Bridging Solution for the MAN: Address Separation

October 2002

Marc Holness, Nortel Networks

Problem Context

- An Ethernet Service Provider supplies Ethernet Layer 2 services among customer sites
- IETF's Provider Provisioned Virtual Private Networks (PPVPN) WG is defining L2-VPNs
  - PPVPN drafts do not assume that bridges, as defined by 802.1, are essential to providing the service
  - They assume that devices which learn MAC addresses and forward packets based on that learned information are essential, but these are not bridges
- How can bridges supply these services?

Objective

- Proposal of an Ethernet bridging solution for the MAN
  - Carrier address space separation
- Solution alignment with 802 and 802.1 Architecture
- Solution interworking with MPLS/IP/EoS network

Carrier Grade Ethernet: Part I - Address Separation

Value Proposition

- Without address separation between customer network and carrier network, interior bridges (i.e., MAC learning devices) of carrier network need to scale with customer address space
  - Non scalable solution
- Address separation allows interior bridges of carrier network to scale with own network
  - More manageable network
  - More predictable scaling attributes
- Service Provider network resources are not tied to the customer's topology

Solution Overview

- Introduce Hierarchy Bridges at the edge of the carrier network, and use Transparent Bridges* as interior network MAC learning devices
- Hierarchy bridges
  - Located at ingress (access) and egress of carrier network
  - Provides address separation between customer address space and carrier address space
  - Requires MAC learning of both customer and service provider address space

* Maximum frame size supported needs to be extended to be greater than 1522 octets. Suggest 1600 octets?

