

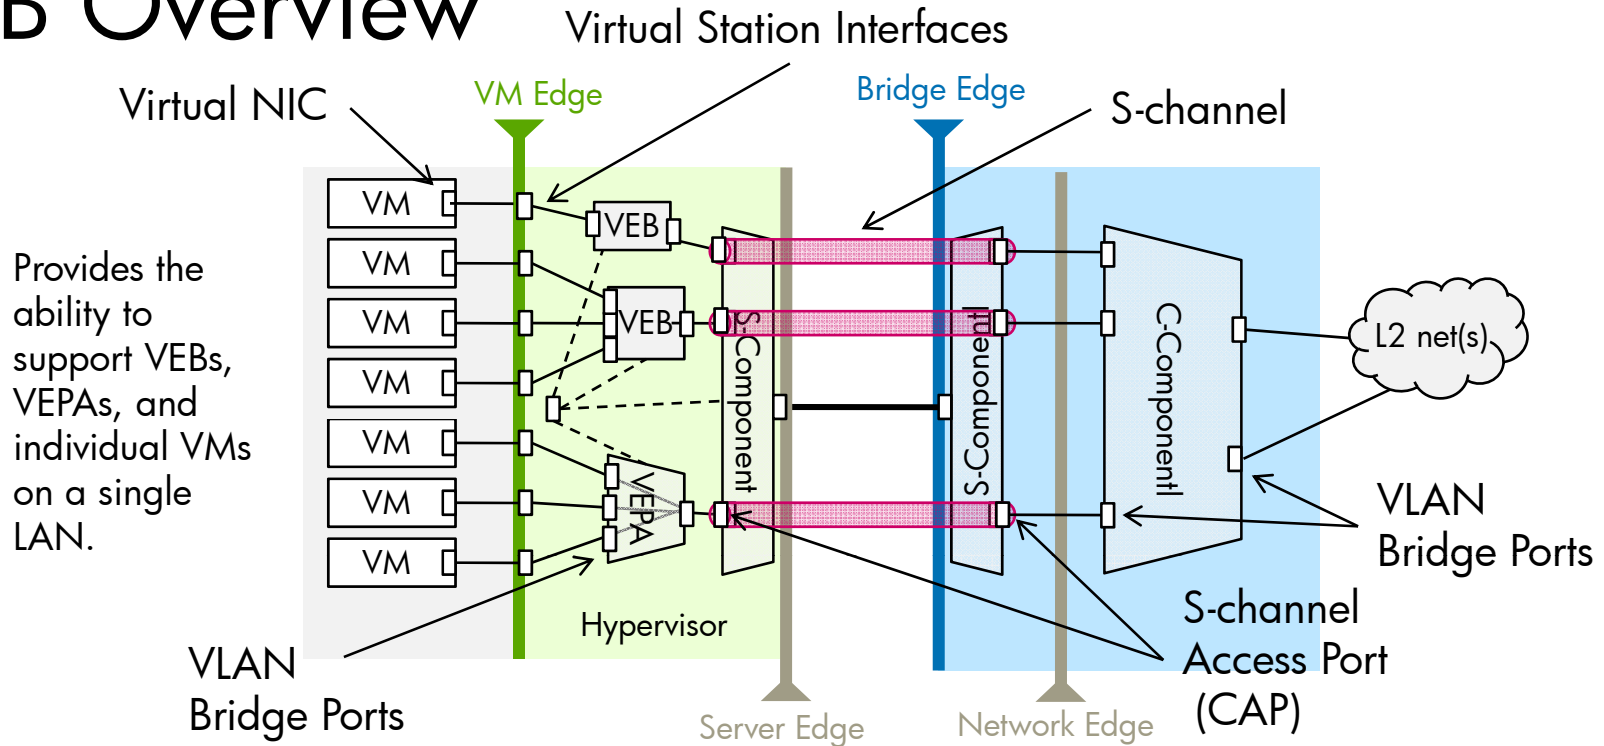
EVB Basic Architecture

V17

March 30, 2010

*Paul Bottorff (HP), Paul Condon(HP), Uri Elzur (Broadcom),
Chuck Hudson (HP), Daya Kamath (BNT), Vivek Kashyap(IBM),
Jay Kidambi(BNT), Mike Krause (HP) , Vijoy Pandey(BNT),
Joe Pelissier (Cisco), Renato Recio(IBM), Rakesh Sharma(IBM) ,
Pat Thaler(Broadcom)*

EVB Overview



•**Virtual NIC (vNIC)** the protocol layers used to implement the End Station functions of the attachment point. The vNIC sources an instance of 802.3 MAC service

•**Virtual Station Interface(VSI)** is an internal LAN which connects between a virtual NIC and a Bridge Port of a VEB or VEPA. A VSI carries a single MAC service instance. The term VSI is also used to reference a point where a Bridge Port of a VEB or VEPA attaches to a vNIC.

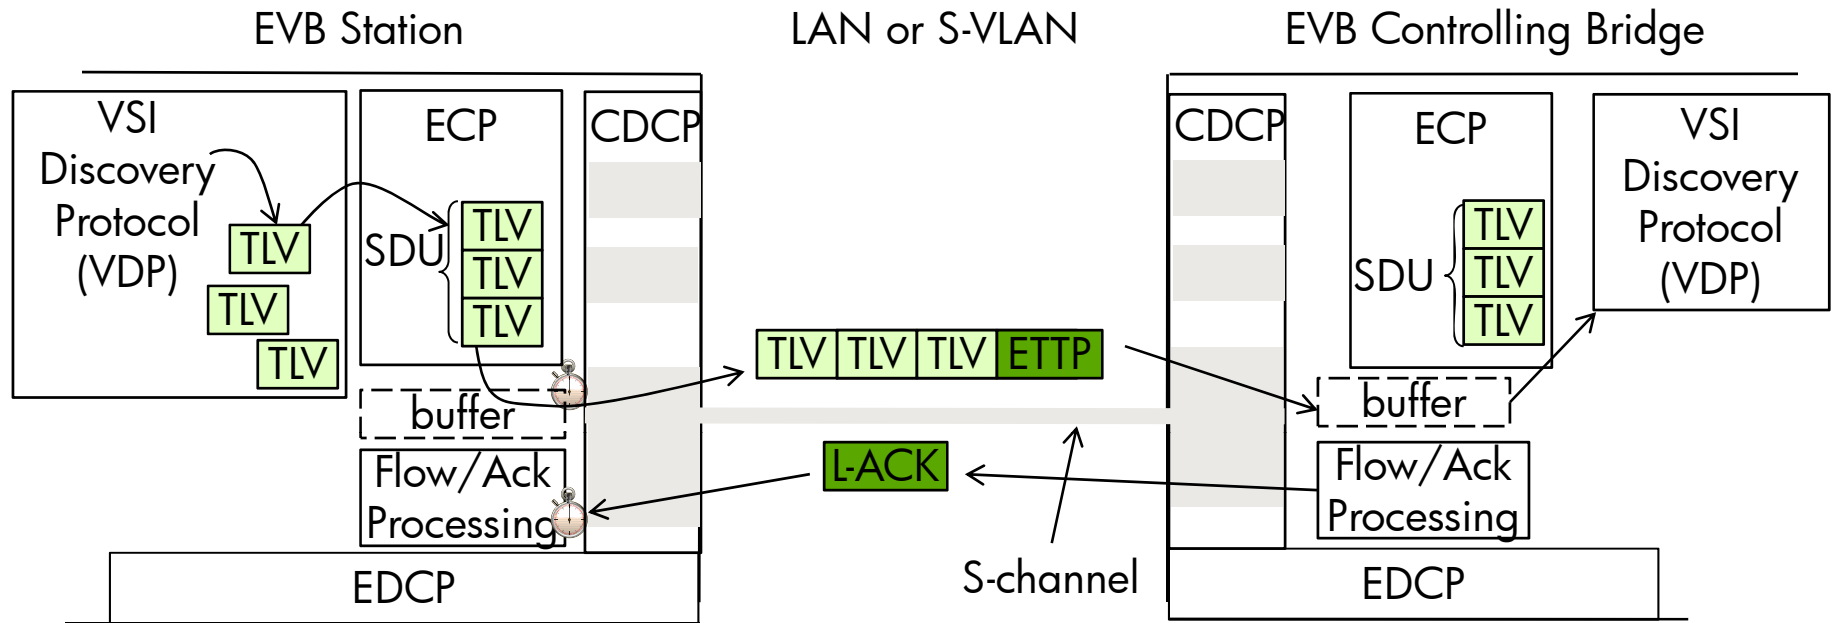
•**Virtual Ethernet Bridges (VEB):** a specialized bridge within the station used for bridging support between multiple virtual stations interfaces and an external network

•**Virtual Ethernet Port Aggregator (VEPA):** a specialized bridge within a station which operates in collaboration with an adjacent, external bridge to provide bridging support between multiple virtual stations interfaces and an external network

•**S-channel** is a Point-to-point S-VLANs which spans over Port-mapping S-VLAN components(802.1Qbc) or M-components(802.1Qbh) and between Provider Access Ports (PAPs). Only a single S-channel on a particular PAP

