IEEE802.1Qay/D0.0

Panagiotis Saltsidis

ERICSSON

TAKING YOU FORWARD

D0.0 – Main items

- Placeholders for future material
- Project's goal statement
- Updates in subclause 8
- Informative annex giving an overview of PBB-TE

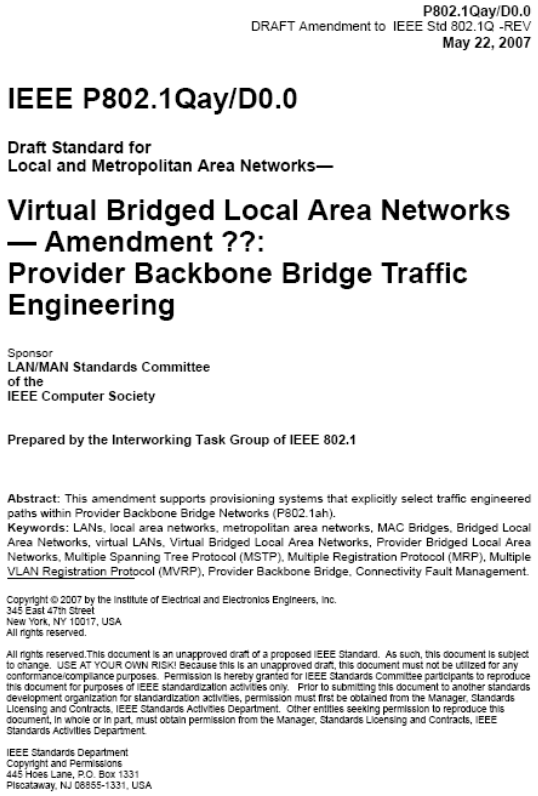

P802.1Qay/D0.0
DRAFT Amendment to IEEE Std 802.1Q -REV
May 22, 2007

IEEE P802.1Qay/D0.0

Draft Standard for
Local and Metropolitan Area Networks—

Virtual Bridged Local Area Networks
— Amendment ??:
Provider Backbone Bridge Traffic
Engineering

Sponsor
LAN/MAN Standards Committee
of the
IEEE Computer Society

Prepared by the Interworking Task Group of IEEE 802.1

Abstract: This amendment supports provisioning systems that explicitly select traffic engineered paths within Provider Backbone Bridge Networks (P802.1ah).
Keywords: LANs, local area networks, metropolitan area networks, MAC Bridges, Bridged Local Area Networks, virtual LANs, Virtual Bridged Local Area Networks, Provider Bridged Local Area Networks, Multiple Spanning Tree Protocol (MSTP), Multiple Registration Protocol (MRP), Multiple VLAN Registration Protocol (MVRP), Provider Backbone Bridge, Connectivity Fault Management.

IEEE Standards Department
Copyright and Permissions
445 Hoes Lane, P.O. Box 1331
Piscataway, NJ 08855-1331, USA

Scope

- Identifies a method for splitting the B-VID space between MSTP and Provisioned control.
- Supports discard of unknown destination addresses.
- Provides extensions to CFM supporting the operation of Continuity Check and Loopback protocols on provisioned paths.
- Provides extensions to PBB supporting one-to-one protection switching.
- Provides any required SNMP MIB management extensions.

Subclause 8

- Active topology enforcement no longer exclusively controlled by Spanning Tree Protocols
- External agent provides the active topology by setting up provisioned Ethernet Switched Paths.
- Forwarding does not imply Learning.
- A Bridge supporting PBB-TE can support any of the spanning tree protocols but can also assign specific VLAN to provisioned ESPs.
- VIDs associated with the PBB-TE ESPs are assigned to a special MSTID (0xFFE)