Bridged System Reference Model

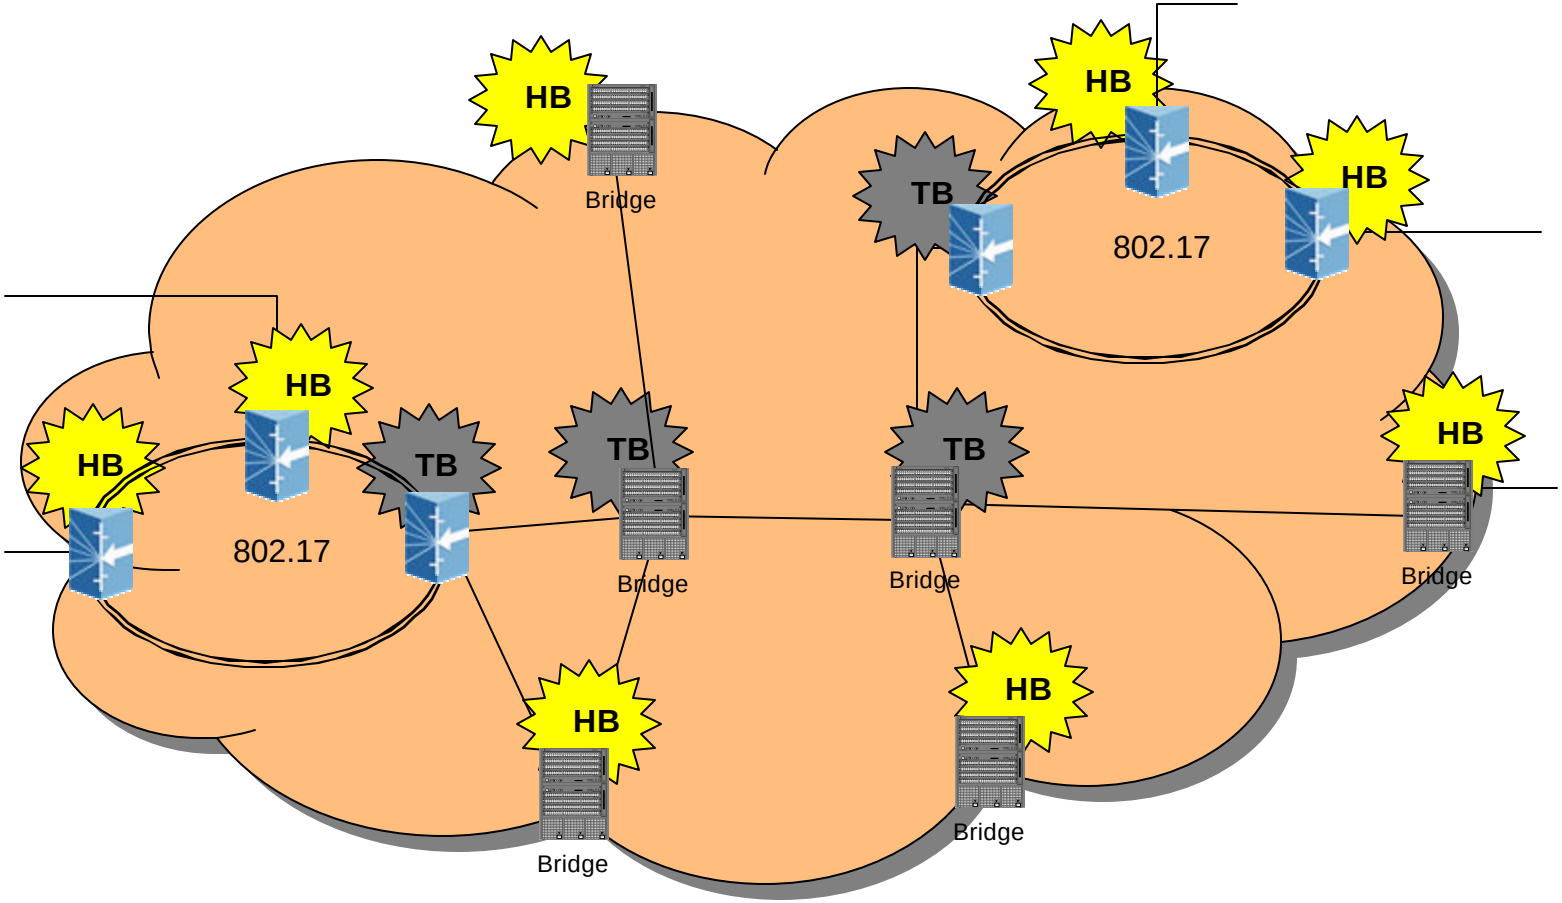

- **NOTE: Hierarchy Bridge (HB) have FDB that scale with the host MAC address space. The Transparent Bridge (TB) have FDB that scale with the Network MAC address space. Consequently, the core of the Network scales (and is 802.1D compliant).**
- **NOTE: Value proposition only realized in a (contained) Network when all edge nodes are HB conformant, and interior nodes are TB (802.1D/Q) conformant.**

Solution Scaling Attributes

	Edge/Access	Interior
802.1D (TB)	• Bridge Relay FDB size scales with number of host MACs serviced by the Network ***O(h_MAC)***	• Bridge Relay FDB size scales with number of host MACs serviced by the Network ***O(h_MAC)***
Hierarchy Bridge (HB)	• Bridge Relay FDB size scales with the number of Carrier MAC address within Service Provider Network ***O(p_MAC)*** • Some sort of FDB is needed to associated host MACs with Service Provider MACs. Consequently, this FDB scales with number of host MACs. ***O(h_MAC)***	• Run Transparent Bridge* in the interior of the network. Consequently, the Bridge Relay FDB size scales with number of Carrier MACs. ***O(p_MAC)***

Assume: Range of h_MAC > p_MAC!!

