and a non-zero number of I-components) at the edges of the PBB-TE region. The B-SAs encapsulating the customer frames contain the MAC addresses of the customer backbone ports (CBPs), which have the capability to map BSIs to B-VLAN and provide the termination points for ESPs. The CBP MACs associated with ESPs are called ESP-MACs.

The external agent, which may be either a management entity or a control plane instance, provisions the ESPs within a PBB-TE region by populating the filtering tables of the corresponding BEBs and BCBs on the paths with table entries for the ESP-MACs and ESP-VIDs.

The ability of a PBB-TE region to utilize an external management or control plane agent is facilitated by PBB because the encapsulating ESP-MACs are allocated by the provider and therefore can all be identified in the provider's topology by the external agent.

The external PBB-TE management/control plane is responsible for maintaining and controlling all the topology information to support point-to-point or point-to-multipoint ESPs over the PBB-TE region. The PBB-TE topology can coexist with active topologies associated with spanning tree protocols.

The agent forms a topology of CBP rooted trees from each CBP belonging to a PBB-TE region to a specific subset of any of other CBPs in this region. These trees define the paths taken by the frames that belong to ESPs within the PBB-TE region. Each such tree is further qualified by the PBB-TE reserved ESP-VID, which enables the construction of independent trees per ESP-VID. The agent routes the ESPs along these trees by explicitly populating the filtering databases (FDBs) in the bridges along a tree with the static filtering entries containing the CBP MAC DA and ESP-VIDs. The agent also manages the bandwidth of all ESPs along each routing tree. For each destination CBP that is part of a PBB-TE routing tree, PBB-TE will maintain tree(s) that provide co-routed reverse path(s) from the CBP MAC DA to the CBP MAC SA. The ESP-VID(s) used in this reverse ESP(s) need not have the same ESP-VID used for the forward ESP. Accordingly each of the provisioned ESPs can be identified by a 3-tuple:

<ESP-MAC DA, ESP-MAC SA, ESP-VID>,

where the three component fields are described as follows:

- The ESP-MAC SA is the address of the port encapsulating the customer service instance, which by configuration is the same as the address of the internally connected source CBP on the ESP initiating IB-BEB.
- The ESP-MAC DA identifies the CBP destination address(es).
- The ESP-VID is a VID value distinguishing among ESPs having the same <destination, origin> pair. It can only take values that are allocated to the PBB-TE region.

Key properties of an ESP are that:

- The ESP is identified at all points along its path by a single identifier <ESP-MAC DA, ESP-MAC SA, ESP-VID>.
- The information referenced for forwarding <ESP-MAC DA, ESP-VID> is contained within the ESP identifier.
- The information referenced for forwarding <ESP-MAC DA, ESP-VID> does not change along the length of the ESP.

Two types of TE service instances are described in the PBB-TE standard:

- A point-to-point (PtP) TE service instance provided by a pair of co-routed point-to-point ESPs that have the same endpoints, forming a bidirectional service and correspondingly by a pair of 3-tuples: < DA1, SA1, VID1>, < SA1, DA1, VID2>. The VLAN identifiers can be the same or different.
- A point-to-multipoint (PtMP) TE service instance which is provided by one multipoint ESP plus n unidirectional PtP ESPs, routed along the leaves of the multicast ESP and correspondingly by $n + 1$ 3-tuples: <DA, SA, VID>, <SA, SA1, VID1>, …,<SA, SAn, VIDn>. The DA used by the root CBP identifies the list of MAC addresses {SA1, SA2, …, SAn}.

Although an ESP is identified by the 3-tuple <ESP-MAC DA, ESP-MAC SA, ESP-VID>, only the ESP-MAC DA, ESP-VID pair is used for forwarding decisions. It is possible to treat the combined (ESP-MAC DA, ESP-VID) field as though it was a single 59-bit address, where 12 bits are the ESP-VID and 47 bits are the ESP-MAC DA (allowing for the local reserved bits in the MAC space). This allows PBB-TE to consider the ESP-VID field as a path selector to the destination CBP rather than a B-VID, allowing up to 2^{12} unique routing trees to any single CBP. Figure 3 depicts a PBB-TE region where some example TE service instances have been configured.