Subclause 8

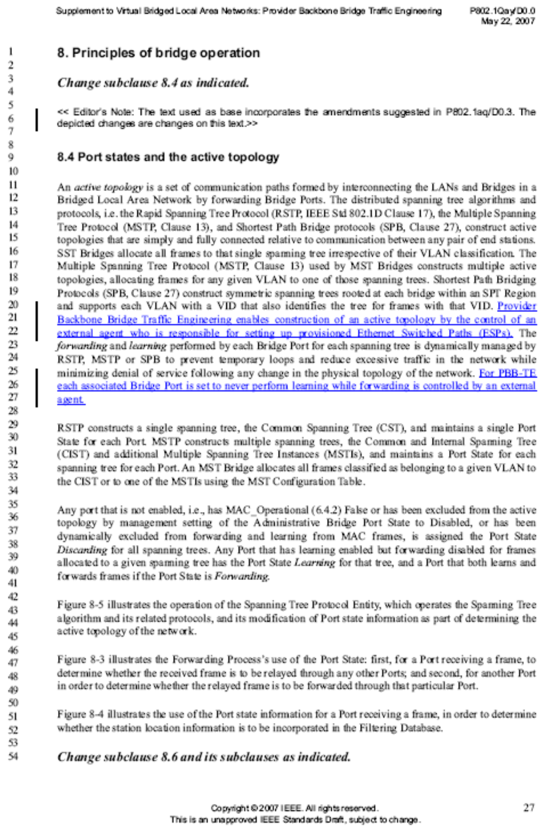

Supplement to Virtual Bridged Local Area Networks: Provider Backbone Bridge Traffic Engineering P802.1Qay/D0.0
May 22, 2007

8. Principles of bridge operation

Change subclause 8.4 as indicated.

<< Editor's Note: The text used as base incorporates the amendments suggested in P802.1aq/D0.3. The depicted changes are changes on this text.>>

8.4 Port states and the active topology

An *active topology* is a set of communication paths formed by interconnecting the LANs and Bridges in a Bridged Local Area Network by forwarding Bridge Ports. The distributed spanning tree algorithms and protocols, i.e. the Rapid Spanning Tree Protocol (RSTP, IEEE Std 802.1D Clause 17), the Multiple Spanning Tree Protocol (MSTP, Clause 13), and Shortest Path Bridge protocols (SPB, Clause 27), construct active topologies that are simply and fully connected relative to communication between any pair of end stations. SST Bridges allocate all frames to that single spanning tree irrespective of their VLAN classification. The Multiple Spanning Tree Protocol (MSTP, Clause 13) used by MST Bridges constructs multiple active topologies, allocating frames for any given VLAN to one of those spanning trees. Shortest Path Bridging Protocols (SPB, Clause 27) construct symmetric spanning trees rooted at each bridge within an SPT Region and supports each VLAN with a VID that also identifies the tree for frames with that VID. Provider Backbone Bridge Traffic Engineering enables construction of an active topology by the control of an external agent who is responsible for setting up provisioned Ethernet Switched Paths (ESPs). The *forwarding* and *learning* performed by each Bridge Port for each spanning tree is dynamically managed by RSTP, MSTP or SPB to prevent temporary loops and reduce excessive traffic in the network while minimizing denial of service following any change in the physical topology of the network. For PBB-TE each associated Bridge Port is set to never perform learning while forwarding is controlled by an external agent.

RSTP constructs a single spanning tree, the Common Spanning Tree (CST), and maintains a single Port State for each Port. MSTP constructs multiple spanning trees, the Common and Internal Spanning Tree (CIST) and additional Multiple Spanning Tree Instances (MSTIs), and maintains a Port State for each spanning tree for each Port. An MST Bridge allocates all frames classified as belonging to a given VLAN to the CIST or to one of the MSTIs using the MST Configuration Table.

Any port that is not enabled, i.e., has MAC_Operational (6.4.2) False or has been excluded from the active topology by management setting of the Administrative Bridge Port State to Disabled, or has been dynamically excluded from forwarding and learning from MAC frames, is assigned the Port State *Discarding* for all spanning trees. Any Port that has learning enabled but forwarding disabled for frames allocated to a given spanning tree has the Port State *Learning* for that tree, and a Port that both learns and forwards frames if the Port State is *Forwarding*.

Figure 8-5 illustrates the operation of the Spanning Tree Protocol Entity, which operates the Spanning Tree algorithm and its related protocols, and its modification of Port state information as part of determining the active topology of the network.

Figure 8-3 illustrates the Forwarding Process's use of the Port State: first, for a Port receiving a frame, to determine whether the received frame is to be relayed through any other Ports; and second, for another Port in order to determine whether the relayed frame is to be forwarded through that particular Port.

Figure 8-4 illustrates the use of the Port state information for a Port receiving a frame, in order to determine whether the station location information is to be incorporated in the Filtering Database.

Change subclause 8.6 and its subclauses as indicated.

