PROVIDER BACKBONE BRIDGE WITH TRAFFIC ENGINEERING: A CARRIER ETHERNET TECHNOLOGY OVERVIEW

By Thierno Diallo, Product Specialist

Originally designed as a local-area network (LAN) communication protocol, Ethernet allows computers and nodes to be interconnected within a small network. But to expand its reach into the core infrastructure, Ethernet needed to offer carrier-grade quality of service (QoS) guarantees and reliability to thousands of computers across metropolitan, national and global distances without affecting its simplicity and cost-efficiency. A new addition to the 802.1 standard known as *Provider Backbone Bridge with Traffic Engineering* (PBB-TE), defined as IEEE 802.1 Qay, is widely regarded as Ethernet's carrier-class solution to this challenge.

This standard is based on the PBB frame format (defined in IEEE 802.1ah; also known as Mac-in-Mac), in which an existing Ethernet frame is encapsulated in another Ethernet frame. This approach allows for the separation of customer and carrier Ethernet domains while maintaining Ethernet-forwarding and error-detection techniques.

In PBB-TE, traffic engineering is based on the forwarding control given to the carrier, a departure from legacy Ethernet. By giving forwarding control to the carrier, PBB-TE provides deterministic forwarding; i.e., routes are preconfigured using external management software.

This application note presents this new technology and its evolution compared to legacy Ethernet, provides details about its format and suggests some of the testing requirements for the deployment of such networks.

PBB-TE: An Evolution from Legacy Ethernet

To increase the adoption of Ethernet as a carrier technology, a few changes were needed to make the standard interesting for metropolitan applications. The following details the various additions that were applied to the Ethernet standard leading up to PBB-TE.

1. 802.1Q: VLAN

Virtual local-area network (VLAN), defined in the standard 802.1Q, identified VLAN tags that are used by switching devices to differentiate traffic with a VLAN identifier between 0 and 4095. This allowed for QoS and traffic separation functions on a shared media.

Destination MAC Address	Source MAC Address	Ether Type (0 x 8100)	VLAN Tag	Ether Type	Customer Data	FCS
6	6	2	2	2	46-1500	4

Priority	CFI	VID
3 bits	1 bit	12 bits

2. 802.1ad: Provider Bridge

VLAN tags have a maximum of 4096 IDs, which is sufficient for LANs, but not enough for carriers with multiple routes and subscribers. To improve the scalability of VLANs, provider bridge (also known as Q-in-Q) defined stacked VLAN tags. An additional tag is inserted in an already tagged frame, creating an inner tag and an outer tag. This allowed carriers to add tags specific to their networks without modifying the tags already in the frame inserted by the service provider.

Destination MAC Address	Source MAC Address	Ether Type 0 x 8100 0 x 9200 0 x 9100 0 x 9300	S-VLAN	Ether Type 0 x 8100 0 x 9200 0 x 9100 0 x 9300	C-VLAN	Ether Type	Customer Data	FCS
6	6	2	2	2	2	2	46-1500	4

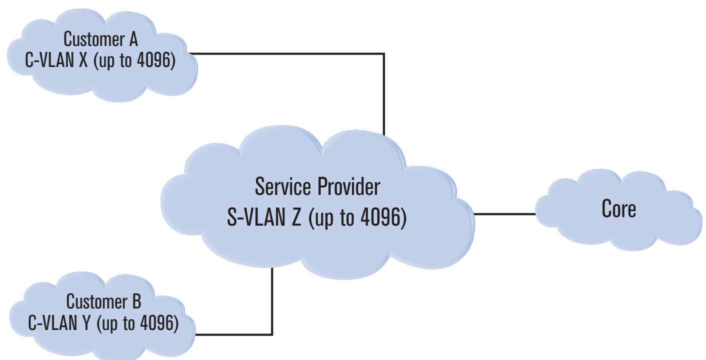

3. 802.1ah: Provider Backbone Bridge (PBB)

As networks grew, core switches needed to manage a greater number of medium access control (MAC) addresses in the forwarding tables. Combined with the limited number of VLANs, this increased the complexity of the networks.