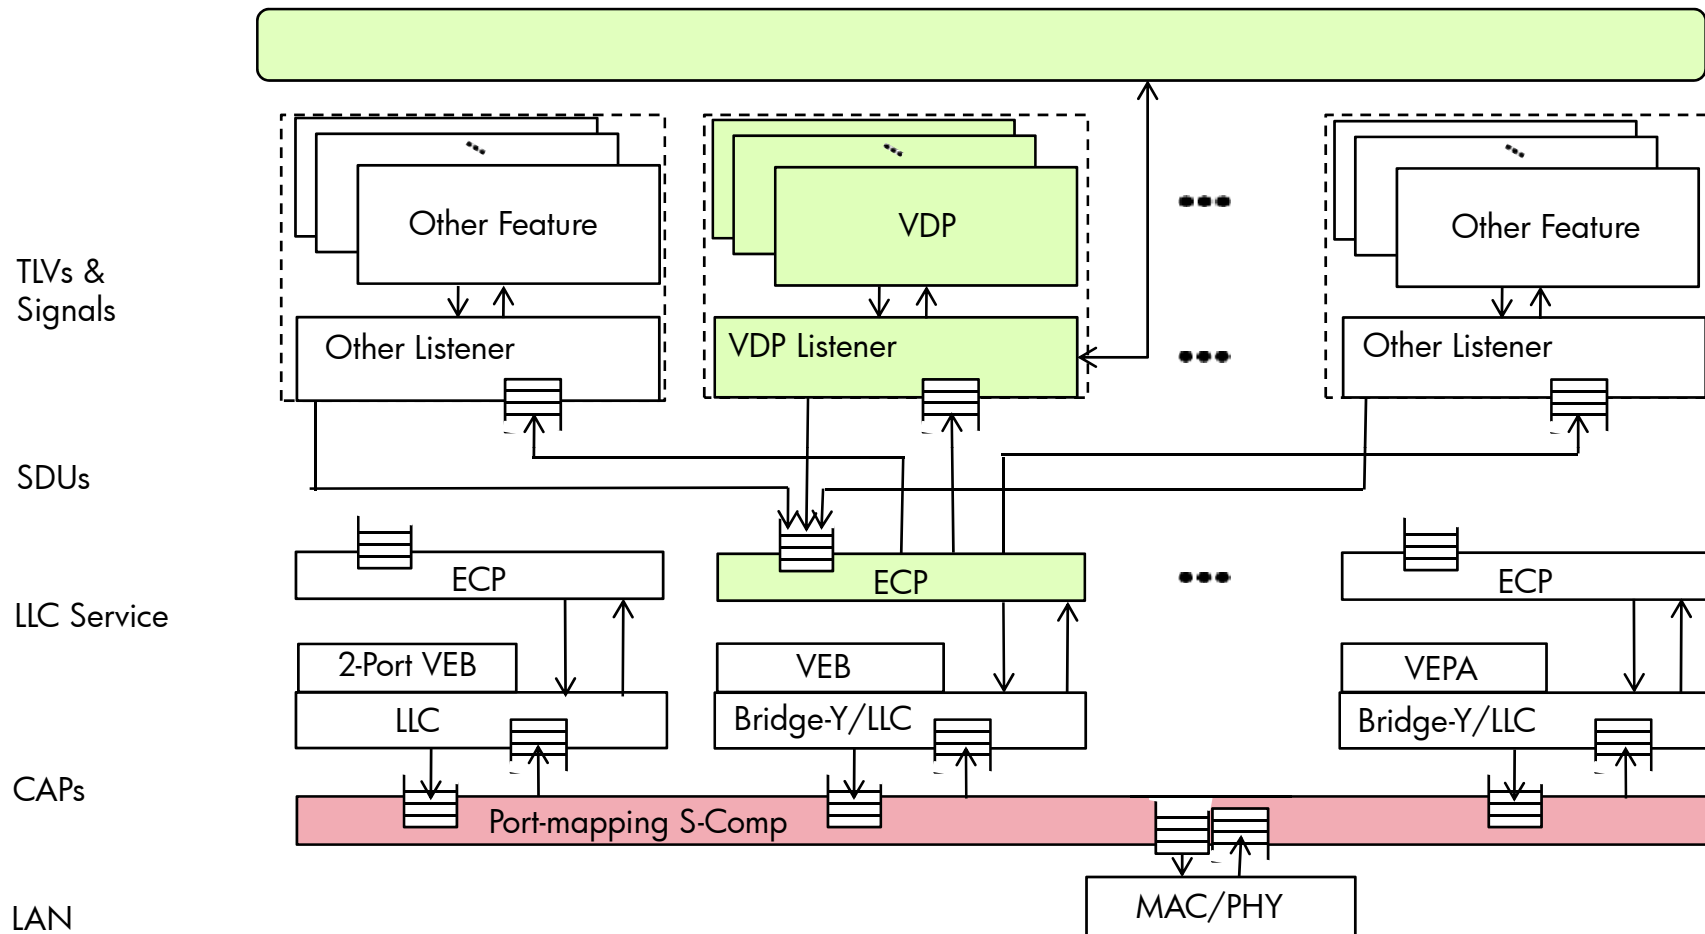
• **S-channel Access Port (CAP)** is the name for the un-S-tagging Bridge Port (Provider Access Port of a Port-mapping S-VLAN component (802.1Qbc)) that terminates an S-channel

EVB VSI Control Overview



- Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP) coordinates network resources for Virtual Machines (VMs) and packs and unpacks VDP TLVs into PDUs which are handed to ECP for delivery.
- Edge Control Protocol (ECP) provides reliable delivery of SDUs between the station and bridge and between the port extender and bridge
- Edge Discovery and Configuration Protocol (EDCP) is an LLDP based protocol used to discover and configure VEPA, ECP and VDP
- S-channel Discovery and Configuration Protocol (CDCP) is an LLDP based protocol used to configure the S-VLANs.

Example VSI Discovery Implementation



- There is one VDP per VSI. The VDP handles all VSI state transitions. It sends and exchanges TLVs and Signals from the VDP listener.
- There is one listener per ULP. The listener handles packing and unpacking TLVs into PDUs.
- There is one ECP per S-channel plus one for the default S-VLAN (not and S-channel). It provides reliable SDU delivery for SDUs which are placed in its transmit queue by the listeners. It also delivers received SDUs to the appropriate listener.

