PBB-TE Metro/Aggregation interoperability with MPLS Core

White Paper

Introduction

Pure-packet technologies are becoming more dominant in the evolving next generation networks and are adopted by many service providers (SPs) and private network operators. While the access network is dominated by Carrier Ethernet-based solutions, the networks' core often utilizes MPLS switches and IP routers to enable connectivity to IP-based services and to ensure service quality across the core network. With the migration of SPs from legacy SONET/SDH to packet transport, Ethernet has evolved to fit the needs of the metro/aggregation transport layer. Provider Bridging[1] (PB), a.k.a Q-in-Q, was introduced to allow separation between customer's VLANs and SP VLANs. However the number of customers supported on an SP network was still limited to approximately 4000 and the SP core switches were overloaded by the need to learn every customer's MAC address. To overcome these limitations, an alternative method known as MAC-in-MAC[2] was developed. With this new scheme, the SP network is only required to carry packets between ports of backbone edge bridges (BEBs) using a special 24-bit Instance Service ID (I-SID). This enables support of up to 16 million services in the SP network and eliminates the need for backbone core bridges (BCB) to learn customers' - MAC addresses (see Fig.1).

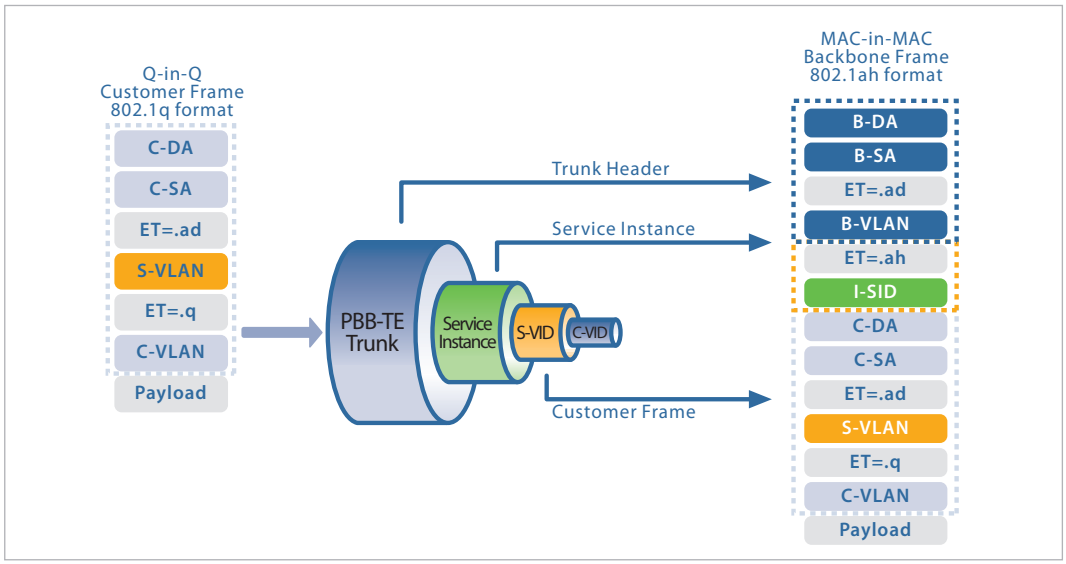

Fig. 1 – *Carrier Ethernet frame formats*

While Ethernet is gaining popularity at the metro/aggregation layer due to its simplicity and low cost, positioning it as an optimal evolution of SONET/SDH, IP/MPLS has become more dominant in the core of the network. The Ethernet domain must therefore be interoperable with the MPLS core to enable services be carried between customers and servers that reside on the core of the network and customers that are connected to different Ethernet domains (see Fig. 2).
Multiple options exist for such interoperability, including:

- UNI Interconnect (based on S-VLAN)
- NNI Interconnect (based on B-VLAN)
- UNI Interworking Function (based on I-SID/MPLS-label translation)