The 802.1ah PBB approach proposed encapsulates a customer's Ethernet frame into a carrier Ethernet frame, complete with its own MAC address space. Within a PBB network, frames are switched according to the destination backbone switch MAC.

The main advantage of this approach is the complete separation of customer and carrier domains, enabling the customer's Ethernet frames to be transparently transported in the carrier's Ethernet frames. This greatly reduces the complexity of the switch-forwarding tables, as the entries are limited to the carrier's network switches.

PBB also added a unique field called I-Tag, which allows the carrier to assign QoS parameters and define a unique customer identifier (I-SID). Therefore, traffic flows are assigned a unique I-Tag per customer, and QoS can be performed per customer instead of per VLAN. Moreover, since the I-SID is 24 bits long, there are up to two million service identifiers.

B-DA	B-SA	Ether Type (0 x 88A8)	B-VID	Ether Type (0 x 88E7)	I-Tag	Customer Ethernet Frame	B-FCS
6	6	2	2	2	4	64-1510	4

Destination MAC Address	Source MAC Address	Ether Type (0 x 8100)	VLAN Tag	Ether Type	Customer Ethernet Frame	FCS
6	6	2	2	2	46-1492	4

4. 802.1Qay: Provider Backbone Bridge with Traffic Engineering (PBB-TE)

This latest standard is based on the Nortel technology known as Provider Backbone Transport (PBT). Essentially based on the PBB frame format, the PBB-TE standard focuses on frame transport within the network as it replaces the existing spanning tree protocol (STP) with a connection-oriented and pre-established path configured by the user.

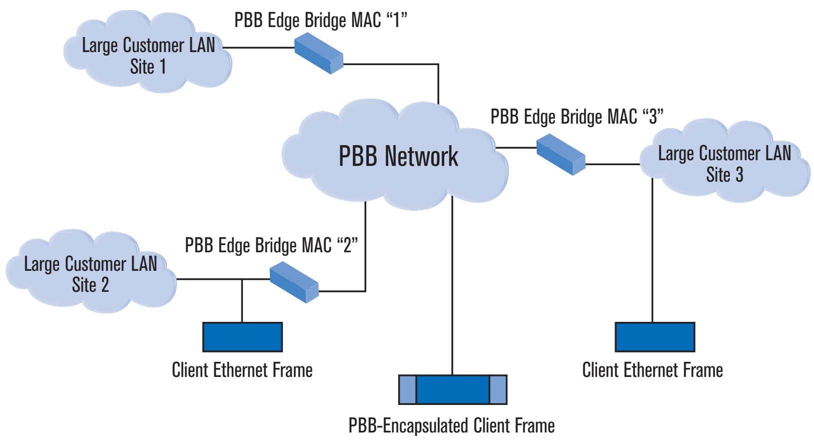

In the example above, the spanning tree is disabled, so the PBB forwarding devices only need to “learn” the MAC addresses of the edge devices to transmit the frames properly.

PBB-TE Basics

1. Architecture

A typical PBB-TE network is composed of two main devices, the PBT edge bridge and the PBT backbone (or core) switch.