EVB Data and Control Entities - Station

Sublayer Service

Interfaces

LLC-SS →

ISS →

ECP-SS →

ISS →

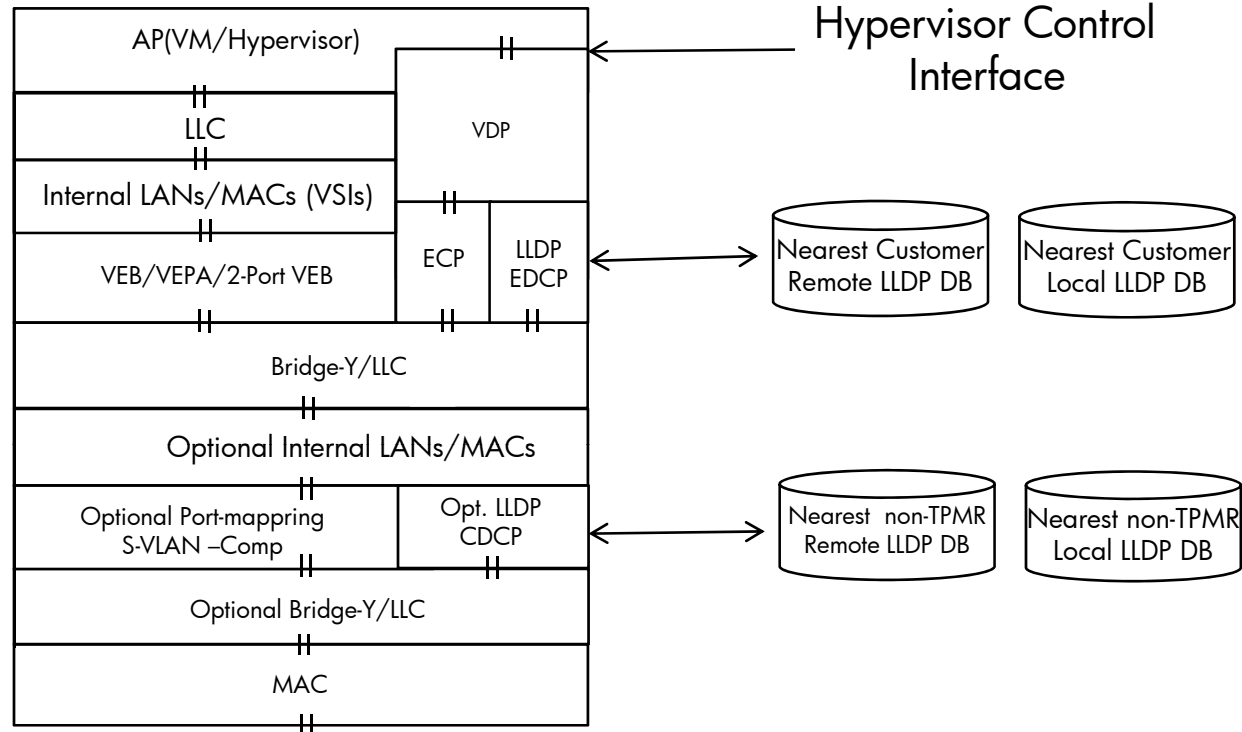
ISS/LLC-SS →

ISS →

ISS →

ISS/LLC-SS →

ISS →

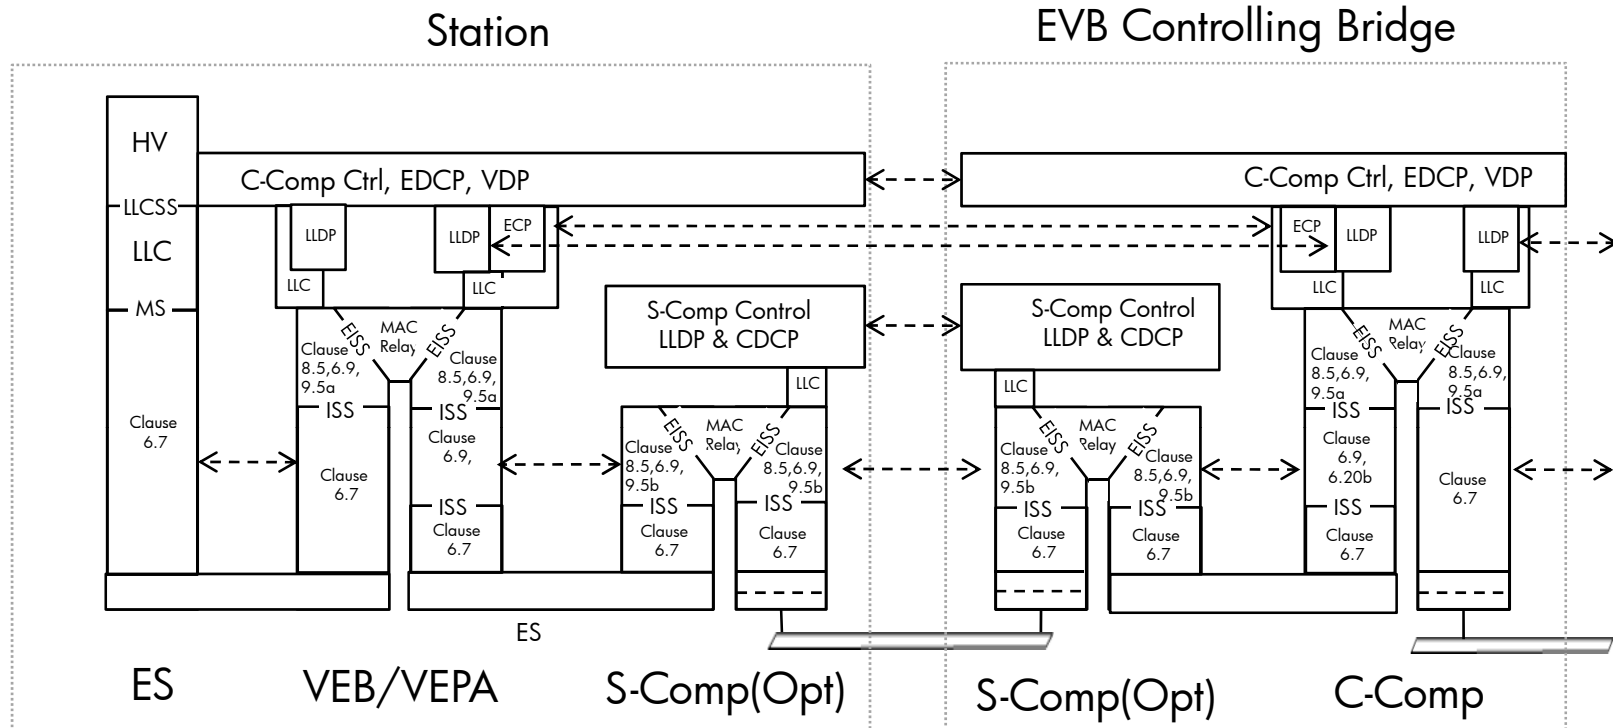


- MAC: Media Access Control 802 2001 subclause 6.2.3 and 802.1Q Rev 2010 subclause 6.1
- ISS: Internal Sublayer Service 802.1Q Rev 2010 subclause 6.6
- LLC: Link Layer Control Protocol see 802 2001 subclause 6.2.2 and 802.2 (note: see 802.1AB 2009 subclause 6.7)
- LLC-SS: Link Layer Control Protocol Sublayer Service 802.1AB-Rev 2009 subclause 6.7
- LLDP: Link Layer Discovery Protocol 802.1AB Rev 2009
- CDCP: S-channel Discovery and Configuration Protocol is an LLDP based S-channel discovery protocol
- EDCP: Edge Virtual Bridge Discovery and Configuration Protocol is an LLDP based EVB discovery protocol
- ECP: Edge Control Protocol new link layer protocol
- ECP-SS: Edge Control Protocol Sublayer Service new service interface for ECP to ULP
- S-Comp: Draft 802.1Qbc Port mapping S-VLAN component subclause 5.10 & S-VLAN component 802.1Q Rev 2010 subclause 5.6
- VEB/VEPA: C-VLAN component 802.1Q Rev 2010 subclause 5.5

Address Choices for LLDP

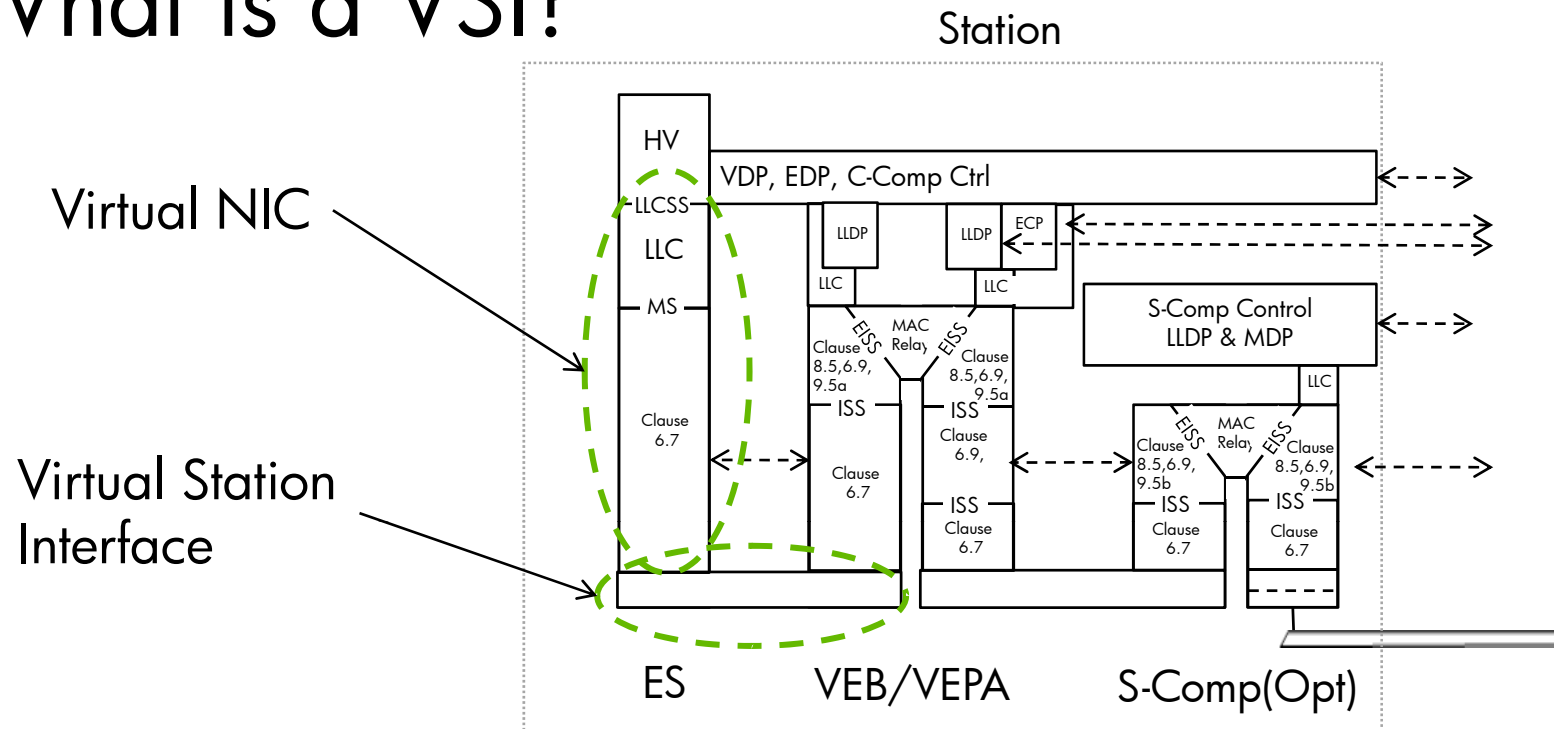
- The Port-mapping S-VLAN component used to implement S-channels is a conformant S-Comp and therefore we use Nearest non-TPMR bridge address (01-80-C2-00-00-03) for LLDP at CDCP.
- The second LLDP database is at an LLC SAP of the ES/VEB/VEPA. Both of these are addresses by the Nearest Customer Address: 01-80-C2-00-00-00.
- If we don't have an S-Comp then (see LLDP clause 7.4)
 - we only send and receive on the nearest customer bridge address
 - we discard any nearest non-TPMR bridge addressed frames
 - we discard any nearest bridge addressed frames
- If we have both S-Comp and ES/VEB/VEPA then we
 - send on both nearest customer bridge and nearest non-TPMR bridge
 - we discard any nearest bridge addressed frames