Carrier Grade Ethernet: Address Space Separation and Hierarchy Bridge Details

802.1 Bridging Architecture

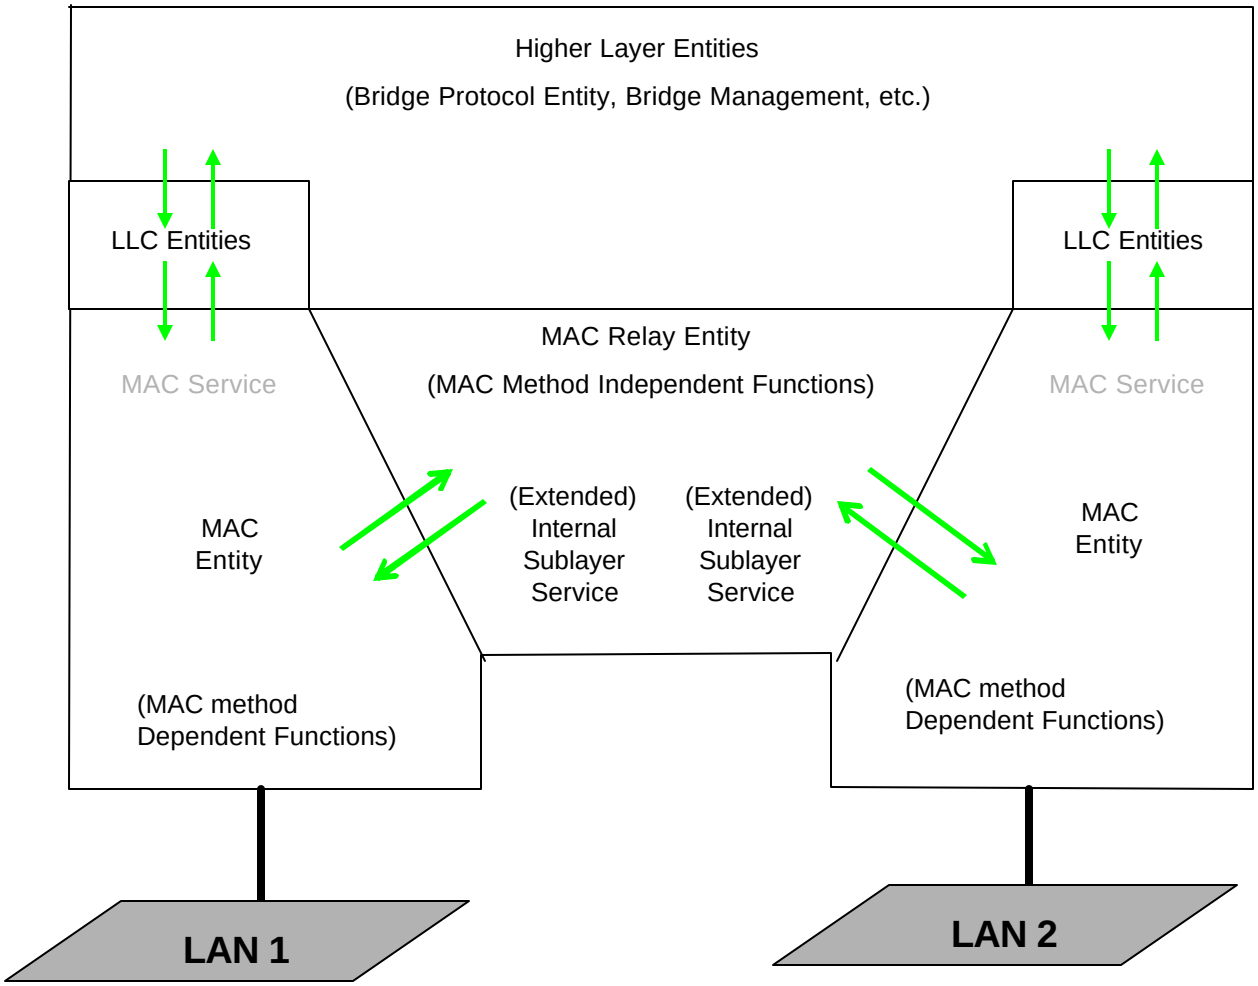

802.1D Bridging Reference

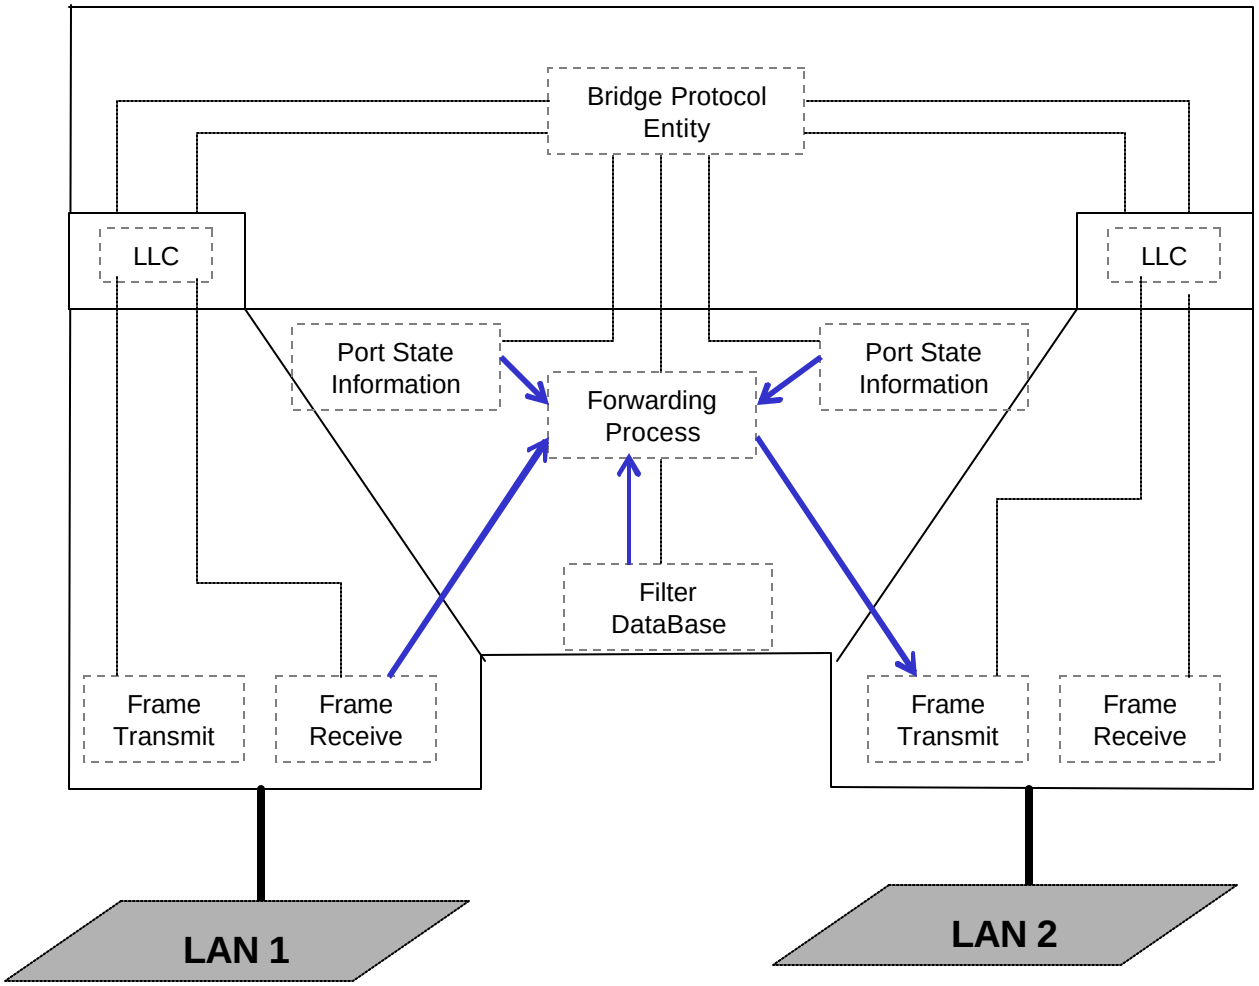

- Denotes Relaying MAC frames
- Denotes reception and transmission of BPDUs

Hierarchy Bridging Architecture

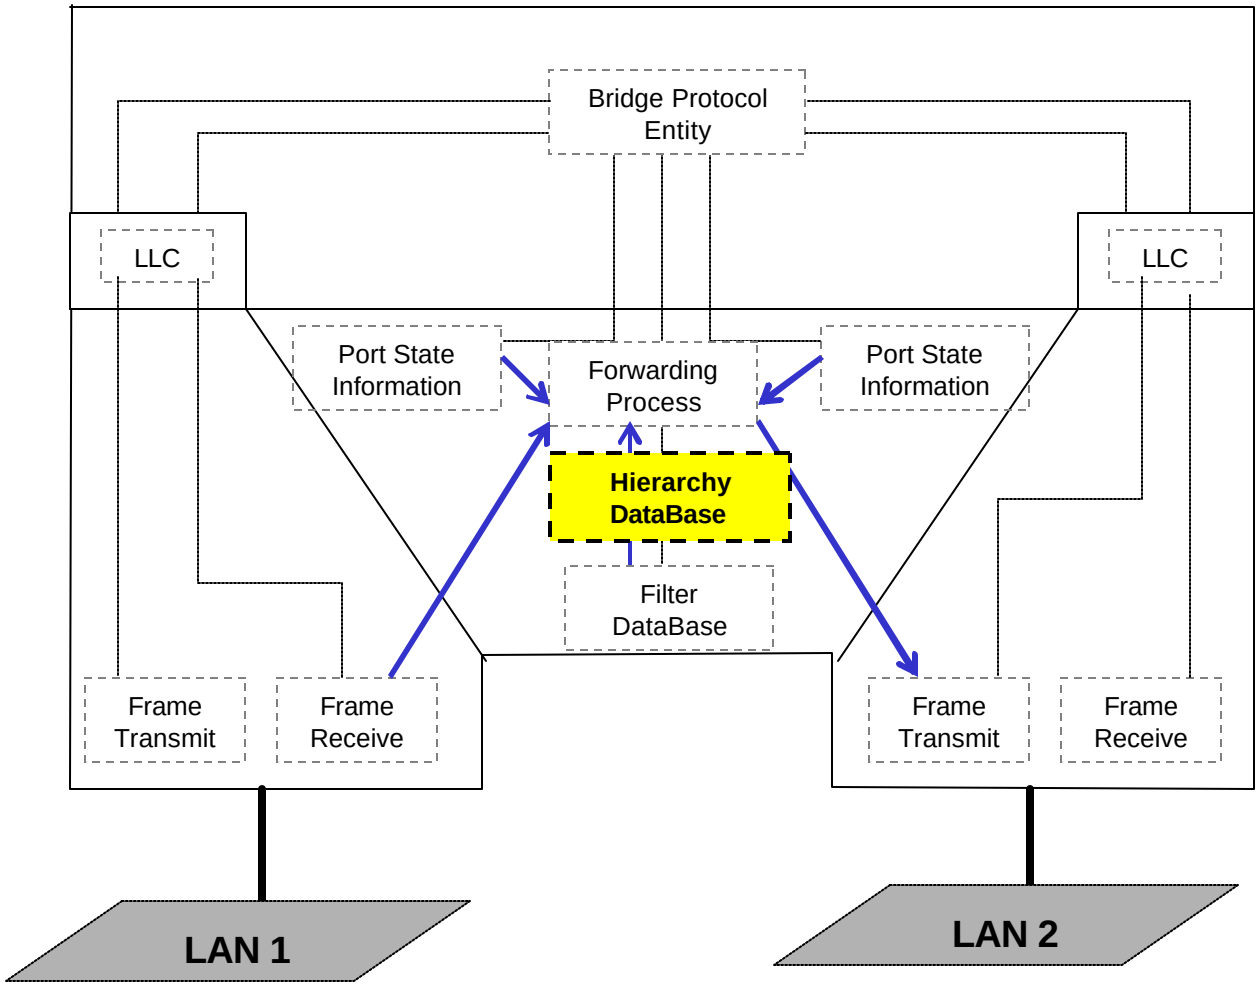

- Denotes Relaying MAC frames
- Denotes reception and transmission of BPDUs

Hierarchy Bridge Relay Entity

- Hierarchy Data Base is introduced. This FDB will associate MAC address in one user space with MAC address in another user space. For example, associate MAC address in the Service Provider space and those of the Host MAC addresses
- Utilizes Filtering Data Base as defined by 802.1D. The addresses stored in this FDB are those of the Customer space (e.g., Host address space)
- Relay Entity makes a decision to encapsulate packet with other address space or de-encapsulate the relayed packet

NOTE: The model of operation is in no way intended to constrain real implementations of a MAC Distributed Bridge. These may adopt any internal model of operation compatible with the externally visible behavior that this proposal specifies. Conformance of equipment to this specification is purely in respect of observable protocol.