27

Placeholders

- **Clauses 3, 4, 5,…**
- **Clause 26**
  - **26.9 Use of Ethernet Switched Paths**
  - **26.10 Ethernet Switched Path Protection**

Annex M

Supplement to Virtual Bridged Local Area Networks: Provider Backbone Bridge Traffic Engineering P802.1Qay/D0.0
May 22, 2007

Insert the following Annex after Annex L

Annex M (informative)

M.1 Provider Backbone Bridges - Traffic Engineering Overview

In a standard Provider Backbone Bridged Network (802.1ah PBBN) traffic engineering is limited as a consequence of the use of IEEE 802.1Q MSTP control plane protocols which control the population of the bridge filtering tables. The underlying 802.1Q/802.1ad/802.1ah bridge relays however don't have any inherent characteristics which prevent full traffic engineering. For example, the IEEE Shortest Path Bridging (802.1aq) project is improving link utilization by replacing the MSTP control plane with a shortest path spanning tree control plane. Provider Backbone Bridges - Traffic Engineering is a method for providing full traffic engineering of point-to-point paths in an 802.1ah network. To do this PBB-TE replaces the MSTP control plane with either a management plane or an external control plane and then populates the bridge filtering tables of the component 802.1ad and 802.1ah bridge relays by creating static filtering table entries (see Figure B-1).

The ability of PBB-TE to utilize an external management or control plane is facilitated by 802.1ah because the B-DAs are all managed by the Provider and therefore can all be discovered and identified in the Provider's topology by the external management or control plane. The description of PBB-TE provided here is based on using PBB-TE within a PBBN, however it is possible to extend the use of PBB-TE to other environments where the Provider controls the SA and DA addresses. An example of such an environment is where the Provider and Customer agree on a specific set of MAC address as part of an SLA. The Provider then only accepts frames with a source and destination addresses as specified in the SLA. Though these arrangements are possible they require some form of static agreement on MAC address which are used between the customer and provider.

The external PBB-TE management/control plane is responsible for maintaining and controlling all the topology information to support point-to-point unidirectional Ethernet Switched Paths (ESP) over the PBBN. The PBB-TE topology can co-exist with the existing active MSTP or with the SPB topology by allocating B-VID spaces to PBB-TE, MSTP, or SPB, or PBB-TE can be stand alone. PBB-TE takes control of a range of B-VIDs from the Backbone Core Bridges (BCB) and Backbone Edge Bridges (BEB) of the PBBN.

The PBB-TE management/control plane forms a topology of B-DA rooted trees. For each <B-DA, B-VID> pair configured by PBB-TE an independent tree is maintained. ESPs are routed by PBB-TE along a tree selected by the B-VID to the destination B-DA. For a single B-DA the number of separate routing trees may be up to the number of PBB-TE reserved B-VIDs. The B-VIDs may be reused for every B-DA, therefore the total number of routing trees maintained by the PBB-TE management/control plane may be up to the number of B-DAs times the number of PBB-TE reserved B-VIDs. The trees maintained by PBB-TE for routing ESPs do not have to be spanning since they only require connectivity to all the source B-MACs which have ESPs to the specific B-DA. Each tree may connect to as many B-SAs as desired with the only limits being implementation imposed table sizes. The PBB-TE management/control plane may use any algorithm desired to select the path for a routing tree thereby providing complete route selection freedom. The PBB-TE management/control plane also manages the bandwidth of all ESPs along each routing tree. For each B-SA which is part of a routing tree maintained by the PBB-TE management/control plane, PBB-TE will maintain a routing tree which provides a co-routed reverse path from the B-DA to the B-SA. The B-

45

8. Principles of bridge operation

Change subclause 8.4 as indicated.

<< Editor's Note: The text used as base incorporates the amendments suggested in P802.1aq/D0.3. The depicted changes are changes on this text.>>

8.4 Port states and the active topology

An *active topology* is a set of communication paths formed by interconnecting the LANs and Bridges in a Bridged Local Area Network by forwarding Bridge Ports. The distributed spanning tree algorithms and protocols, i.e. the Rapid Spanning Tree Protocol (RSTP, IEEE Std 802.1D Clause 17), the Multiple Spanning Tree Protocol (MSTP, Clause 13), and Shortest Path Bridge protocols (SPB, Clause 27), construct active topologies that are simply and fully connected relative to communication between any pair of end stations. SST Bridges allocate all frames to that single spanning tree irrespective of their VLAN classification. The Multiple Spanning Tree Protocol (MSTP, Clause 13) used by MST Bridges constructs multiple active topologies, allocating frames for any given VLAN to one of those spanning trees. Shortest Path Bridging Protocols (SPB, Clause 27) construct symmetric spanning trees rooted at each bridge within an SPT Region and supports each VLAN with a VID that also identifies the tree for frames with that VID. Provider Backbone Bridge Traffic Engineering enables construction of an active topology by the control of an external agent who is responsible for setting up provisioned Ethernet Switched Paths (ESPs). The *forwarding* and *learning* performed by each Bridge Port for each spanning tree is dynamically managed by RSTP, MSTP or SPB to prevent temporary loops and reduce excessive traffic in the network while minimizing denial of service following any change in the physical topology of the network. For PBB-TE each associated Bridge Port is set to never perform learning while forwarding is controlled by an external agent.