Bridge “Baggy Pants” Model



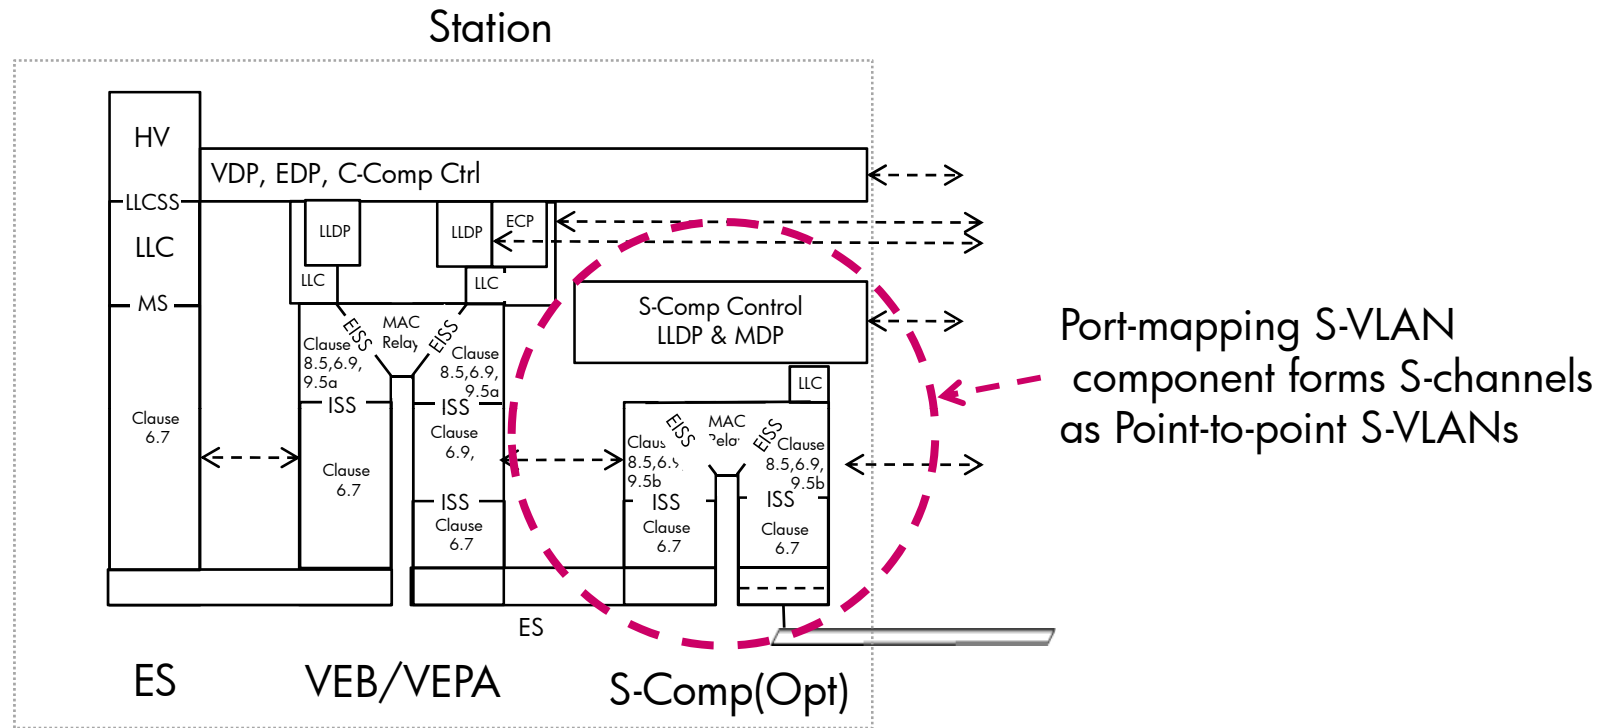
- A station uses a dual relay. The outside relay is an S-VLAN aware component. The inside relay is a VEB or VEPA. Stations also have an ES layer providing LLC service to the Hypervisor and VMs.
- The S-Comp relay is optional. When an S-channel is connected to a single ES a 2-port VEB is used which provides C-tag functions for the VM or hypervisor.
- An S-channel is implemented using the peered Port-mapping S-Comps (803.1Qbc). An LLDP database exists on each exterior facing leg of the S-Comp
- Each VEB, VEPA or 2-Port VEB has an LLDP database on it's exterior facing legs
- Each VEB VEPA or 2-Port VEB may have an LLDP database facing at the Bridge Port facing the VM. It also will filter all LLDP multicast addresses.

What is a VSI?



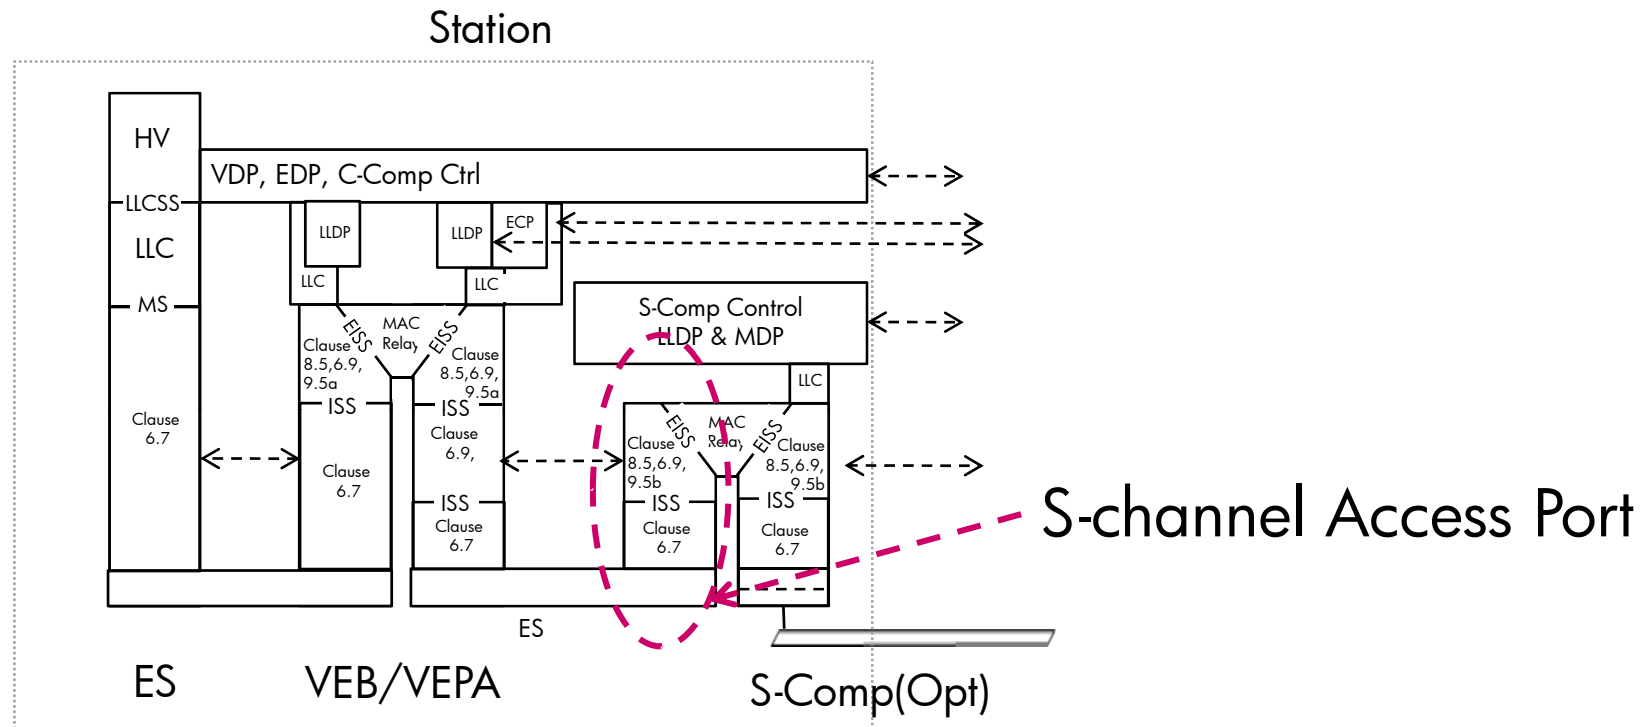
- A Virtual Station Interface (VSI) is an internal LAN which connects between a virtual NIC and a Bridge Port of a VEB, VEPA, or 2-Port VEB. A VSI carries a single MAC service instance. The term VSI is also used to reference a point where a Bridge Port of a VEB or VEPA attaches to a vNIC.
- vNIC
 - Source and Sink for a single instance of 802.3 MAC service
 - Includes functions:
 - 1)ES LLC used to form the SAP for attaching the VM
 - 2)ES MAC address filter for an station and group addresses
 - 4)VDP control
 - 5)Management for VSI functions

What is a S-channel?



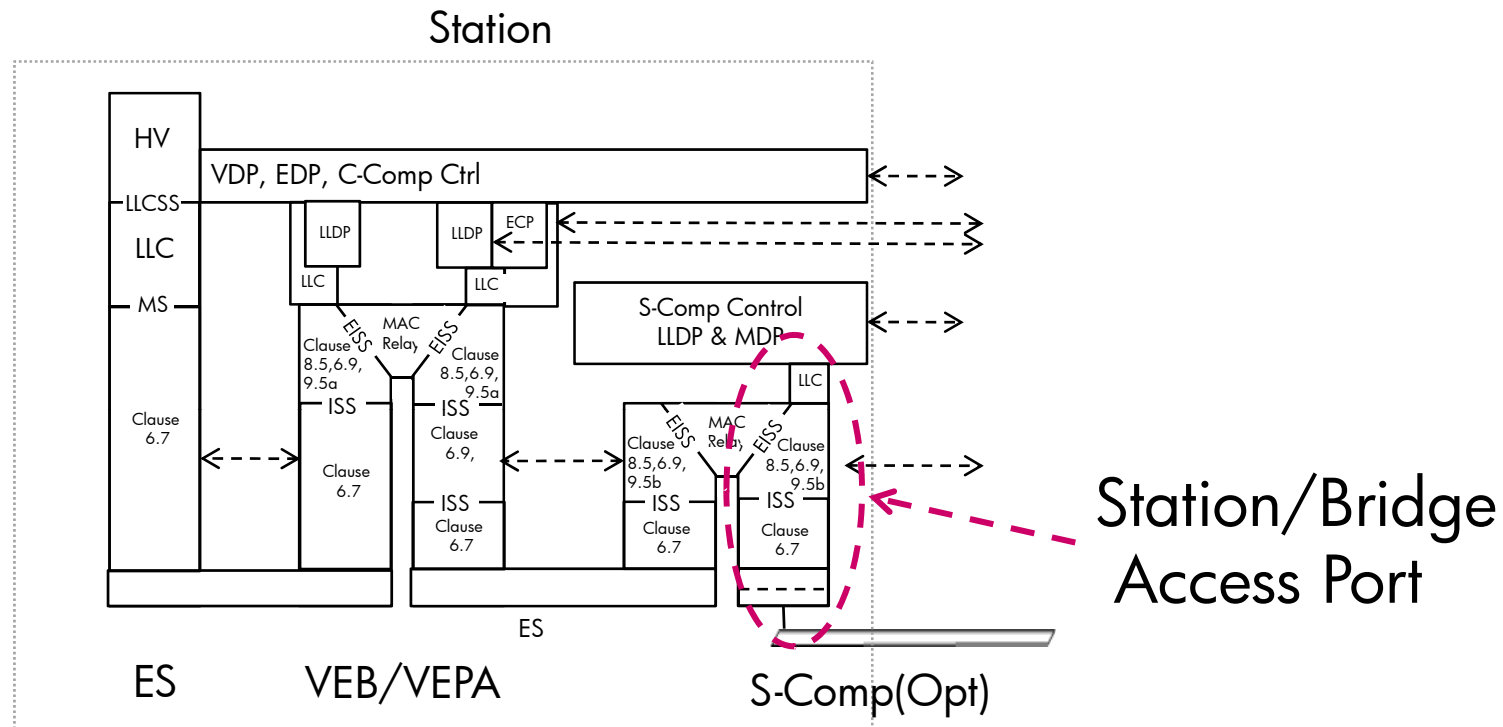
- S-channels:
 - 1) are pt-pt S-VLANs formed by Port-mapping S-VLAN components
 - 2) begin and end at an S-channel Access Port (CAP)
 - 3) frames are always S-tagged in the channel and always un-S-tagged by the CAP
 - 4) Only a single S-channel may terminate (or begin) at a each CAP
 - 5) Multiple S-channels pass through the Station Bridge Access Port to the external LAN

What is a S-channel Access Port (CAP)?