Hierarchy Bridging Internal Sublayer Service

- The Hierarchy Bridge ISS uses the same ISS as defined by 802.1D/Q

HB-ISS Parameter Mapping

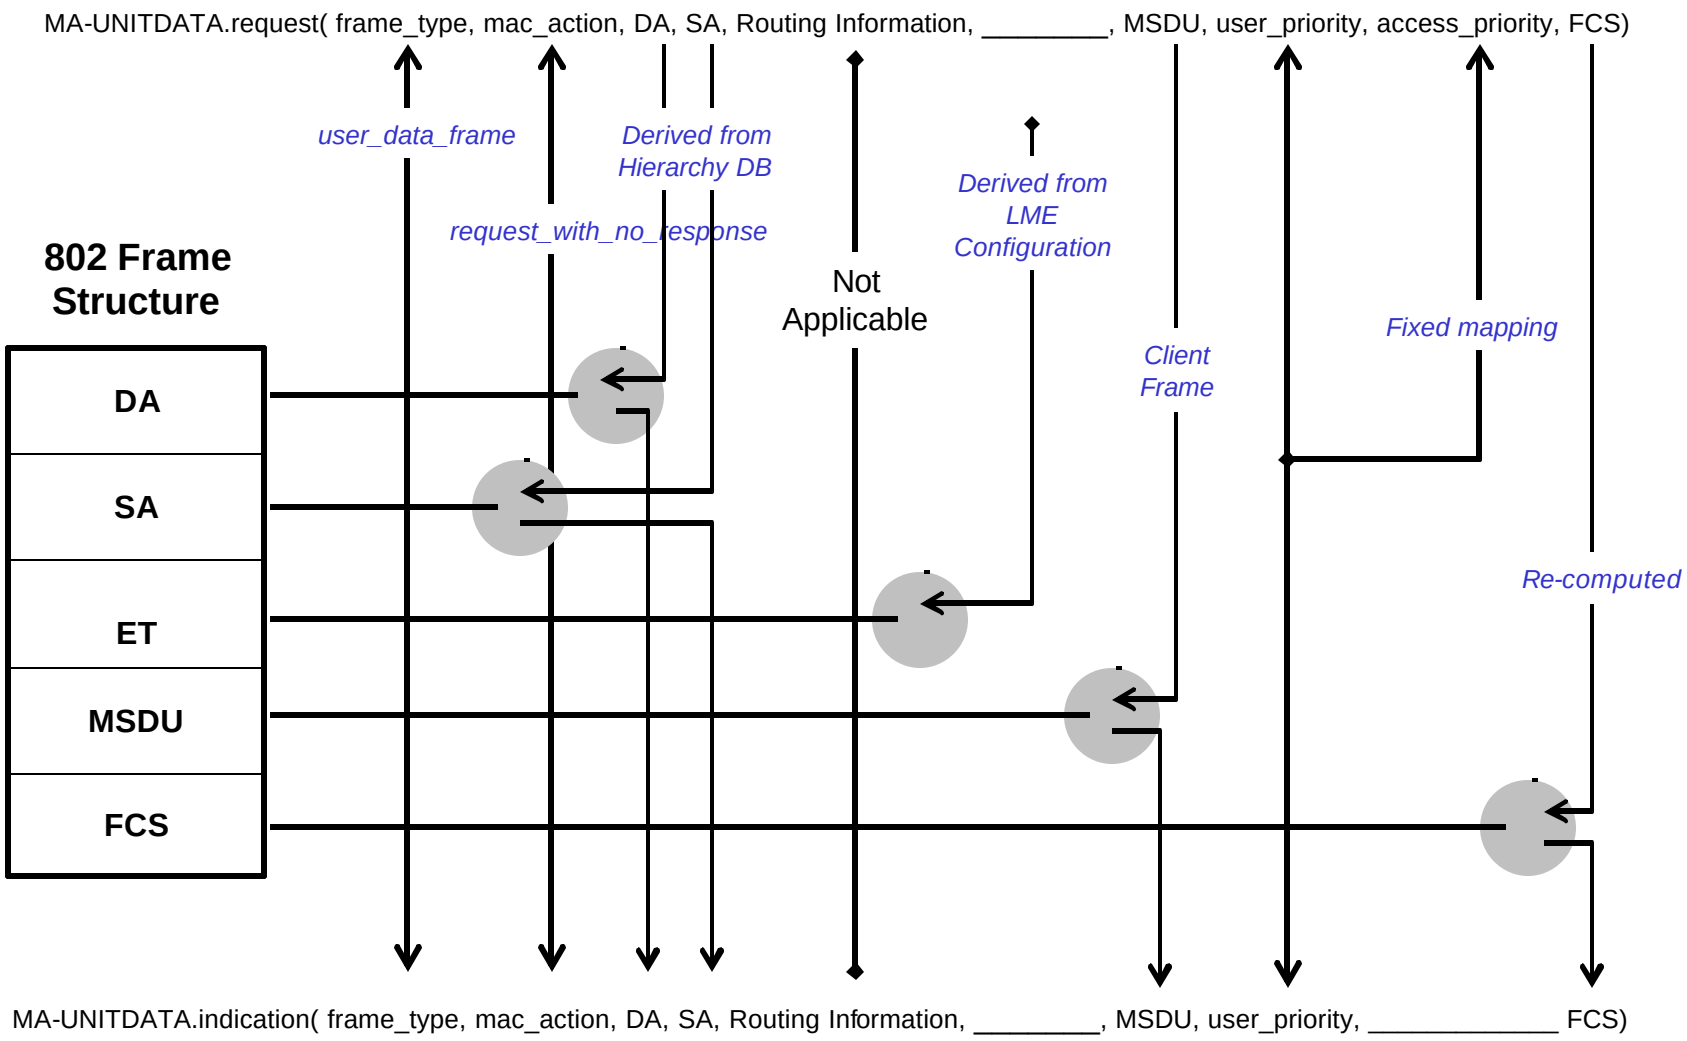

802 Frame Format

- Use the EtherType (ET) field to indicate that a client frame payload
  - Additional ET value is being proposed
  - New value that is being considered is `EtherFrame`