If the ESP-VID range delegated is the full 4094 possible values, each CBP termination can sink 2^{12} routing trees, and the theoretical network maximum is about 2^{59} ESPs. Thus, PBB-TE allows scaling to a very large number of ESPs. The number is constrained by issues related to coordinated management, independent management of groups of nodes or independent routing choices through parts of the path, limit through intermediate nodes, and so on.

The following requirements will be met by a bridge supporting PBB-TE:

- Flooding on the ESP-VID space may result

The same network that provides traditional LAN services can be enhanced to provide Traffic Engineered (TE) services. The first step in adding TE capabilities to a PBBN is to remove a range of VIDs from the control of MSTP and assign them to PBB-TE.

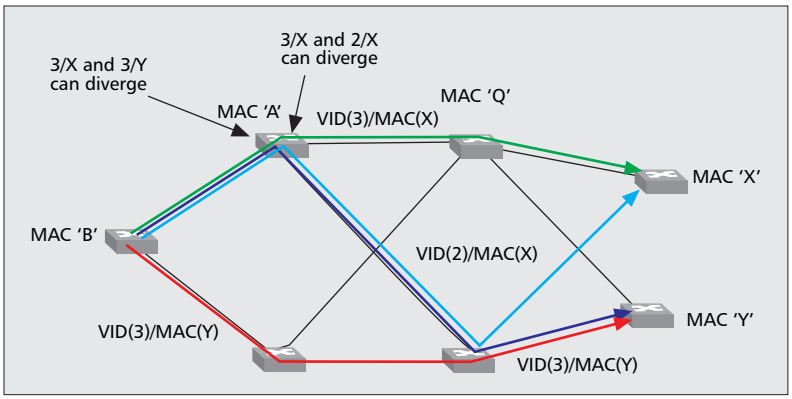

Figure 3. *Ethernet switched path.*

in unbounded looping and replication as the ESP-VID is not associated with a spanning tree. For this reason, frames associated with the ESP-VID are discarded when a static filtering entry corresponding to the (ESP-MAC DA, ESP-VID) is not found in the FDB.

- Where ESPs have been correctly provisioned, learning would not interfere with the proper operation of bridges, but in the presence of provisioning errors, learning may result in undesired behavior. Correspondingly MAC learning is not performed on receipt of a frame carrying an ESP-VID.
- The active topology for the ESP-VIDs allocated to PBB-TE is no longer controlled by spanning tree protocols but by an external agent responsible for setting up the ESPs.

In a network supporting both VLAN (i.e., PBB) and ESP (PBB-TE) traffic, the establishment of an ESP and the policing of traffic at the ingress of the ESP are necessary but not sufficient conditions to ensure that traffic associated with an ESP receives the QoS intended by the network operator. The operator should further ensure that the priority of VLAN traffic does not exceed that of ESP traffic, and that sufficient network capacity is available to accommodate broadcast-unknown traffic, spanning tree control traffic, and other traffic associated with VLAN usage.

The static and provisioned nature of an ESP provides support for TE. These properties of the ESP, together with policing and queuing functions supported by a bridge, can provide per ESP QoS.

Protection Switching of Traffic Engineered Paths

PBB-TE provides a scalable end-to-end resiliency mechanism that offers bidirectional end-to-end 1:1 linear protection capabilities for PtP TE service instances in a PBB-TE region. A dedicated protection PtP TE service instance is established for one particular working PtP TE service instance, and the traffic is automatically switched from the working (primary) TE service instance to the protection (backup) TE service instance when a failure occurs on the primary entity. The protection entity is pre-established, enabling availability of the resources when a defect is detected and a corresponding sub-50-ms switchover.

The PBB-TE linear protection scheme can be configurable to be "revertive" or "non-revertive," where traffic reception (and transmission where applicable) reverts, or not, to the working path automatically once operations, administraion, and maintenance (OAM) [4] indicates that the fault or defect has cleared. It also incorporates holdoff and wait to restore timers. The holdoff timer allows the fault to be protected by a lower layer or upstream protection switching mechanism. This obviously slows the overall recovery time for a fault within the protection domain. The wait to restore timer ensures that the performance of the working path is fully restored before switching back to it. Finally, PBB-TE provides protection with load sharing, which allows some of the services to use the working entity and others to use the protection entity during periods when both entities are available. In the presence of a failure, all traffic is moved to the remaining good TE service instance.