1. *IEEE 802.1ad*
2. *PBB – IEEE 802.1ah and PBB-TE – IEEE 802.1Qay*

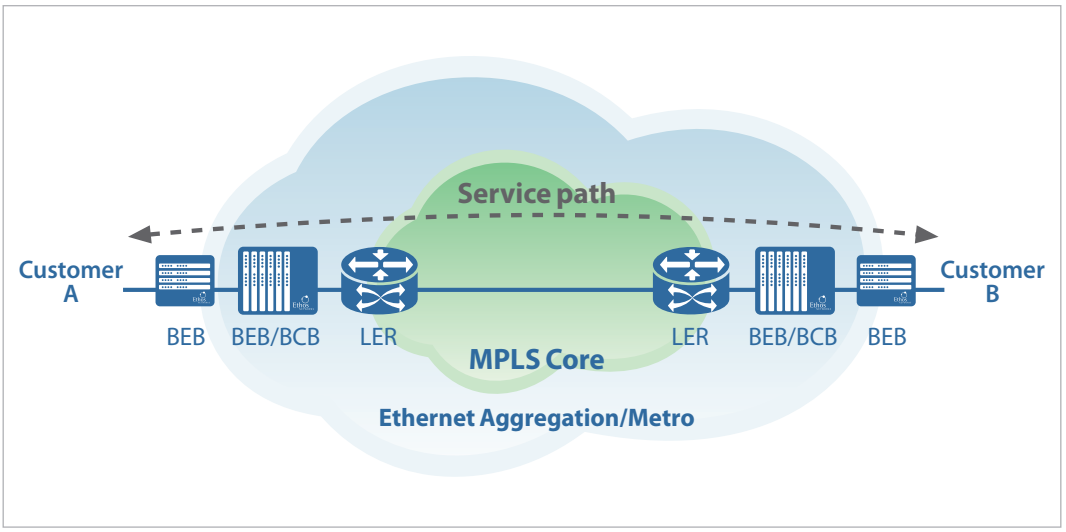

Fig. 2 – *Hierarchical Ethernet/MPLS network*

UNI Interconnect

With UNI Interconnect, the Ethernet and the MPLS networks are connected using a User-Network Interface (UNI). As illustrated in Fig. 3, traffic arrives from the Ethernet domain, the BEB strips down the backbone header, leaving a standard QiQ (802.1ad) header. The MPLS Label Edge Router (LER) receives this standard packet, pushes an MPLS label onto it and then forwards it into the appropriate Label Switched Path (LSP). At the other edge of the MPLS domain, the appropriate LER pops off the MPLS header and hands off a standard QiQ frame to the Ethernet BEB.

This straightforward approach uses the lowest common interface and is therefore fully interoperable with all standard MPLS routers. The combination of PBB/PBB-TE and MPLS also overcomes the scalability issues of QiQ, i.e. the 4000 VLANs limitation is now associated with each MPLS LER port, rather than with the whole access /aggregation network.

This allows (4000 x Number_of_MPLS_ports) service instances in the network, and the number-of-customers limitation is practically eliminated.

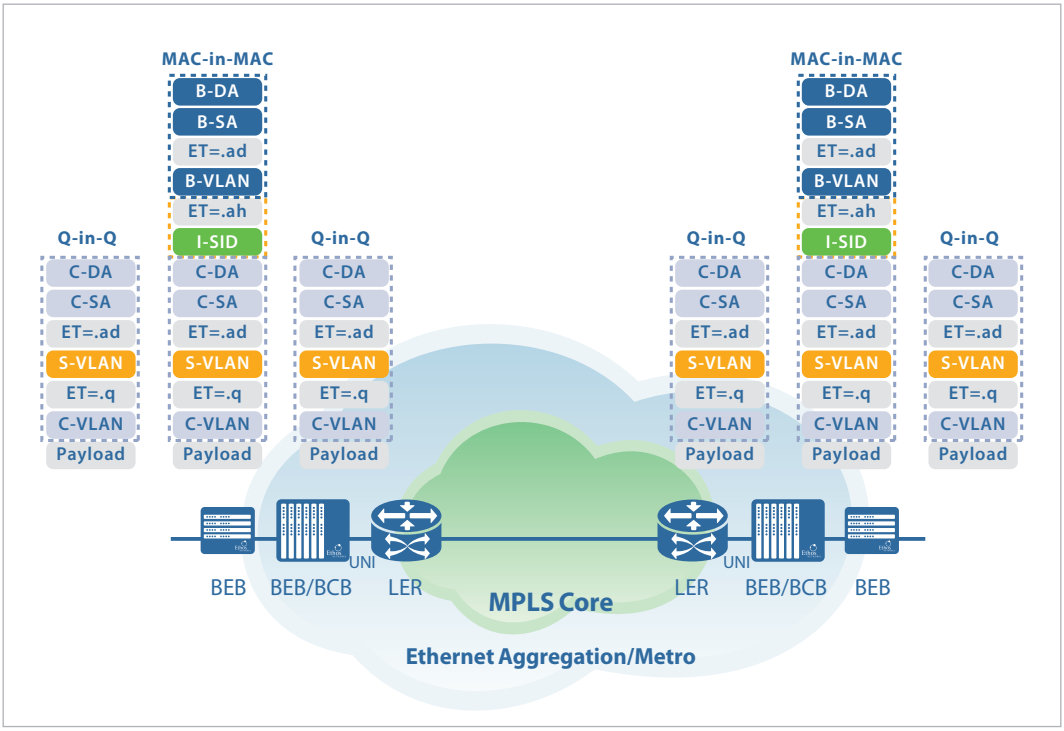

Fig. 3 – *UNI Interconnect*

NNI Interconnect

A more sophisticated approach for interconnecting the domains utilizes Ethernet Network-to-Network Interfaces (NNI) to connect to the MPLS network (Fig. 4). With this approach, the Ethernet BCB treats the whole LSP across the MPLS domain as a single link so it hands off to the MPLS LER a full backbone frame (802.1ah-compliant). The LER analyzes the frame and picks up its Backbone MAC (B-MAC) addresses and Backbone VLAN ID (B-VID) as if they were the customer MAC (C-MAC) and the Service VLAN ID (S-VID) of an 802.1ad frame.