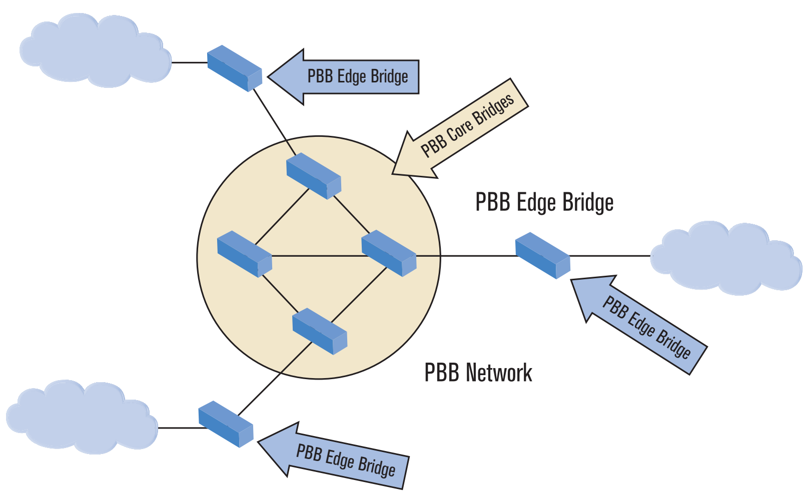

The edge bridge is the interface between the customer network and the service provider network. This device is responsible for the encapsulation or de-encapsulation of the customer’s Ethernet frames with PBB headers as well as the insertion of the proper backbone and I-Tag.

These fields are mandatory, as the frames will be switched within the PBB-TE network based on the backbone destination MAC address (B-DA) and the backbone VLAN ID (B-VID). The I-Tag is used to identify QoS levels as well as the customer carried by the PBB frames via the service ID.

The backbone switch is responsible for the forwarding of PBB frames within the PBB-TE network using predefined routes according to the B-VID. These switches differ from normal Ethernet switch in that they lack STP (and its variant, rapid spanning tree protocol, or RSTP) measures.

2. Frame Forwarding

In legacy Ethernet, switches perform spanning tree and flooding to discover neighboring switches, thus creating dynamic forwarding routes. In Learning mode, a switch forwards frames with unknown destination MAC addresses to all ports. These addresses are then added to the switch table once an incoming frame with that source MAC is detected at a port. Broadcast frames are processed and forwarded by the switch to all ports.

In PBB-TE, switches are configured with static routes by the network operator, ensuring that frames take predetermined paths within the network. The user must configure all the backbone switches in the forwarding table using external management software. In such a situation, frames with destination MAC addresses not yet associated in the port-MAC table will be dropped; they will not be forwarded. Since broadcast frames are not supported in PBB-TE networks, they will also be dropped by the backbone switches.