Figure 4 depicts a network where two PtP TE service instances have been provisioned. The service instance at the top of the figure consists of two ESPs each identified by one 3-tuple: <CBP-B, CBP-A, VID-1> depicted in dark green and <CBP-A, CBP-B, VID-2> in light green. The service instance at the bottom consists of two other ESPs identified by the 3-tuples <CBP-B, CBP-A, VID-3> and <CBP-A, CBP-B, VID-4> and depicted in black and gray, respectively. A set of BSIs is assigned to the protection group consisting of the two TE service instances, and have TE-1 assigned as their working TE service instance and TE-2 assigned as their protecting TE service instance.

In Fig. 4 the VLAN membership of the bridge ports on the BEBs is depicted by the color of the boxes that are placed inside each of the B-components. For example, the upper right PNP-1 port on the west B-component is a member of the VID-1 (dark blue), while the CBP port on the same component is a member of VID-2 (light blue) and VID-4 (orange). Correspondingly, in this example only VID-1 VLAN-tagged frames can egress the top right PNP-1, and only VID-2 and VID-4 VLAN-tagged frames can egress from the CBP-A port on the same west B-component.

Each of the TE service instances is monitored by an independent maintenance association (MA). Two up maintenance association endpoints (MEPs), associated with each of these MAs, are configured on the CBPs that terminate the TE service instance in question. Each of these MEPs has its own primary VID, VID-1 for the MEP on the west B-component associated with the top TE service instance and VID-2 for the MEP on the east B-component. In this configuration each MEP receives frames tagged with any of the VIDs in the MA, but sends frames tagged only with that MEP's primary VID. In particular, in the depicted example the MEP for the top entity on the west B-component can send only VID-1 tagged CCMs, while the corresponding MEP on the east component can only send VID-2 tagged CCMs. Both MEPs can receive CCM frames that are either VID-1 or VID-2 tagged. In Fig. 4 the primary VID of each MEP is depicted by the color of the MEP.

Data traffic is mapped to a TE service

instance by configuring the backbone service instance table on the CBP. The CBP backbone service instance identifier is used to allow specific service instances to be carried by the TE service instance, while the CBP's B-VID column in the backbone service instance table is used to map the identified service instances to specific ESPs. The CBP's B-VID value is depicted in Fig. 4 by the color of the bars just outside the boxes representing the B-components. If all the services in the protection group are mapped to the <CBP-B, CBP-A, VID-1> ESP at CBP-A or the <CBP-A, CBP-B, VID-2> ESP at CBP-B, TE-1 will correspond to the working entity and TE-2 to a standby protection entity. CCM frames are always exchanged on both working and protection TE service instances to regularly check the provided connectivity.

If a fault occurs at any of the ESPs, the MEP on the receiving end will be notified. In particular, if a fault on the <CBP-B, CBP-A, VID-1> ESP occurs, as shown in Fig. 4, the MEP on the east B-component will declare the remote MEP defect by setting the corresponding defect parameter. This is done when a timer that is equal to three times the CCM interval expires. The declaration of such a defect will result in a change of the appropriate B-VID entry in the backbone service instance table from VID-2 to VID-4, which is the ESP-VID of the associated provisioned ESP on the protection TE service instance.

All subsequent CCMs sent by the east MEP on TE-1 will have the RDI field set as long as proper CCMs are not received by the MEP. Reception of a CCM frame with the RDI field set will cause a change of the appropriate B-VID entry in the backbone service instance table of the CBP on the west B-component to the preconfigured value of the protection ESP. The result will be to move the affected service instances to the protection TE service instance, as depicted in Fig. 4.