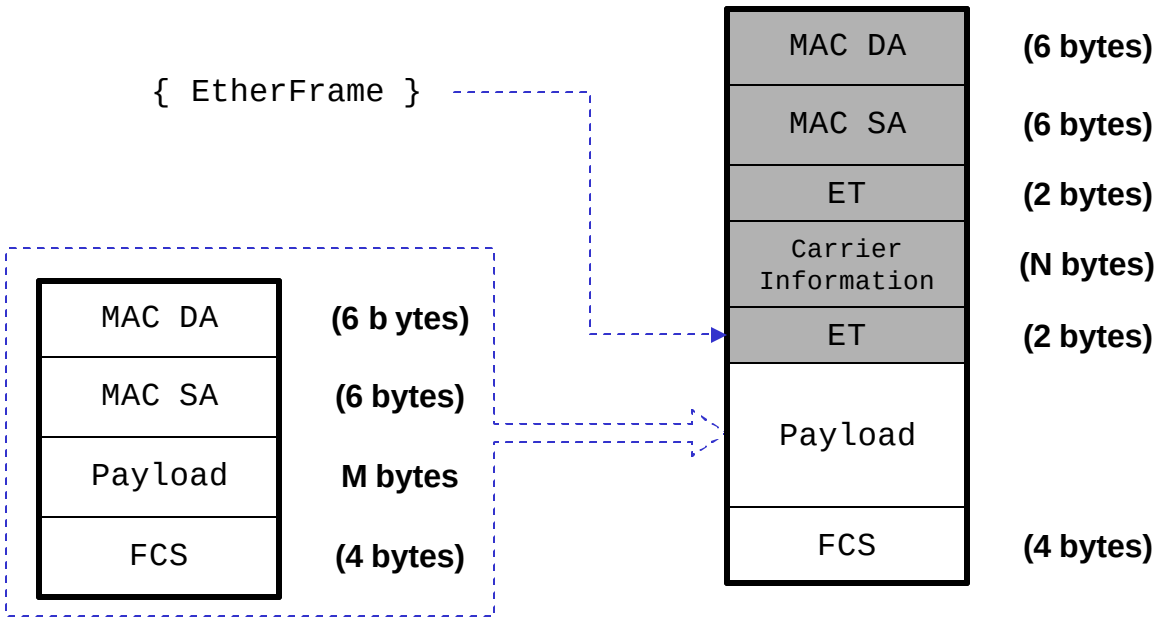

Observations

- Hierarchy bridge is a simple extension to 802 defined bridges
- No changes required to existing service interfaces (e.g., ISS/E-ISS)
- Scaling property of carrier network is achieved

Back Up Charts

Example: High Level Walk-Thru

Station S1 and S3 are configured to be Distributed Bridges. Station S2 is configured to be a Transparent Bridge (802.1D). Assume no DB learning has occurred.

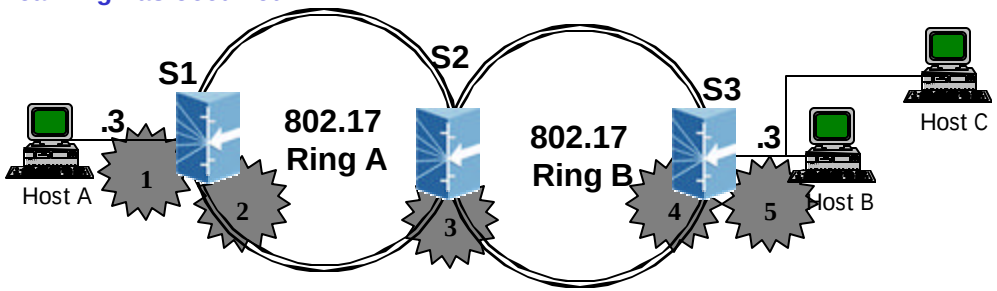

Step	Description	Packet Format	FDB	Hierarchy DB
1	HostA sends packet destined to HostB	B | A	N/A	N/A
2	RPR S1 receives packet and Bridges/Relays onto RingA.	FF | S1 | B | A	*[A, p1]*	*[S1, A]*
3	RPR S2 receives packet and process as a TB.	FF | S1 | B | A	*[S1, PortA]*	N/A
4	RPR S3 receives packet.	FF | S1 | B | A	-	*[S1, A]*
5	HostB receives packet.	B | A	N/A	N/A

Example: High Level Walk-Thru

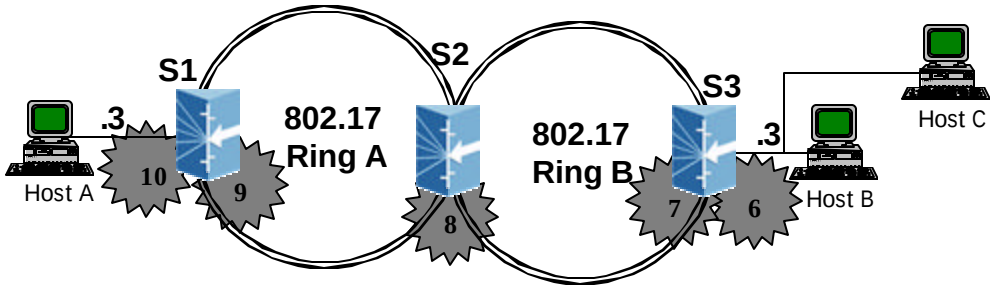

Step	Description	Packet Format	FDB	Hierarchy DB
6	HostB responds to HostA.	A | B	N/A	N/A
7	RPR S3 receives packet and Bridges/Relays onto RingB.	S1 | S3 | A | B	*[B, p3]*	[S1, A] *[S3, B]*
8	RPR S2 receives packet and process as a TB.	S1 | S3 | A | B	[S3, PortB] *[S1, PortA]*	N/A
9	RPR S1 receives packet.	S1 | S3 | A | B	[A, p1]	[S3, B] *[S1, A]*
10	HostB receives packet.	B | A	N/A	N/A