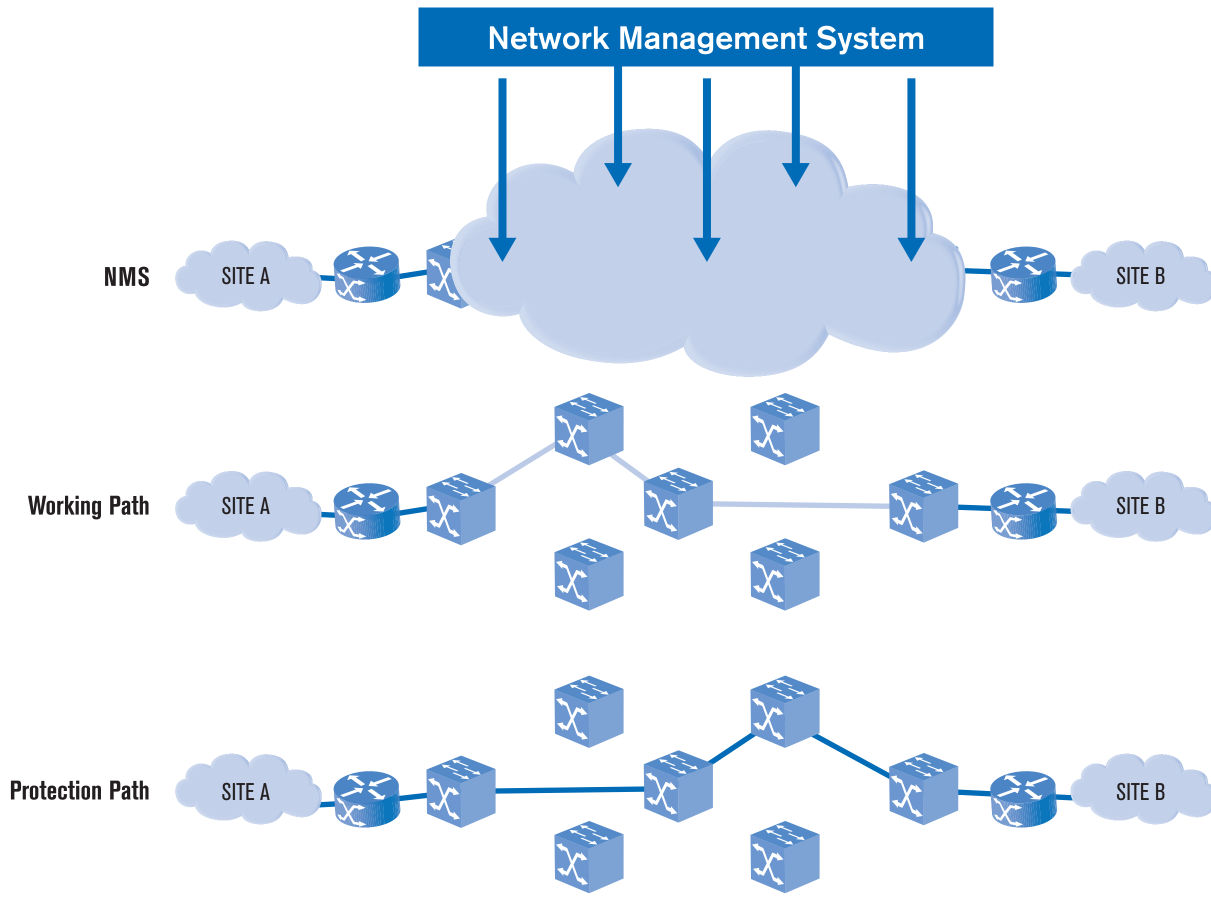

3. Network Resiliency

In legacy Ethernet, very few mechanisms existed to determine if a network link failed. This function was ensured by the spanning tree learning mechanism, which was automatically performed at specific intervals. In a situation where a link between two devices failed, the spanning tree learning process would automatically rediscover the network interruption and elect new routes based on its new calculations.

These mechanisms disrupt the flow of traffic until the learning and calculation process is completed, typically 3 to 5 seconds with RSTP. This also causes arbitrary selection of routes. The fastest link is not always selected, only the one with the lowest cost is, which may not be to the advantage of the customer.

In PBB-TE, since spanning tree is disabled and routes are configured by the network operators, network resiliency also becomes a configurable item. The network can contain at least two paths, a working path and a protection path. The working path consists of the main path taken by frames under normal operation while the protection path is the backup path taken by frames in case the main path is broken.

Path assignment is based on the B-VID assigned to the frames during their encapsulation at the edge switch. Therefore, the network operator must determine the working and protection VLANs as well as configure the routes that each VLAN must take on the network.

For PBB-TE, path issues are determined by the monitoring of 802.1ag connectivity fault management (CFM) continuity check messages, which should be received at specific intervals. In the event that a backbone or core device does not receive a CFM message after a specific interval, a failure in the link is assumed. Frames are then automatically forwarded using the protection path within 50 ms.

The PBB-TE recovery times are therefore within the sub-50 ms requirement of traditional transport networks while providing deterministic paths for frames and added value for carrier networks.

4. PBB Frame Format

The PBB frame format consists of a supplementary MAC layer that encapsulates another Ethernet frame.

B-DA	B-SA	Ether Type (0 x 88A8)	B-VID	Ether Type (0 x 88E7)	I-Tag	Customer Ethernet Frame	B-FCS
6	6	2	2	2	4	64 - 1510	4

The following fields are used:

- Backbone destination address (B-DA): represents the address of the destination edge switch of the PBB network
- Backbone source address (B-SA): represents the MAC address of the originating edge switch of the PBB network
- Ether type (1): indicates that a backbone VLAN is present in the PBB frame and has the value 0 x 88A8
- Backbone VLAN ID (B-VID): represents the VLAN ID applied to the frame and is used to ensure that frames take the proper path; this field is optional
- Ether type (2): indicates the content of the payload. In this case, 0 x 88E7 indicates Ethernet payload preceded by an I-Tag
- Instance Tag (I-Tag): is systematically applied to all PBB frames and contains QoS parameters as well as the 24 bits service ID (SID) field used to uniquely identify customers

Testing Needs

Testing a network usually involves ensuring connectivity and resiliency, as well as measuring performance. However, PBB-TE networks require further testing in order to prove connectivity and ensure the proper switching of data across the backbone switches. The following is a series of tests specially recommended for PBB-TE networks.