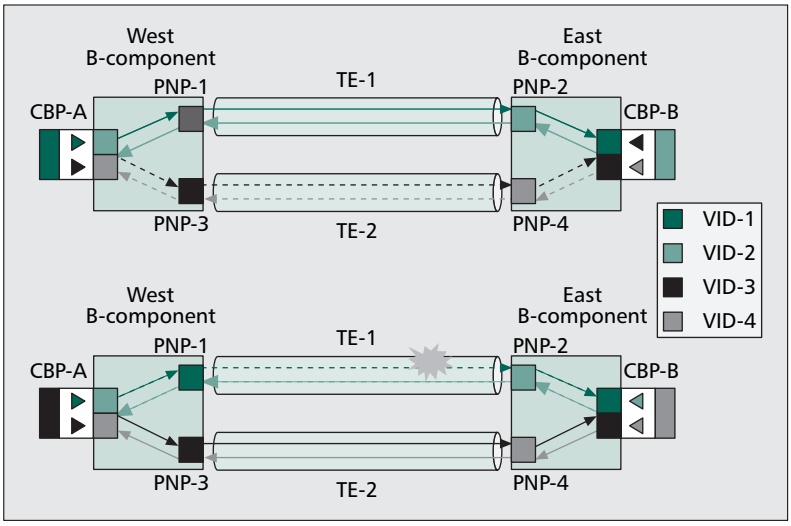

Figure 4. *PBB-TE protection switching.*

Conclusions

A PBBN together with PBB-TE and CFM provides a highly scalable network core with provider oriented features while retaining the simplicity of Ethernet bridging and compatibility with PB. The addition of PBB-TE to PBBNs provides key features for traffic engineered services that are controlled by a central management system or an external control plane rather than MSTP. Each TE service instance is used as a transport for many customer service instances and may provide rapid (sub-50 ms) protection. Protection channels along with working channels are continuously monitored using CFM, providing full service operation. With the addition of IEEE PBB-TE to the existing complement of PB, PBBN, and CFM, Ethernet is now equipped to build a large-scale full-featured provider network.

References

[1] L. Fang *et al.*, "The Evolution of Carrier Ethernet Services: Requirements and Deployment Case Studies," *IEEE Commun. Mag.*, Mar. 2008, pp. 69–76.

[2] IEEE 802.1ad-2005, "IEEE Standard for Local and Metropolitan Area Networks — Virtual Bridged Local Area Networks-Amendment 4: Provider Bridges," 2005.

[3] IEEE P802.1ah-2008, "Standard for Local and Metropolitan Area Networks — Virtual Bridged Local Area Networks — Amendment 6: Provider Backbone Bridges."

[4] IEEE 802.1ag-2007, "IEEE Standard for Local and Metropolitan Area Networks — Virtual Bridged Local Area Networks-Amendment 5: Connectivity Fault Management," 2007.

[5] IEEE P802.1Qay, "Standard for Local and Metropolitan Area Networks — Virtual Bridged Local Area Networks — Amendment: Provider Backbone Bridge Traffic Engineering."

[6] IEEE 802.1Q-2005, "IEEE Standard for Local and Metropolitan Area Networks-Virtual Bridged Local Area Networks," 2005.

[7] S. Salam and A. Sajassi, "Provider Backbone Bridging and MPLS: Complementary Technologies for Next Generation Carrier Ethernet Transport," *IEEE Commun. Mag.*, Mar. 2008, pp. 77–83.

[8] D. Fedyk and D. Allan, "Ethernet Data Plane Evolution for Provider Networks," *IEEE Commun. Mag.*, Mar. 2008, pp. 84–89.

[9] G. Parsons, "Ethernet Bridging Architecture," *IEEE Commun. Mag.*, Dec. 2007, pp. 112–19.

Biographies

PAUL BOTTORFF (pbottorf@nortel.com) has been a technologist and visionary in the communication industry for the last 25 years. Starting in the early 1980s with the development of the first Ethernet products at Bridge Communications, Inc, he went on to advance the state of the art by working on the development and standardization of FDDI, ATM, 10 GE WAN, and metro Ethernet. For the past 10 years he has led the development of metro and wide area Ethernet through both innovation and standards development. He is an inventor of a number of new Ethernet transport technologies including PBB, provider backbone transport, provider link state bridging, and 10 GE WAN. He is the chief editor of the IEEE 802.1ah PBB standard and has served in the past as editor for the IEEE development of 10GE Ethernet. He was a founder of the Metro Ethernet Forum, and has served on the board and as technical co-chair since it began in 2001. He received his B.S. from the University of Wisconsin-Madison. He is the holder of six patents in the area of metro and wide area Ethernet.

PANAGIOTIS SALTSIDIS (panagiotis.saltsidis@ericsson.com) holds a Ph.D. in high energy theoretical physics from the University of Stockholm and has worked on string theory in the Department of Applied Mathematics and Theoretical Physics at the University of Cambridge. Since 2001 he has worked for Ericsson Research on Ethernet Standardization. He is one of the main contributors to a number of IEEE 802.1 standards including CFM, PBBs, Multiple Registration Protocol, two-port MAC relay, data driven and data dependent CFM, and others. His is the sole editor of IEEE 802.1Qay (PBB-TE) and holds 12 patents in the area of metro Ethernet.