RSTP constructs a single spanning tree, the Common Spanning Tree (CST), and maintains a single Port State for each Port. MSTP constructs multiple spanning trees, the Common and Internal Spanning Tree (CIST) and additional Multiple Spanning Tree Instances (MSTIs), and maintains a Port State for each spanning tree for each Port. An MST Bridge allocates all frames classified as belonging to a given VLAN to the CIST or to one of the MSTIs using the MST Configuration Table.

Any port that is not enabled, i.e., has MAC_Operational (6.4.2) False or has been excluded from the active topology by management setting of the Administrative Bridge Port State to Disabled, or has been dynamically excluded from forwarding and learning from MAC frames, is assigned the Port State *Discarding* for all spanning trees. Any Port that has learning enabled but forwarding disabled for frames allocated to a given spanning tree has the Port State *Learning* for that tree, and a Port that both learns and forwards frames if the Port State is *Forwarding*.

Figure 8-5 illustrates the operation of the Spanning Tree Protocol Entity, which operates the Spanning Tree algorithm and its related protocols, and its modification of Port state information as part of determining the active topology of the network.

Figure 8-3 illustrates the Forwarding Process's use of the Port State: first, for a Port receiving a frame, to determine whether the received frame is to be relayed through any other Ports; and second, for another Port in order to determine whether the relayed frame is to be forwarded through that particular Port.

Figure 8-4 illustrates the use of the Port state information for a Port receiving a frame, in order to determine whether the station location information is to be incorporated in the Filtering Database.

Change subclause 8.6 and its subclauses as indicated.

8.6 The Forwarding Process

8.6.1 Active topology enforcement

The Forwarding Process allocates each received frame to ~~a spanning tree~~ an active topology. If the reception Port State for that ~~spanning tree~~ active topology is ~~Forwarding or~~ Learning, the source address and VID are submitted to the Learning Process. If the reception Port State is Forwarding, each Bridge Port, other than the reception Port, with a Port State of Forwarding for that ~~tree~~ active topology is identified as a potential transmission Port.

An SST Bridge allocates all frames to a single spanning tree, the Common Spanning Tree (CST).

An MST Bridge allocates all frames with a given VID to the CIST or to a Multiple Spanning Tree Instance (MSTI). The allocation can be controlled by configuration of the MST Configuration Table (8.9.1) maintained by the Forwarding Process, subject to constraints (if any) imposed by the allocation of VIDs to FIDs (8.8.7). VIDs allocated to different spanning trees shall also be allocated to different FIDs. VIDs allocated to a given spanning tree may share the same FID.

A Bridge supporting PBB-TE can allocate all frames to a single spanning tree or all frames with a given VID to an MSTI but has also the capability of allocating frames with a given VID to provisioned ESPs. The allocation can be controlled by configuration of the MST Configuration Table (8.9.1) maintained by the Forwarding Process. VIDs allocated to PBB-TE ESPs will be identified by the PBB-TE MSTID.

Change subclause 8.8 and its subclauses as indicated.

8.8 The Filtering Database

8.8.1 Static Filtering Entries

A Static Filtering Entry specifies

a) A MAC Address specification, comprising
   1) An Individual MAC Address; or
   2) All Individual Addresses, for which no more specific Static Filtering Entry exists; or
   3) A group MAC Address; or
   4) All Group Addresses, for which no more specific Static Filtering Entry exists; or
   5) All Unregistered Group Addresses, for which no more specific Static Filtering Entry exists.

b) A VLAN identifier specification, comprising:
   1) The VID of a specific VLAN to which the static filtering information applies; or
   2) The wildcard VID (see Table 9-2), indicating that the static filtering information applies to all VLANs for which no specific Static Filtering Entry exists.

c) A Port Map, containing a control element for each outbound Port, specifying that a frame with a destination MAC Address and VID that meets this specification is to be
   1) Forwarded, independently of any dynamic filtering information held by the Filtering Database; or
   2) Filtered, independently of any dynamic filtering information; or
   3) Forwarded or filtered on the basis of dynamic filtering information, or on the basis of the default Group filtering behavior for the outbound Port (8.8.6) if no dynamic filtering information is present specifically for the MAC Address.