1. Connectivity across the Network (Edge-to-Edge)

Since PBB-TE requires the network operator to configure the path taken by frames across the network, this important test ensures that all nodes of the PBB network can be reached from any other node. More specifically, it ensures connectivity between edge switches, intrinsically measuring the route configured by the user.

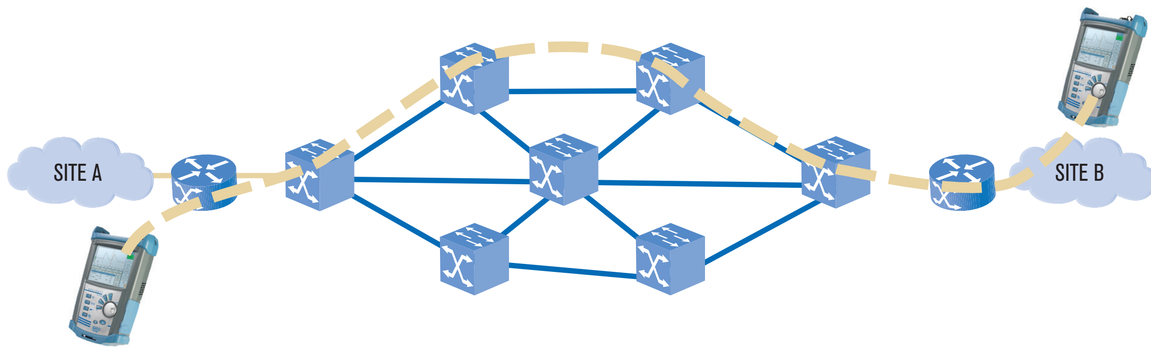

This test is typically performed in-service and involves a unit generating frames to an analyzer at the edge of the network and measuring statistics and errors.

2. VLAN Discrimination

As protection or working paths are used according to VLAN settings, it is important to ensure that the edge and core switches are able to discriminate between specific VLANs and forward them correctly. Recognized VLANs must be forwarded as defined by the customer, while unrecognized VLANs should be discarded.

VLAN discrimination testing ensures that frames can travel along the working and protection paths assigned to specific VLANs. This test should be executed by sending frames with valid VLAN tags along the working or protection paths to ensure they reach their destination. The test should be repeated using frames with unknown VLAN tags to ensure that they are discarded by the first network device they reach.