Example: High Level Walk-Thru

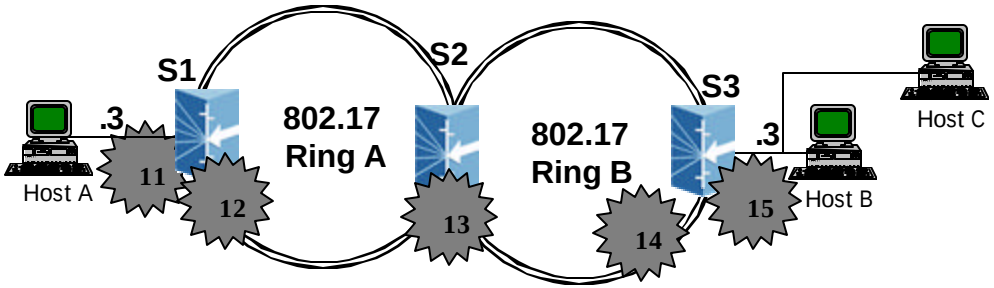

Step	Description	Packet Format	FDB	Hierarchy DB
11	HostA responds to HostB	B | A	N/A	N/A
12	RPR S1 receives packet and Bridges/Relays onto RingA.	S3 | S1 | B | A	[A, p1]	[S3, B] [S1, A]
13	RPR S2 receives packet and process as a TB.	S3 | S1 | B | A	[S3, PortB] [S1, PortA]	N/A
14	RPR S3 receives packet.	S3 | S1 | B | A	[B, p3]	[S1, A] [S3, B]
15	HostB receives packet.	B | A	N/A	N/A

Example: High Level Walk-Thru

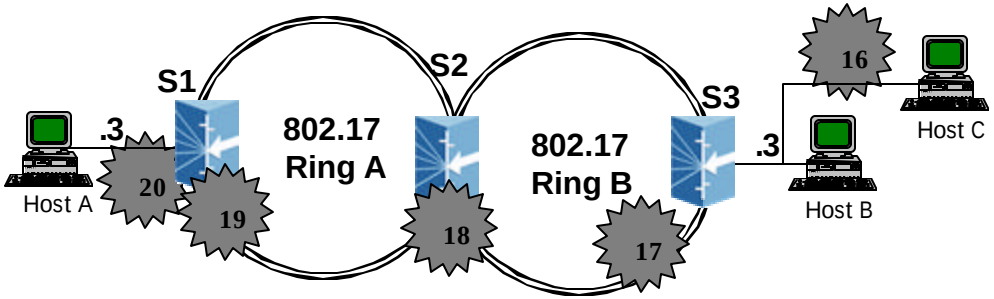

Step	Description	Packet Format	FDB	Hierarchy DB
16	HostC sends packet destined to HostA.	A | C	N/A	N/A
17	RPR S3 receives packet and Bridges/Relays onto RingA.	S1 | S3 | A | C	[B, p3] *[C, p3]*	[S1, A] [S3, B] *[S3, C]*
18	RPR S2 receives packet and process as a TB.	S1 | S3 | A | C	[S3, PortB] [S1, PortA]	N/A
19	RPR S1 receives packet.	S1 | S3 | A | C	[A, p1]	[S3, B] [S1, A] *[S3, C]*
20	HostA receives packet.	A | C	N/A	N/A

Hierarchy Bridge Pseudo-Code

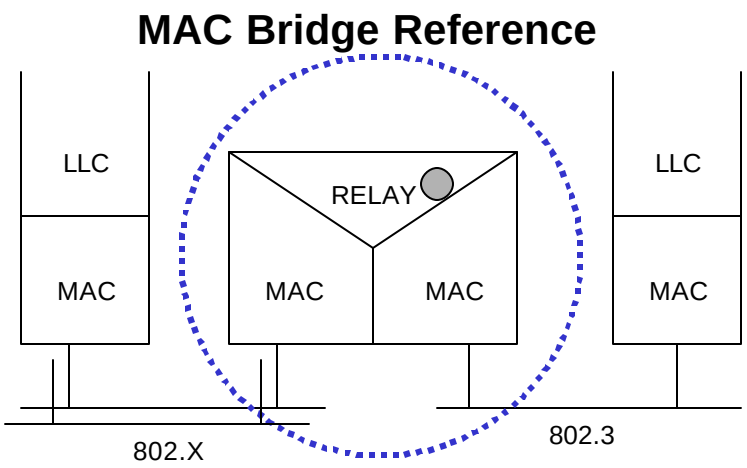


```
Case LME.Bridge in:
{ Hierarchy Bridge }:
   If ( M_UNITDATA.primitive == Indication )
      If ( ET == EtherFrame )
         /* Decapsulate Frame */
         Indication.MSDU( EtherEncap, SA ) -> Request.SA
         HierarcyDB.Update( Indication.SA, Request.SA, age )
         Indication.MSDU( EtherEncap, DA ) -> Request.DA
         Indication.MSDU( EtherEncap, Payload ) -> Request.MSDU
         Indication.MSDU ( EtherEncap, FCS ) -> Request.FCS
      Else
         /* Encapsulate Frame */
         FDB.Update( Indication.SA, port, age )
         HierarchyDB.Update( Device.Addr, Indication.SA, age )
         { EtherEncap } -> Request.ET
         Request.MSDU( Indication.SA, Indication.DA,
         Indication.MSDU, Indication.FCS )
         HierarcyDB.Index( Indication.DA ) -> Addr
         If ( addr == NIL )
            #FFFFFF -> Request.DA
         Else
            Addr -> Request.DA
         End If
         Device.Addr -> Request.SA
         Request.ReComputeFCS()
      End If
   End If
{ Transparent Bridge }:
   /* Perform 802.1D/Q operations */
End Case
```