- S-channel Access Port (CAP)
 - 1)The Bridge Port of the Port-mapping S-VLAN component where an S-channel terminates
 - 2)Always un-S-tagged by the CAP

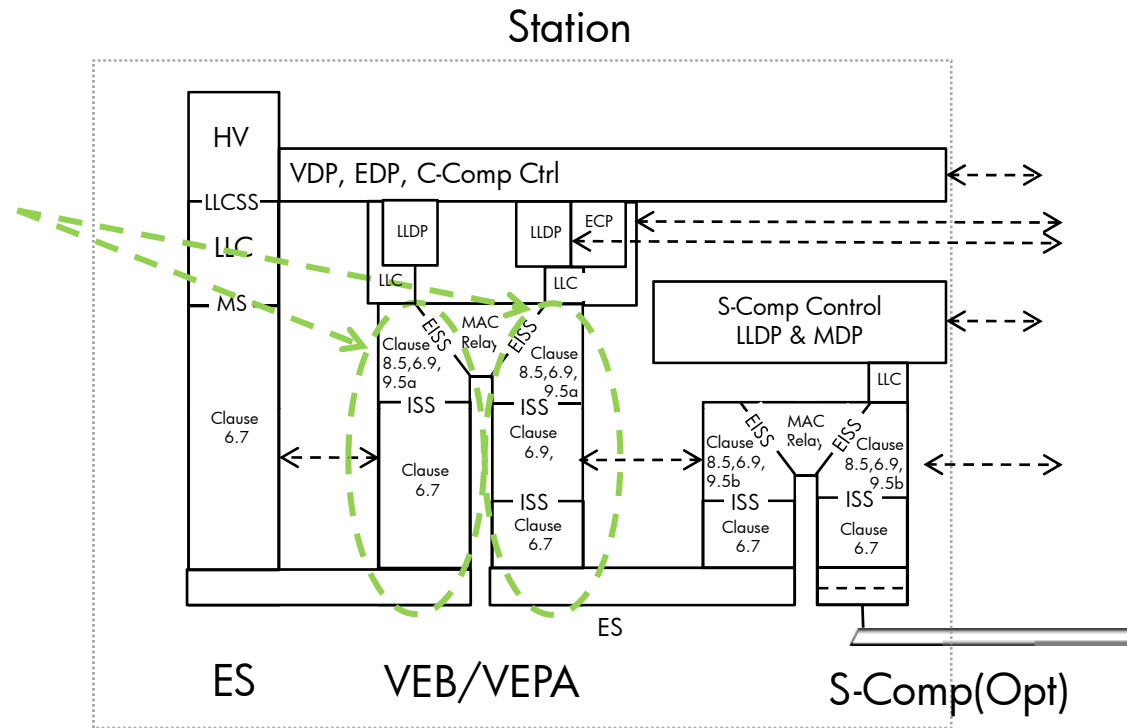
What is a Station/Bridge Access Port(SBAP)?



- Station Bridge Access Port (SBAP)
 - 1)The Bridge Port of the Port-mapping S-VLAN component where an S-channel connects between the station and Bridge
 - 2)A S-channel is always S-tagged at the SBAP

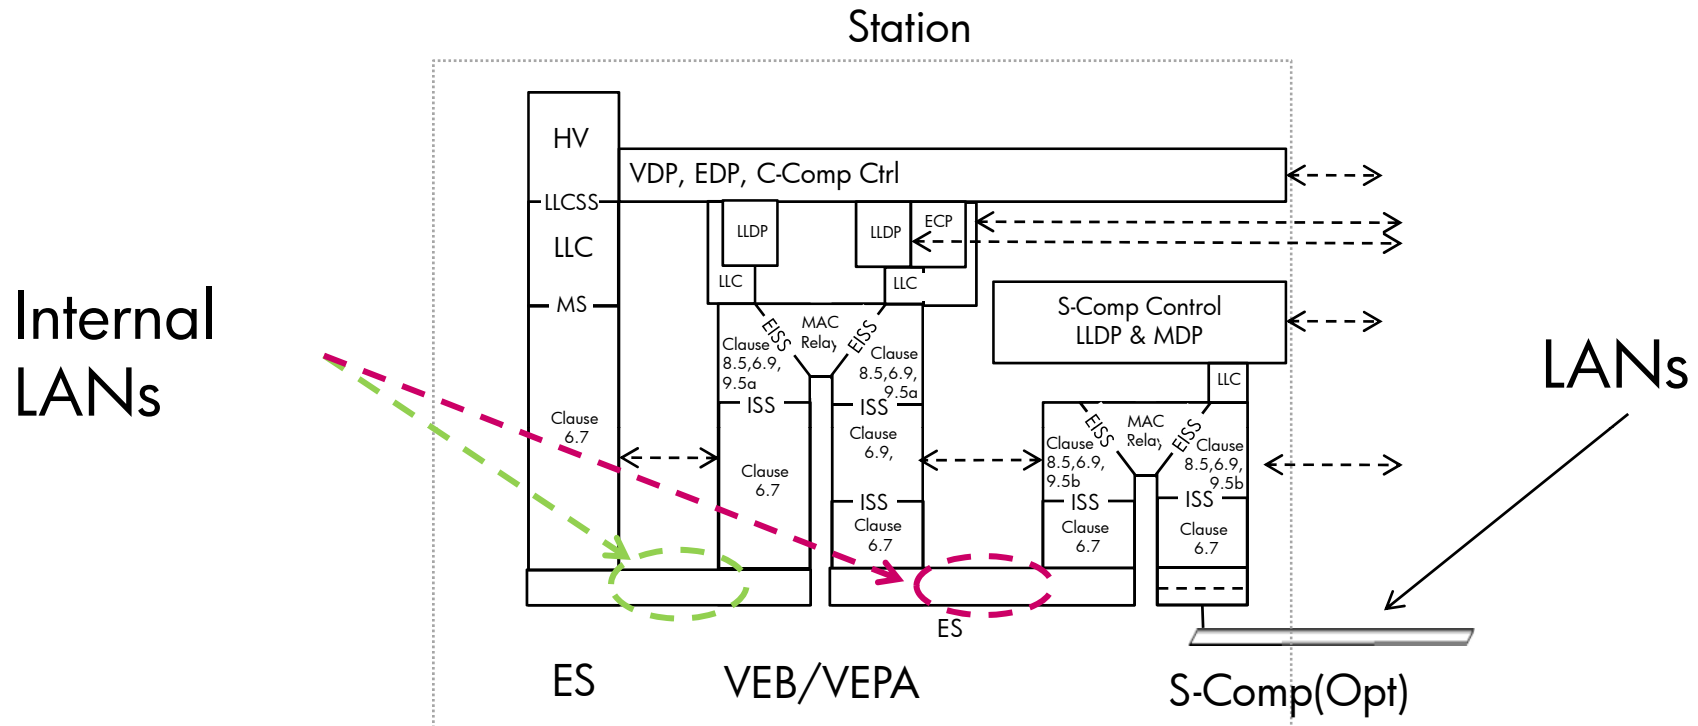
What are the VEB/VEPA ports?

Bridge Ports
(reduced
management?)



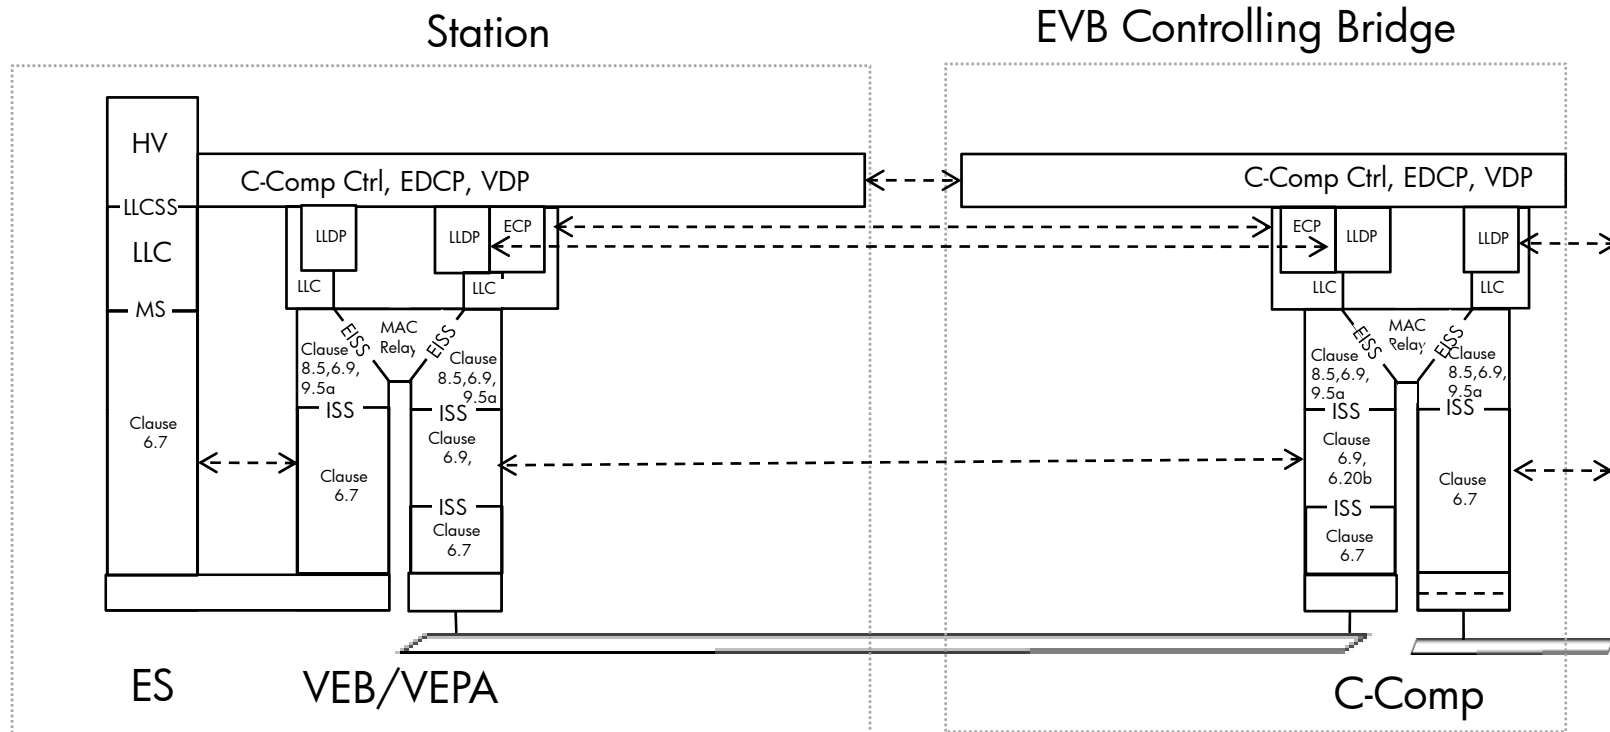
- VLAN Bridge Ports
 - 1) Just regular VLAN aware Bridge Ports