All Bridges shall have the capability to support the first two values for the MAC Address specification, both values of the VLAN identifier specification, and all three values for each control element for all Static Filtering Entries (i.e., shall have the capability to support a1, a2, b1, b2, c1, c2, and c3 above).

A Bridge that supports Extended Filtering Services shall have the capability to support all four values for the MAC Address specification and all three control element values for all Static Filtering Entries.

For a given MAC Address specification, a separate Static Filtering Entry with a distinct Port Map may be created for each VLAN from which frames are received by the Forwarding Process.

The Port Map may contain a connection_identifier (8.8.11) for each outbound Port.

In addition to controlling the forwarding of frames, Static Filtering Entries for group MAC Addresses provide the Registrar Administrative Control values for the GMRP protocol (Clauses 10, 12, and 12.9.1 of IEEE Std 802.1D). Static configuration of forwarding of specific group addressed frames to an outbound port indicates Registration Fixed on that port: a desire to receive frames addressed to that Group even in the absence of dynamic information. Static configuration of filtering of frames that might otherwise be sent to an outbound port indicates Registration Forbidden. The absence of a Static Filtering Entry for the group address, or the configuration of forwarding or filtering on the basis of dynamic filtering information, indicates Normal Registration.

Change subclause 8.9 and its subclauses as indicated.

8.9 MST, SPT and ESP configuration information

In order to support multiple spanning trees, an MST Bridge has to be configured with an unambiguous assignment of VIDs to spanning trees. This is achieved by:

a) Ensuring that the allocation of VIDs to FIDs (8.8.7) is unambiguous; and
b) Ensuring that each FID supported by the Bridge is allocated to exactly one Spanning Tree.

The first of these requirements is met by configuring a set of VLAN learning constraints and/or fixed VID to FID mappings that are self-consistent, and which define an I Constraint, an S Constraint, or a fixed VID to FID allocation for all VIDs supported by the Bridge.

The second requirement is met by means of the FID to MSTI Allocation Table (8.9.3).

The combination of the VID to FID allocations and the FID to MSTIDs allocations defines a mapping of VIDs to MSTIDs, represented by the MST Configuration Table (8.9.1).

An SPT Bridge can allocate any given VLAN to the CIST, or to one of a number of MSTIs, as well as supporting shortest path bridging for other VLANs within an SPT Region. The assignment of VLANs and VIDs to MSTIs is subject to exactly the same considerations and constraints as for MST Bridges. In addition to the Base VID used to identify frames for the VLAN when transmitted on the CST outside the region, each shortest path VLAN is supported by a number of shortest path VIDs (SPVIDs), one for each SPT Bridge within the region. SPVIDs are dynamically allocated to identify a shortest path VLAN (as identified by the Base VID) and the SPT (rooted at a particular SPT Bridge) as required.

Dynamic SPVID allocation is supported by SPB protocols and allows the addition of bridges to existing SPT Regions without disruptive configuration changes, as well as allowing supporting autoconfiguration of shortest path bridging in simple networks where all user data is assigned to VLAN identified by the default PVID of 1 (<Table 9-2>). However, prior to SPVID allocation, all bridges in an SPT Region have to agree which VLANs are to be shortest path bridged, and which SPT Set is to be used for each of those VLANs. By default all VLANs whose Base VID is allocated to the MSTI will be shortest path bridged using the SPT Primary Set. The SPT configuration information can also identify specific SPT Sets, by specifying that frames for a VLAN with a Base VID that has been associated with a specific MSTI be supported by that set instead of by the MSTI.

An SPT Bridge allows a configuration choice of the protocol used to support MSTIs. That choice includes MSTP as specified in clause 13. If no protocol is chosen the MSTIs will not be constructed but the MSTID can still be used to select SPT Sets.

NOTE—The use of MSTIDs to support allocation of Base VIDs to SPT Sets allows the per VLAN components of the necessary information to be conveyed in the MST Configuration Digest, thus avoiding any need for a separate digest.

A Bridge supporting PBB-TE can support any of the spanning tree protocols but can also assign specific VLANs to provisioned ESPs. This is achieved by allocating the VIDs associated with the PBB-TE ESPs to a special MSTID.