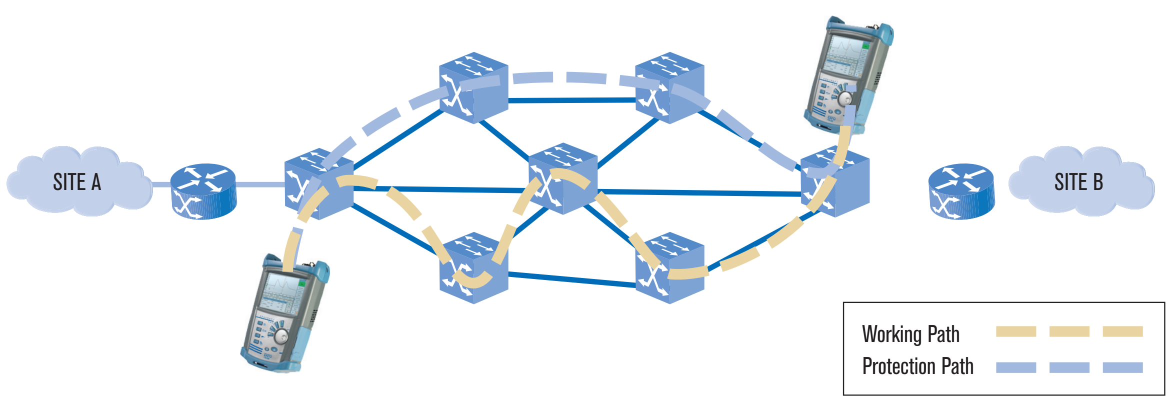

3. Protection Switching Assurance

This test ensures that edge or core switches react accordingly to the loss or lack of reception of continuity check messages or of messages with errors. A switch that does not receive these messages after a certain time should assume failure in the network and assign frames to the protection VLAN.

This test can consist of either suppressing the link between switches of a working path or measuring the service disruption time.

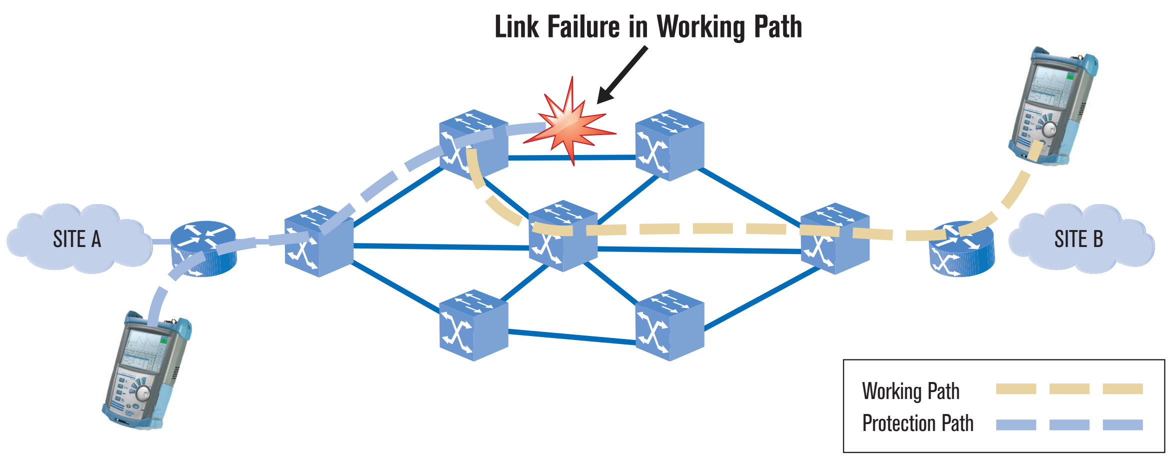

4. Performance Measurement

This key test should be mandatory for any carrier applying QoS to its forwarding. This type of test will typically consist of throughput and latency measurements based on different priority and S-IDs as each carrier assigns QoS. Some of the performance measurements that should be studied are:

- Throughput or the maximum bandwidth available per QoS parameter or per S-ID; these parameters define the bandwidth profile available for the flows.
- Latency measurement or the time it takes for frames to reach the destination; this important SLA parameter is essential for time-dependent applications such as real-time audio/video.

Performance tests should be carried out within the PBB network or from one customer premises to another customer premises. The RFC test between customer premises allows carriers to gather performance measurements as experienced by the subscriber.

Within the PBB network, testing should be executed over both protection and working VLAN paths to validate performance.