Connection to VEPA/VEB?



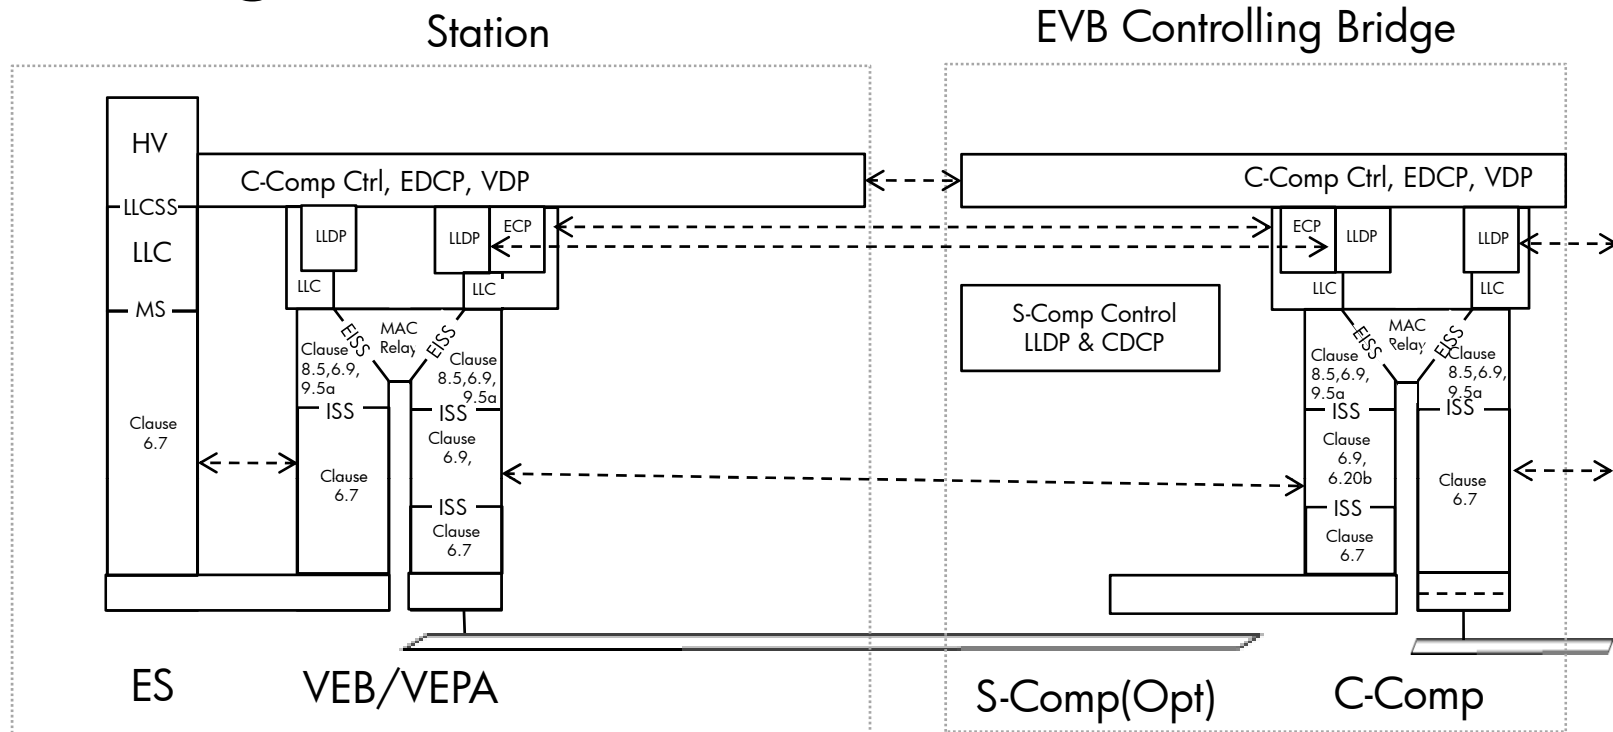
- Two internal LANs
 - 1) Internal LAN between the interior facing VLAN Bridge Ports or VEB/VEPA and the vNIC
 - The point where the internal LAN attaches is the Virtual Station Interface
 - 2) Internal LAN between the CAPs and the exterior facing VLAN Bridge Ports of the VEB/VEPA

Configuration Without S-channels



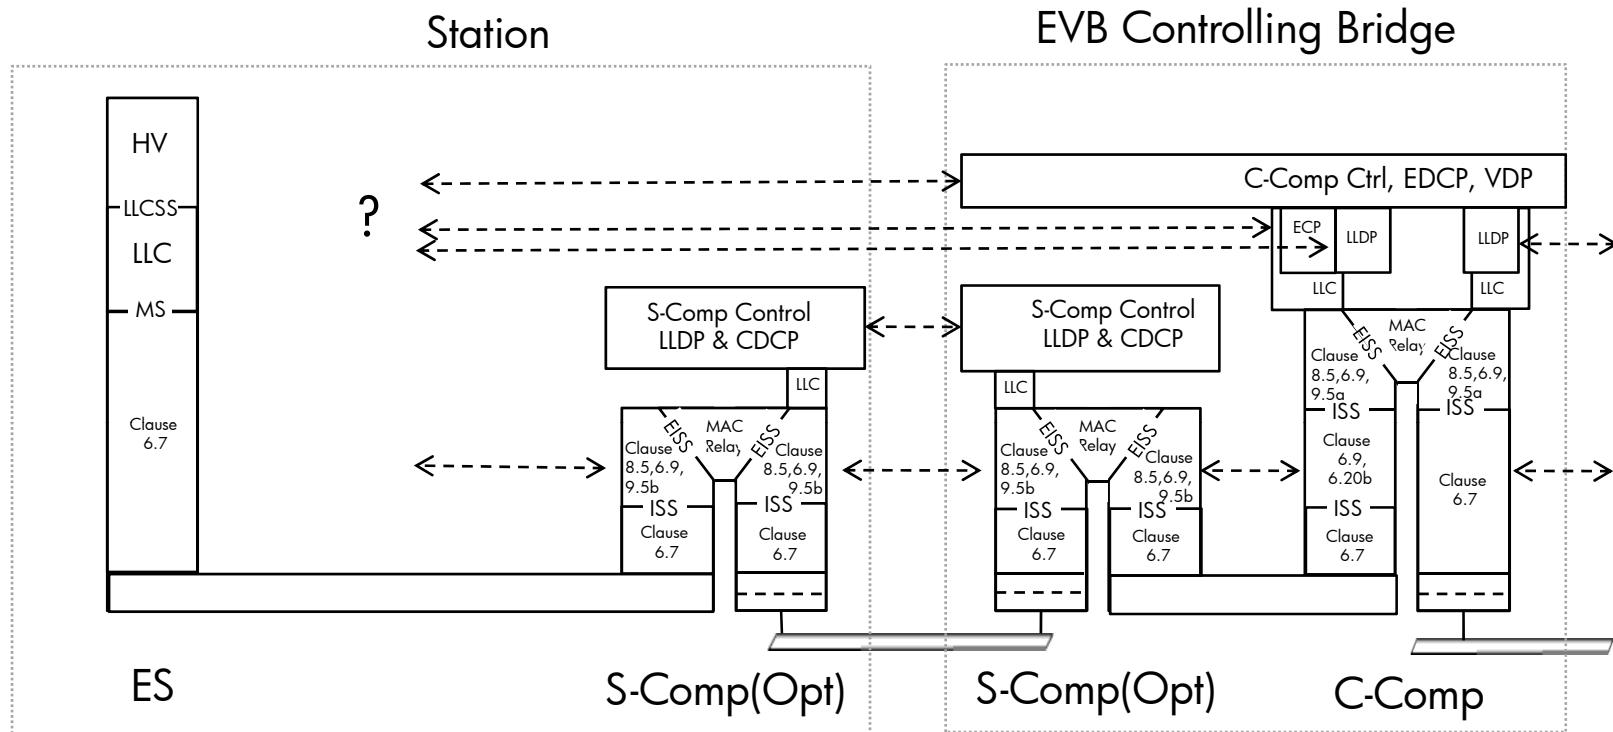
- No Port-mapping S-VLAN components
 - Station supports only a single VEB/VEPA/ES
 - LLDP database uses nearest Customer address
 - VDP/ECP both supported