8.9.1 MST Configuration Table

The MSTI Configuration Table specifies an MSTID for each possible VID.

In an MST Bridge each MSTID identifies the MSTI to which the VID is allocated, and an MSTID of zero identifies the CIST.

In an SPT Bridge, MSTIDs that do not appear in the SPT Set List (8.9.5) also identify an MSTI to which the VID is allocated or the CIST. MSTIDs that do appear in the SPT Set List identify the SPT Set for the VID. An MSTID of hex FFF identifies VIDs that can be used as SPVIDs for any SPT Set.

In a Bridge that supports PBB-TE an MSTID of hex FFE identifies VIDs that can be used by ESPs.

The MST Configuration Table cannot be configured directly; configuration of the table occurs as a consequence of configuring the relationships between VIDs and FIDs (8.8.7) and between FIDs and MSTIDs (8.9.3).

8.9.2 MST configuration identification

For two or more MST Bridges to be members of the same MST Region (3.32), it is necessary for those Bridges to be directly connected together (i.e., interconnected only by means of LANs, without intervening Bridges that are not members of the region), and for them to support the same MST Region configuration. Two MST Region configurations are considered to be the same if the correspondence between VIDs and spanning trees is identical in both configurations and they use the same information to identify the configuration, including the same Configuration Name.

NOTE 1—If two adjacent MST Bridges consider themselves to be in the same MST Region despite having different mappings of VIDs to spanning trees, then the possibility exists of undetectable loops arising within the MST Region.

In order to ensure that adjacent MST Bridges are able to determine whether they are part of the same MST Region, the MST BPDU supports the communication of an MST Configuration Identifier (13.7).

NOTE 2—As the MST Configuration Identifier is smaller than the mapping information that it summarizes, there is a small but finite possibility that two MST Bridges will assume that they have the same MST Region Configuration when this is not actually the case. However, given the size of the identifier, this standard assumes that this possibility is sufficiently small that it can safely be ignored. Appropriate use of the Configuration Name and Revision Level portions of the identifier can remove the possibility of an accidental match between MST Configuration Identifiers that are derived from different configurations within a single administrative domain (see 13.7).

8.9.3 FID to MSTID Allocation Table

The FID to MSTID Allocation Table defines, for all FIDs that the Bridge supports, the MSTID to which the FID is allocated. An MSTID of zero is used to identify the CIST.

NOTE—MSTIDs that are present in the MSTI List (12.12) identify spanning tree instances supported by MSTP. MSTIDs that are in the SPT Set List (8.9.5) identify (indirectly) VLANs that are supported by shortest path bridging.

Insert the following Annex after Annex L

Annex M (informative)

M.1 Provider Backbone Bridges - Traffic Engineering Overview

In a standard Provider Backbone Bridged Network (802.1ah PBBN) traffic engineering is limited as a consequence of the use of IEEE 802.1Q MSTP control plane protocols which control the population of the bridge filtering tables. The underlying 802.1Q/802.1ad/802.1ah bridge relays however don’t have any inherent characteristics which prevent full traffic engineering. For example, the IEEE Shortest Path Bridging (802.1aq) project is improving link utilization by replacing the MSTP control plane with a shortest path spanning tree control plane. Provider Backbone Bridges - Traffic Engineering is a method for providing full traffic engineering of point-to-point paths in an 802.1ah network. To do this PBB-TE replaces the MSTP control plane with either a management plane or an external control plane and then populates the bridge filtering tables of the component 802.1ad and 802.1ah bridge relays by creating static filtering table entries (see Figure B-1).

The ability of PBB-TE to utilize an external management or control plane is facilitated by 802.1ah because the B-DAs are all managed by the Provider and therefore can all be discovered and identified in the Provider’s topology by the external management or control plane. The description of PBB-TE provided here is based on using PBB-TE within a PBBN, however it is possible to extend the use of PBB-TE to other environments where the Provider controls the SA and DA addresses. An example of such an environment is where the Provider and Customer agree on a specific set of MAC address as part of an SLA. The Provider then only accepts frames with a source and destination addresses as specified in the SLA. Though these arrangements are possible they require some form of static agreement on MAC address which are used between the customer and provider.