The advantage of using NNI Interconnect is that it can enhance network resiliency by using redundant tunnels between the two Ethernet BCBs. Two tunnels can be defined as two protected links, each utilizing a different LSP, enabling sub-50ms protection switching on the Ethernet level in case of an LSP failure. The two BCBs in both edges of the MPLS domain will exchange Connectivity Check Messages (CCM) to ensure link health, and will initiate protection switching in case CCM messages are not received within a predefined period. (Note: such protection is not available on a UNI interface and therefore cannot be supported on the UNI Interconnect scheme).

The NNI Interconnect has an additional advantage of enabling end-to-end service performance management. Since the Ethernet BCBs at the edge of the MPLS domain treat the MPLS LSP as a simple link connecting them, the end-to-end service can be viewed as a single management domain (MD) and service OAM can be implemented from the ingress customer interface to the egress interface on the other side of the network. This is an improvement from the UNI Interconnect scheme where the UNI interface at the BEB terminated the Ethernet MD before handing off the packets to the MPLS LER.

It should be noted though that as a prerequisite for the NNI Interconnect scheme to work, the MPLS LER must be configured to ignore Customer VLAN ID (C-VID). Should the LER look for the C-VID and its associated EtherType (802.1q), it would not find any meaningful value and drop the packet. Therefore, if the LER cannot be configured to ignore C-VID then the NNI interconnect scheme should not be used.

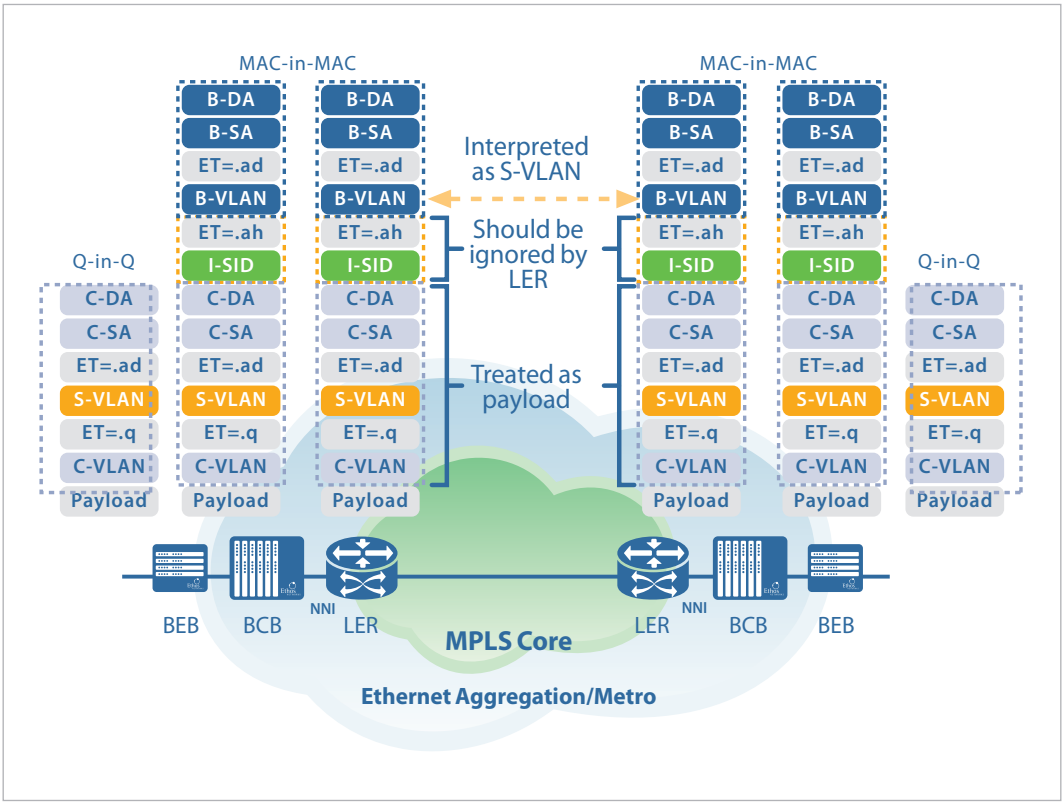

Fig 4 - *NNI Interconnect*

UNI IWF

Using UNI Interworking Function (IWF) requires bidirectional translation between Ethernet tunnels to MPLS LSPs. This function can be implemented in the Ethernet BCB as part of the UNI interface (as illustrated in Fig. 5), or in the MPLS LER. In this scenario, the I-SID is translated to an MPLS label at the edge of the Ethernet domain, and vice versa at the other side of the MPLS domain.