Configuration Without S-channels #2



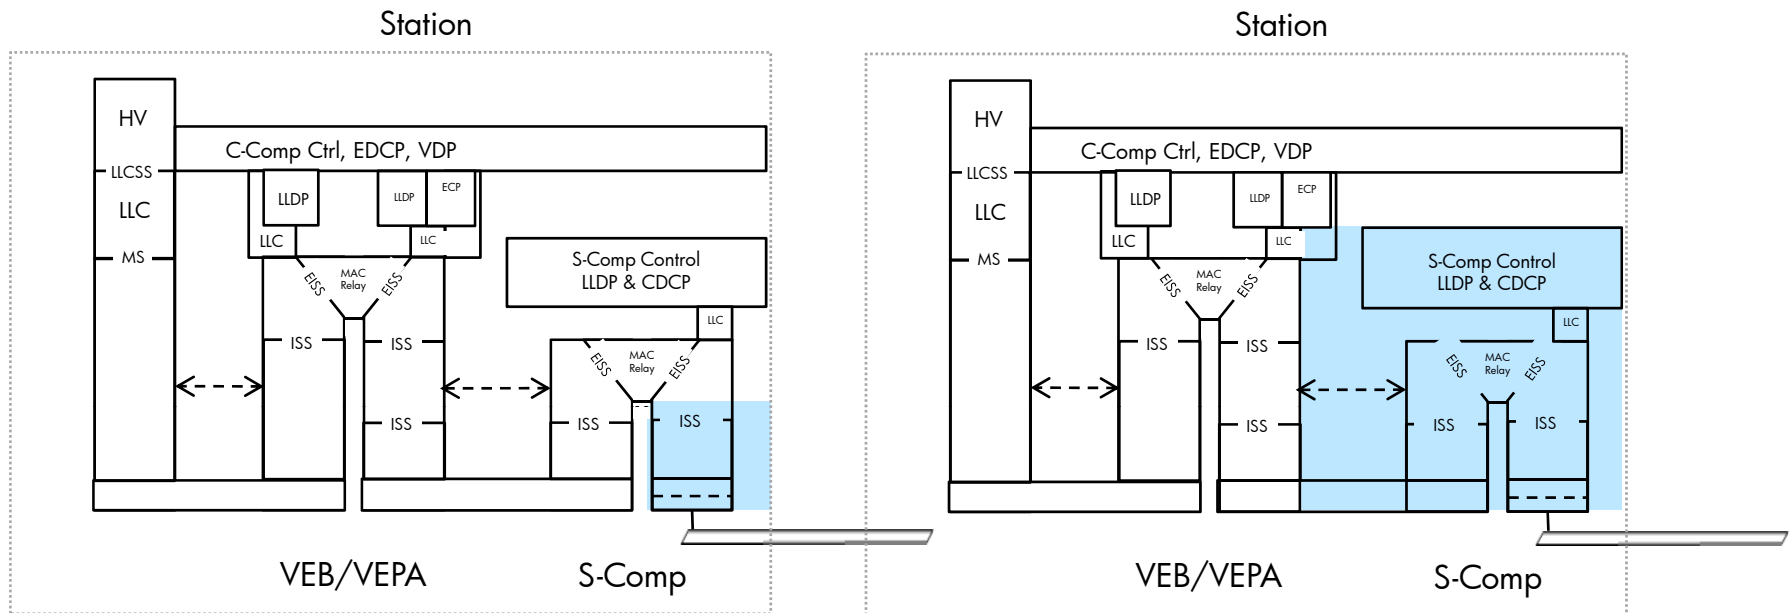
- No Port-mapping S-VLAN components on station
 - No S-channel support
 - Bridge does not see nearest non-TPMR LLDP address and so CDCP does not configure S-channels
 - Bridge S-Comp passes through VEB/VEPA frames on default S-VLAN as un-S-tagged
 - All VEB/VEPA frames land at a single Bridge C-Comp Bridge Port
 - Nearest Customer LLDP database builds as normal
 - EDCP/ECP/VDP operate as normal
 - Operation similar for case where Station has S-Comp, but bridge does not
 - Don't want VM to talk directly to controlling Bridge (must filter all LLDP frames)

No VEPA/VEB in Direct Attach Case?



- vNIC attached directly to S-channel
 - LLDP nearest Customer Bridge address without peer?
 - EDCP/ECP/VDP also have no peers?
 - No C-tag manipulations available?
- Alternative is 2 port VEPA which only needs to C-tag/untag and manipulate priority

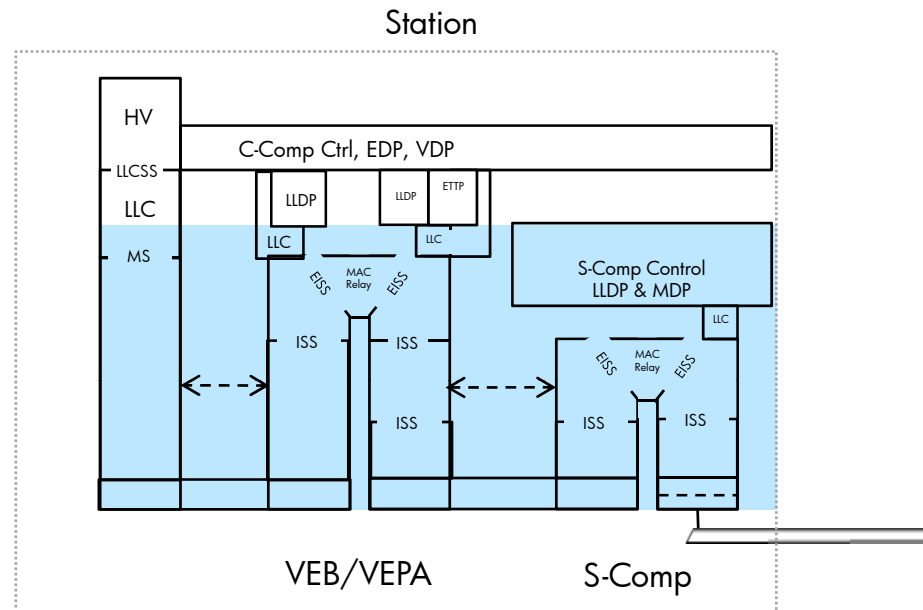
Some Adapter Types



- Link Adapter Provides
 - Ethernet MAC
 - Frame IO PF/VF

- S-channel Adapter Provides
 - Ethernet MAC
 - Frame IO PF/VF
 - S-Comp (Simplified S-VLAN Relay)
 - LAN Link Layer Discovery (LAN-LLDP)
 - S-channel Discovery and Configuration Protocol (CDCP)

Some Adapter Types Continued

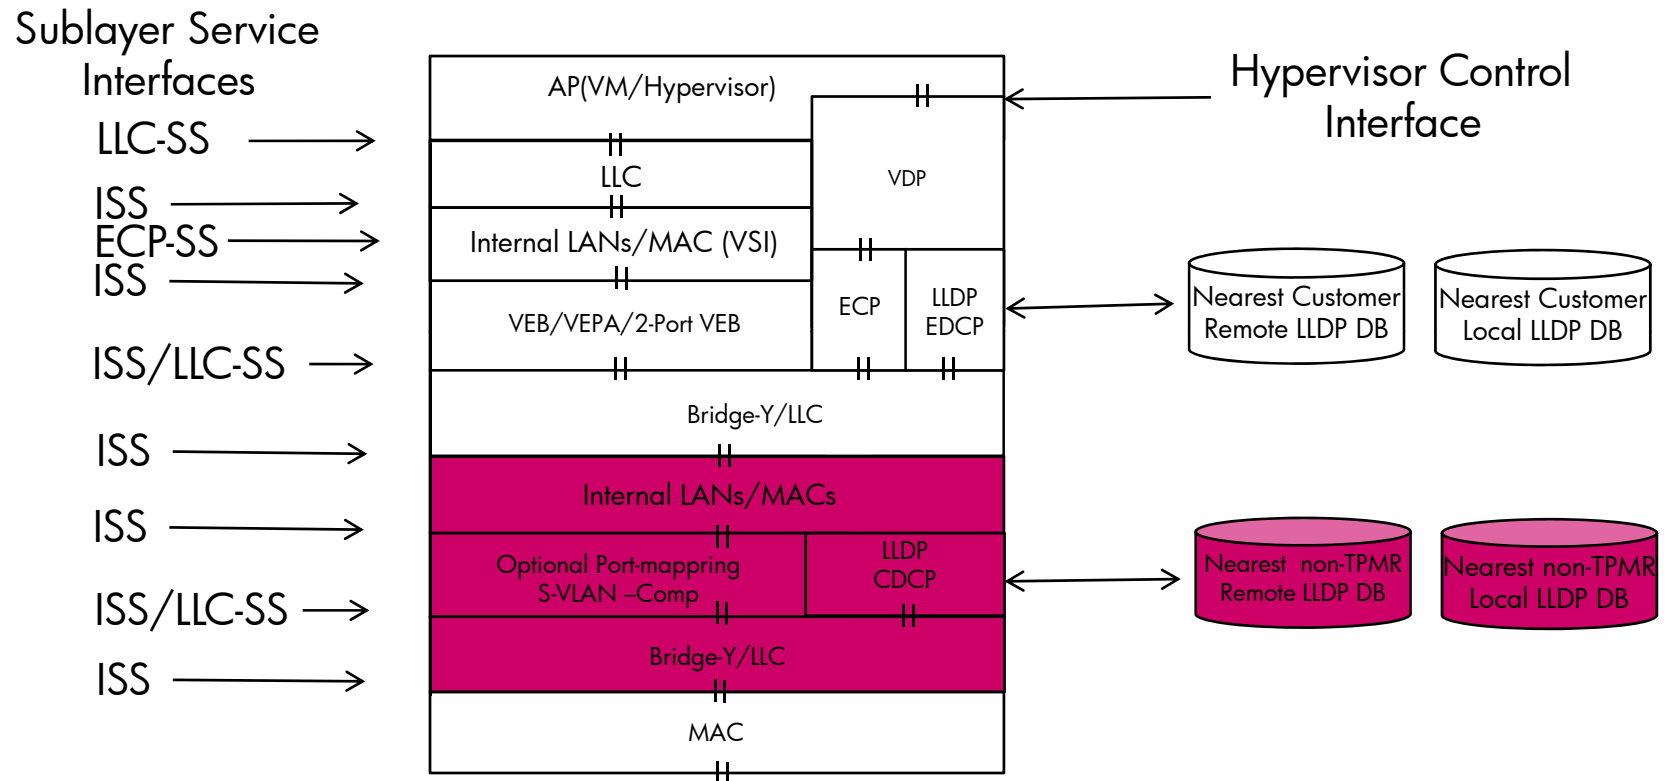


- S-channel/VEB Adapter Provides
 - Ethernet MAC
 - Frame IO PF/VF
 - S-Comp (Simplified S-VLAN Relay)
- Continued
 - LAN Link Layer Discovery (LAN-LLDP)
 - S-channel Discovery and Configuration Protocol (CDCP)
 - VEB (or VEPA)