The external PBB-TE management/control plane is responsible for maintaining and controlling all the topology information to support point-to-point unidirectional Ethernet Switched Paths (ESP) over the PBBN. The PBB-TE topology can co-exist with the existing active MSTP or with the SPB topology by allocating B-VID spaces to PBB-TE, MSTP, or SPB, or PBB-TE can be stand alone. PBB-TE takes control of a range of B-VIDs from the Backbone Core Bridges (BCB) and Backbone Edge Bridges (BEB) of the PBBN.

The PBB-TE management/control plane forms a topology of B-DA rooted trees. For each <B-DA, B-VID> pair configured by PBB-TE an independent tree is maintained. ESPs are routed by PBB-TE along a tree selected by the B-VID to the destination B-DA. For a single B-DA the number of separate routing trees may be up to the number of PBB-TE reserved B-VIDs. The B-VIDs may be reused for every B-DA, therefore the total number of routing trees maintained by the PBB-TE management/control plane may be up to the number of B-DAs times the number of PBB-TE reserved B-VIDs. The trees maintained by PBB-TE for routing ESPs do not have to be spanning since they only require connectivity to all the source B-MACs which have ESPs to the specific B-DA. Each tree may connect to as many B-SAs as desired with the only limits being implementation imposed table sizes. The PBB-TE management/control plane may use any algorithm desired to select the path for a routing tree thereby providing complete route selection freedom. The PBB-TE management/control plane also manages the bandwidth of all ESPs along each routing tree. For each B-SA which is part of a routing tree maintained by the PBB-TE management/control plane, PBB-TE will maintain a routing tree which provides a co-routed reverse path from the B-DA to the B-SA. The B-

VID used in this reverse ESP does not have to be the same one used for the forward ESP. The reverse ESP is used by CFM management to monitor the ESPs.

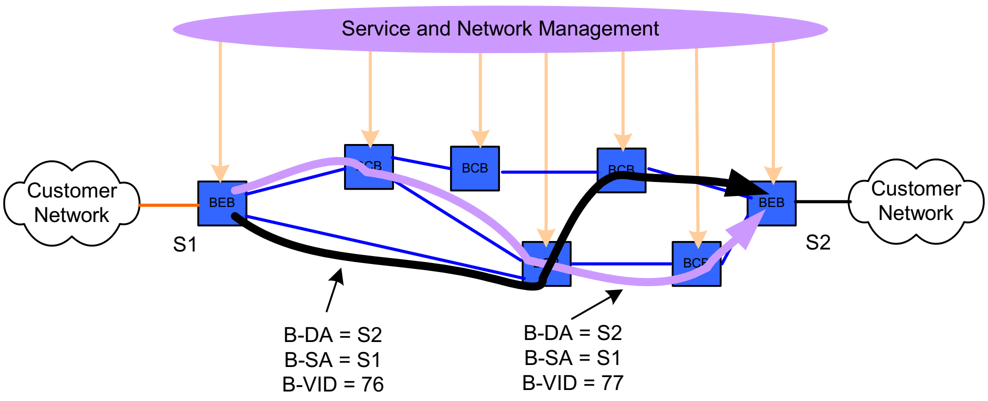

Figure B-1—PBB-TE Network

The relay functions used by PBB-TE are standard bridge forwarding which use a <B-DA, B-VID> tuple rather the B-DA alone for filtering and forwarding decisions. It is possible to treat the combined <B-DA, B-VID> tuple as though it was a single 58 bit address, where 12 bits are the B-VID and 46 bits are the B-DA (allowing for the multicast and local reserved bits in the MAC space). This fact allows PBB-TE to consider the B-VID part of the address as a path selector to the B-DA rather than a B-VLAN ID, allowing up to 4094 unique routing trees to any single B-DA. Typically only a small number of B-VIDs are needed for PBB-TE since it is normally not necessary to have even tens of alternate paths to a single destination. In Figure 1 two paths are configured to reach S2. These two paths are separated by using a different B-VID in combination with the B-DA for a second path.

PBB-TE requires no B-VID translation and operates on most 802.1ah unicast forwarding hardware. PBB-TE allows scaling to an almost infinite number of Ethernet Switched Paths. If the B-VID range delegated is the full 4094 possible values, then each B-DA termination can sink 2^{12} routing trees, and the theoretical network maximum is about 2^{58} ESPs. The forwarding of frames over these ESPs is easily achieved by most existing Ethernet equipment without significantly re-specifying the hardware and management to achieve traffic management.