IETF PPVPN Reference Model

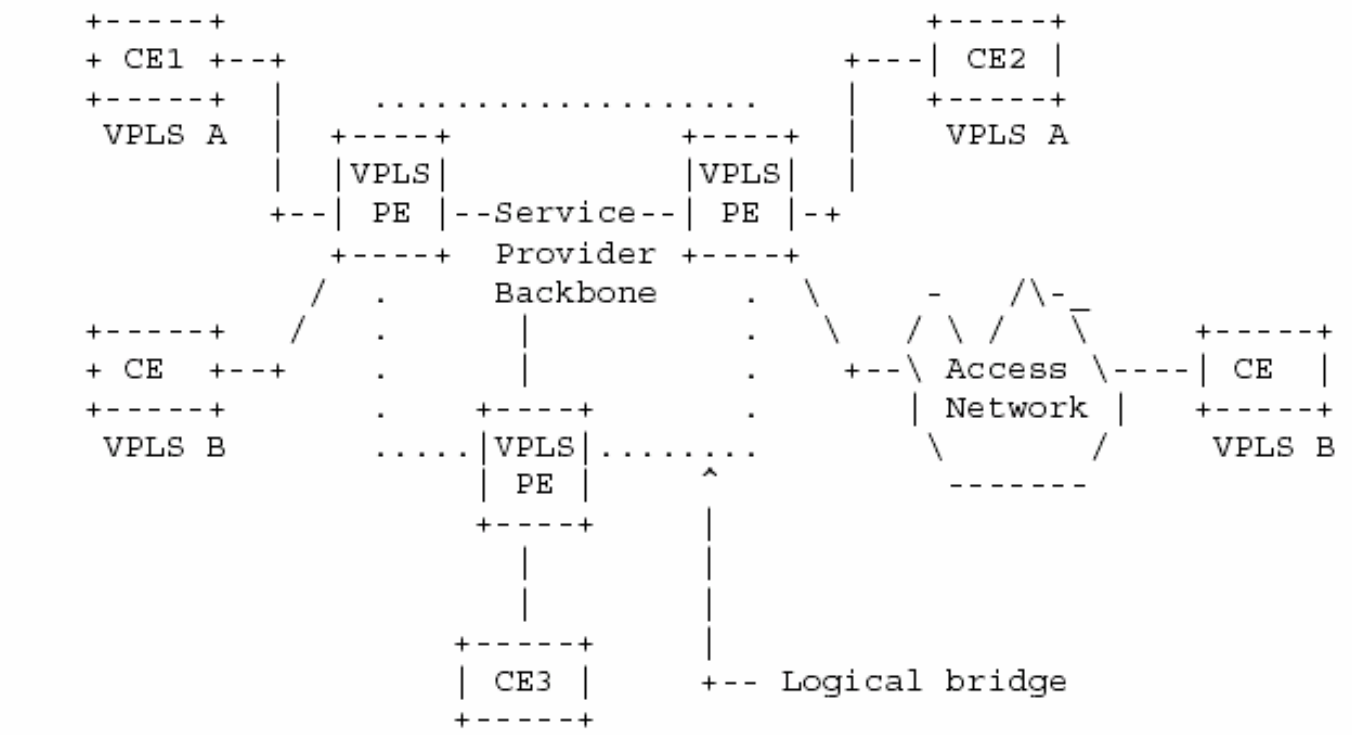

- Carrier domain is modeled as a logical (or distributed) bridge

Alternative IETF PPVPN Reference Model

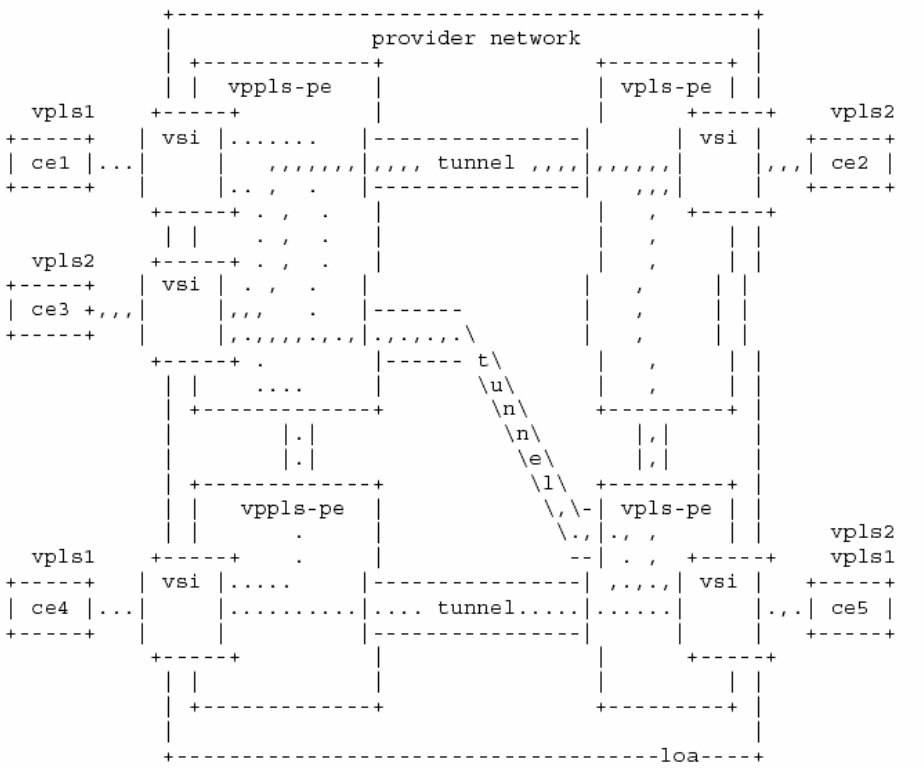

- Carrier domain is modeled as a collection of 802.1 bridges