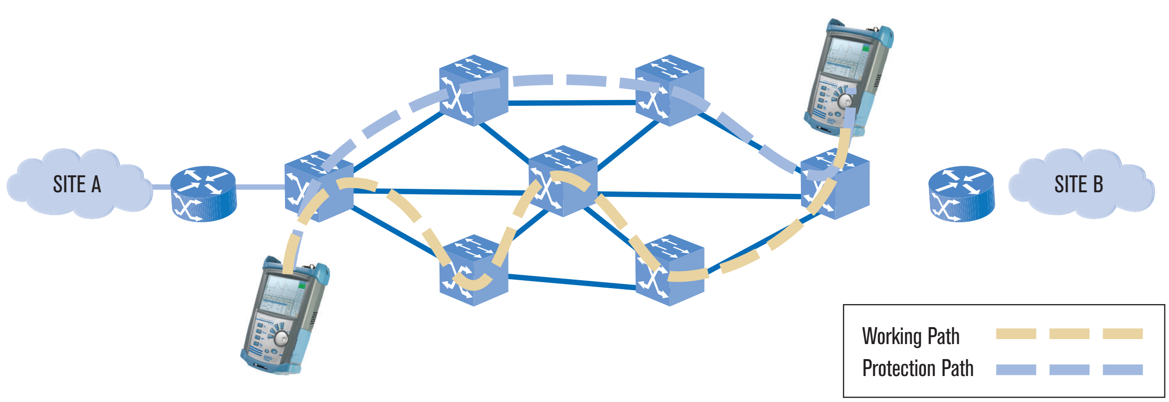

Conclusion

The legacy limitations of Ethernet can now be overcome via PBB-TE technology, bringing the highly sought cost-effectiveness and simplicity of the technology to core carrier networks. The deployment of PBB-TE provides carriers with the scalability, reliability and granularity of service they require while taking advantage of the wide acceptance and low operating costs of Ethernet.

Some of the key elements are:

	Legacy Ethernet	PBB-TE
Network Appliances	– Switches and bridges can only handle legacy Ethernet traffic – Switches need to "learn" all MAC addresses on the network, which can be extensive, depending on network complexity	– Core and edge bridges can handle both legacy and PBB traffic – Core and edge devices only need to "learn" the edge MAC addresses to forward frames
Network Control	– No control over the network	– Network management system provides full control to the customer
Link Error Detection and Recovery	– Spanning tree protocol is used to calculate routes; network is down during calculations – Recovery time is variable and can take seconds, depending on network complexity	– Packets are forwarded either on working or protection paths routes that are predefined by the network operator – Link error detected by 802.1ag CFM – Error recovery within 50 ms
QoS and Scalability	– Provided through VLANs – Limited to 4095 instances – Seven levels of priority	– Provided through I-Tag – Up to 16 million instances – Seven levels of priority

The deployment of PBB-TE, however, requires expanded testing scenarios that assess the performance and the reliability of the network to guarantee service to the customers. EXFO test solutions help carriers assess this performance.

EXFO Corporate Headquarters > 400 Godin Avenue, Quebec City (Quebec) G1M 2K2 CANADA | Tel.: +1 418 683-0211 | Fax: +1 418 683-2170 | info@EXFO.com

Toll-free: +1 800 663-3936 (USA and Canada) | www.EXFO.com

EXFO America	3701 Plano Parkway, Suite 160	Plano, TX 75075 USA	Tel.: +1 800 663-3936	Fax: +1 972 836-0164
EXFO Asia	151 Chin Swee Road, #03-29 Manhattan House	SINGAPORE 169876	Tel.: +65 6333 8241	Fax: +65 6333 8242
EXFO China	No. 88 Fuhua First Road Central Tower, Room 801, Futian District	Shenzhen 518048 P. R. CHINA	Tel.: +86 (755) 8203 2300	Fax: +86 (755) 8203 2306
	Beijing New Century Hotel Office Tower, Room 1754-1755 No. 6 Southern Capital Gym Road	Beijing 100044 P. R. CHINA	Tel.: +86 (10) 6849 2738	Fax: +86 (10) 6849 2662
EXFO Europe	Omega Enterprise Park, Electron Way	Chandlers Ford, Hampshire S053 4SE ENGLAND	Tel.: +44 2380 246810	Fax: +44 2380 246801
EXFO Service Assurance	285 Mill Road	Chelmsford, MA 01824 USA	Tel.: +1 978 367-5600	Fax: +1 978 367-5700

APPNOTE210.1AN

Printed in Canada 09/03

EXPERTISE REACHING OUT