Service protection and performance management shall be implemented in such a case per domain. The Ethernet devices shall ensure service continuity and performance management within the Ethernet domain, while the MPLS devices shall perform the same within the MPLS domain.

This implementation has the additional benefit over UNI Interworking of supporting up to one million services at each MPLS port (limited by the MPLS 20-bit label). However it requires either the Ethernet BEB or the MPLS LER to support this functionality.

In case the IWF is implemented in the BEB, Label Distribution Protocol (LDP) should be used, to coordinate the mapping of Ethernet tunnels and MPLS paths and manage protection switching on this interface.

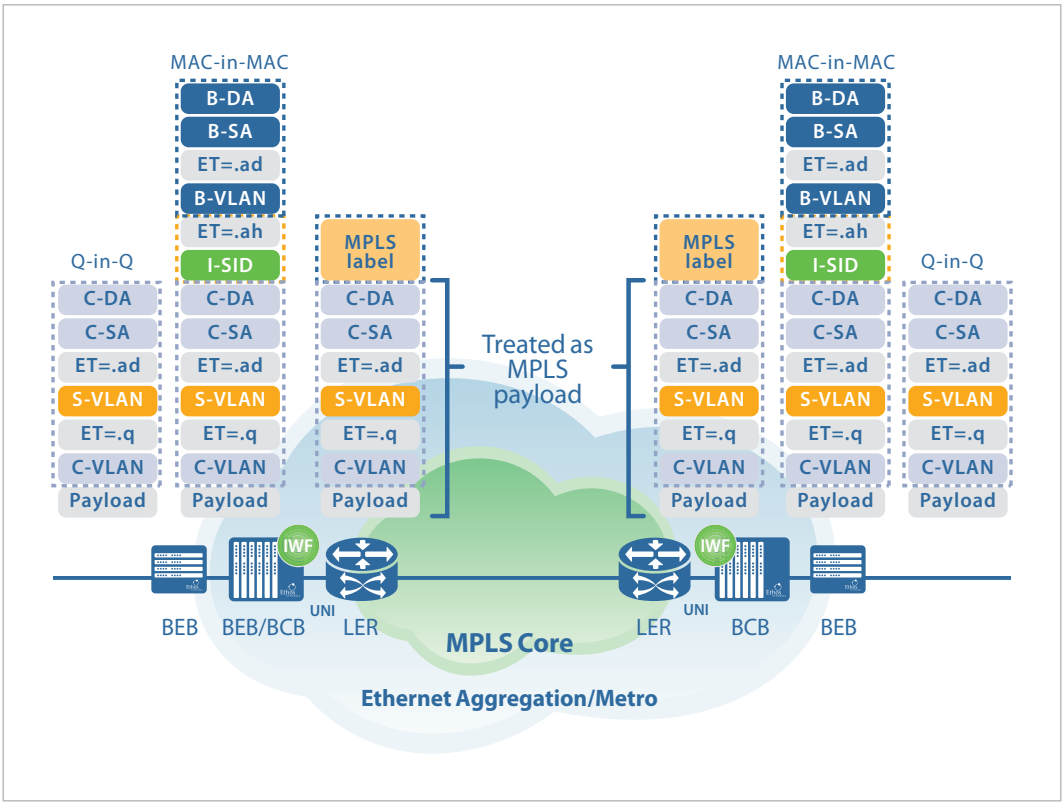

Fig. 5 – *NNI IWF (with IWF function implemented in Ethernet BCB)*

	# of service instances	Service protection	Performance management	Prerequisites
UNI interconnect	4,000 per port	None	None	None
NNI interconnect	4,000 per port	Ethernet protection	Ethernet PM	LER to ignore C-VLAN
UNI IWF	1 million per MPLS network	Ethernet/ MPLS protection	Ethernet/MPLS PM	IWF in BEB or LER

Summary

With Ethernet proliferation in the access and transport/aggregation layers and MPLS dominance at the core, interoperability and proper handoff of traffic are of the essence. Multiple alternatives exist to connect between the two domains:

- UNI Interconnection - simple and ubiquitous, but with some limitations on protection and performance management.
- NNI Interconnection - offers more capabilities, including protection and performance management.
- UNI Interworking Function - MPLS oriented interconnection, but requires development either at the MPLS LER or at the Ethernet BEB (availability should be verified per product).

For more information about Ethos Networks, please visit: www.ethos-networks.com