To make PBB-TE robust, a few aspects of standard 802.1Q bridge forwarding need consideration:

1) Discontinuities in forwarding table configuration for an ESP will result in packets being flooded as “unknown”. As there is no loop free topology for the delegated B-VID range, this will result in unbounded flooding, looping and replication. For this reason flooding of packets with unknown destinations must be disabled for the B-VID range allocated to PBB-TE. Similarly, broadcast and multicast traffic that would be flooded must be filtered at the ingress to the relay function. This function is already supported by 802.1Q.
2) B-SA learning is not required, and may interfere with management/control population of the forwarding tables when combined with the potential for configuration errors. For this reason B-SA learning is disabled for the delegated B-VID range.
3) This approach bypasses spanning tree for the delegated B-VID range, so spanning tree or multiple spanning trees are used only for the non-delegated (traditional operation) B-VID range.
4) When used in conjunction with a best effort spanning tree, traffic on Call Admission Control or “CAC'd” paths requires a higher priority than best effort traffic. When engineering the network,

a reserve bandwidth must be set aside for best effort traffic to allow for spanning tree reconfiguration.

This approach has a number of useful properties:

The use of a global path identifier (the <B-DA, B-VID> tuples) directly for forwarding and with no translation is inherently more robust than alternatives. Any mis-configuration or forwarding table errors resulting in the deviation of a PDU from the intended path will self identify immediately. There is no possibility of collision with other identifier spaces that can mask the fault. This also suggests that no or minimal changes are required to existing CFM and Y.1731 constructs to successfully instrument Ethernet paths.

The ability to explicitly route and pin paths across the network can be combined with call admission control and 802.1Q class-based queuing in order to provide per path QoS. The call admission control function can be enforced by the external management/control plane without any changes to existing Ethernet bridges.

M.2 Example Ethernet Switched Paths

Figure B-2 shows an example Provider Backbone Bridged Network running PBB-TE over the PBBN core. The PBBN B-VID space has been partitioned between MSTP and with B-VIDs 7 and 8 allocated to PBB-TE. PBB-TE has taken over the port forwarding state machines in the Backbone Edge Bridges and Backbone Core Bridges for the allocated B-VID range and controls frame forwarding by adding static entries to the filtering databases of the switches within the PBBN core. MSTP operates normally in parallel to PBB-TE on the other B-VIDs.

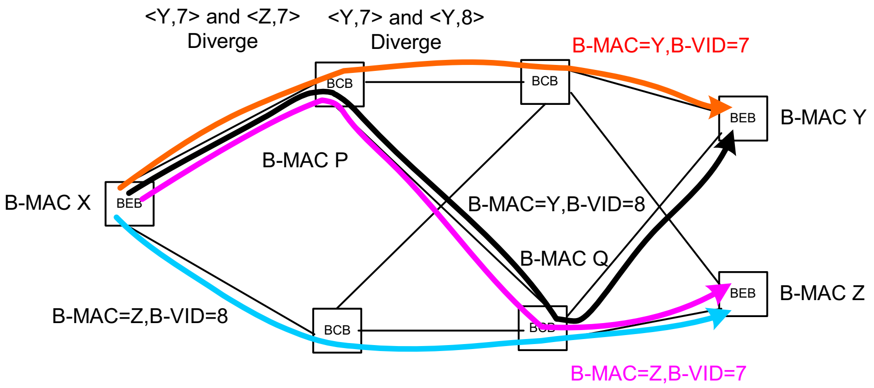

Figure B-2—PBBN with B-VIDs 7 and 8 allocated to PBB-TE operation

The B-VIDs allocated to PBB-TE are not treated as VLAN IDs, instead they are individual instance identifiers for one of a maximum of ‘n’, where ‘n’ is the number of allocated B-VIDs, or possible routing trees to the destination B-DA address. B-VIDs in the allocated range may be re-used for many routing trees within the PBBN as long as the <B-DA, B-VID> tuple is unique. This PBB-TE addressing arrangement results in a 58 bit globally unique destination address that may be shared by any number of B-SAs.

The example in Figure B-2 illustrates the complete route freedom of configured forwarding in bridges. In the example a total of 4 ESPs use 2 B-VIDs to forward traffic to 2 B-DA terminations. At node 'P' above, despite collisions in both the B-DA and the B-VID space, the forwarding properly resolves because the switch uses the B-DA and B-VID together to establish route uniqueness. At node ‘P’ the red and purple

ESPs diverge even though they have the same B-VID because they are addressed to different B-DAs. Likewise at node ‘P’ the red and black ESPs diverge even though they have the same B-DAs because they have different route selector B-VIDs.