New Service Interface for ECP

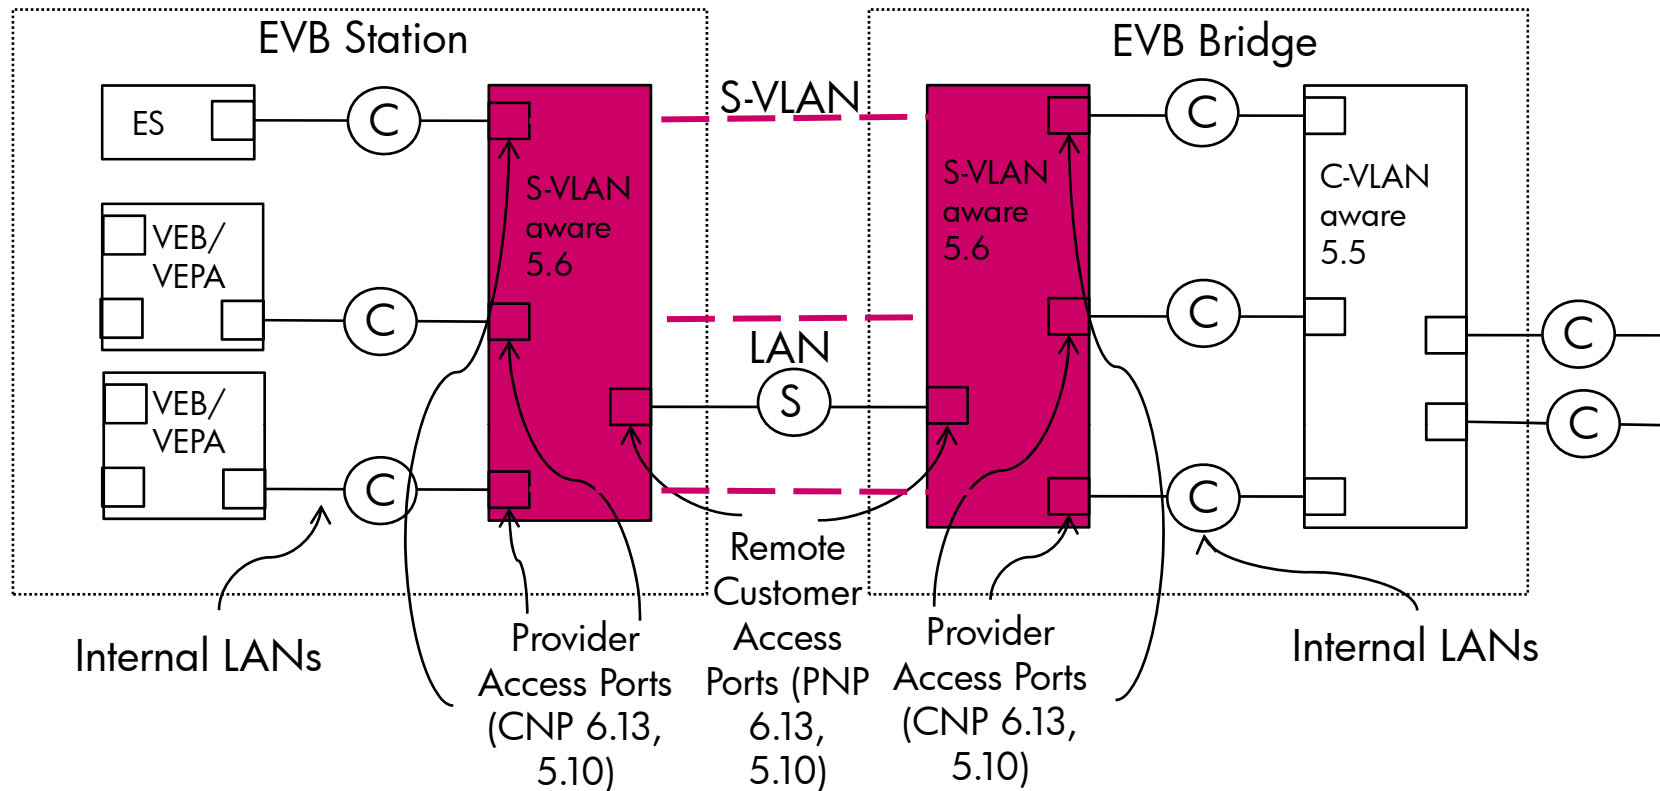
- Parameters:
 - List of TLVs (7 bit ULP type, 9 bit length, tlv-list)
 - TLV list contains TLVs from a single ULP
 - Signals: None
- Primitives (Unconfirmed Service)
 - ECP_UNITDATA.request (ulptype, ulptlv-list)
 - ECP_UNITDATA.indicate (ulptype, ulptlv-list)
- No other new service interfaces
 - VDP Listener to VDP is implementation dependent
 - Hypervisor to VDP is also implementation dependent
 - All other service interfaces are ISS, EISS or LLC-SS

S-channel Architecture



- S-Comp uses 802.1Qbc Port-mapping S-VLAN component and LAN Level LLDP to run MDP

EVB Bridge Component

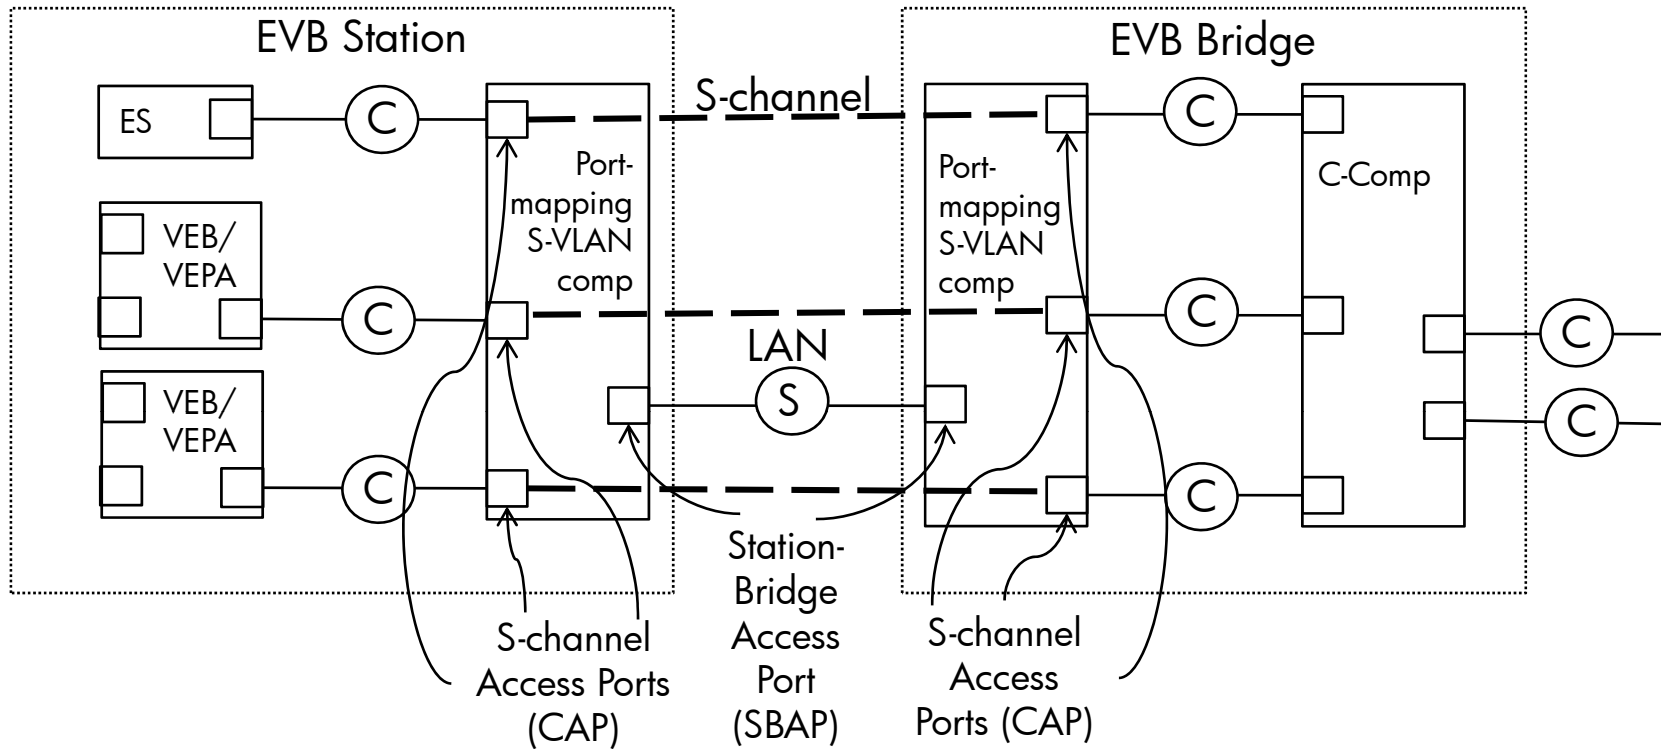


- A Port-mapping S-VLAN component (802.1Qbc) is used to implement S-channels
- Components disable spanning tree and MAC learning
- 1-1 relationship between PAP to C-Comp Bridge Ports of the EVB Bridge
- 1-1 relationship between PAP to VEB/VEPA Bridge Ports or End-Station ports of the EVB Station

Thoughts on Naming?

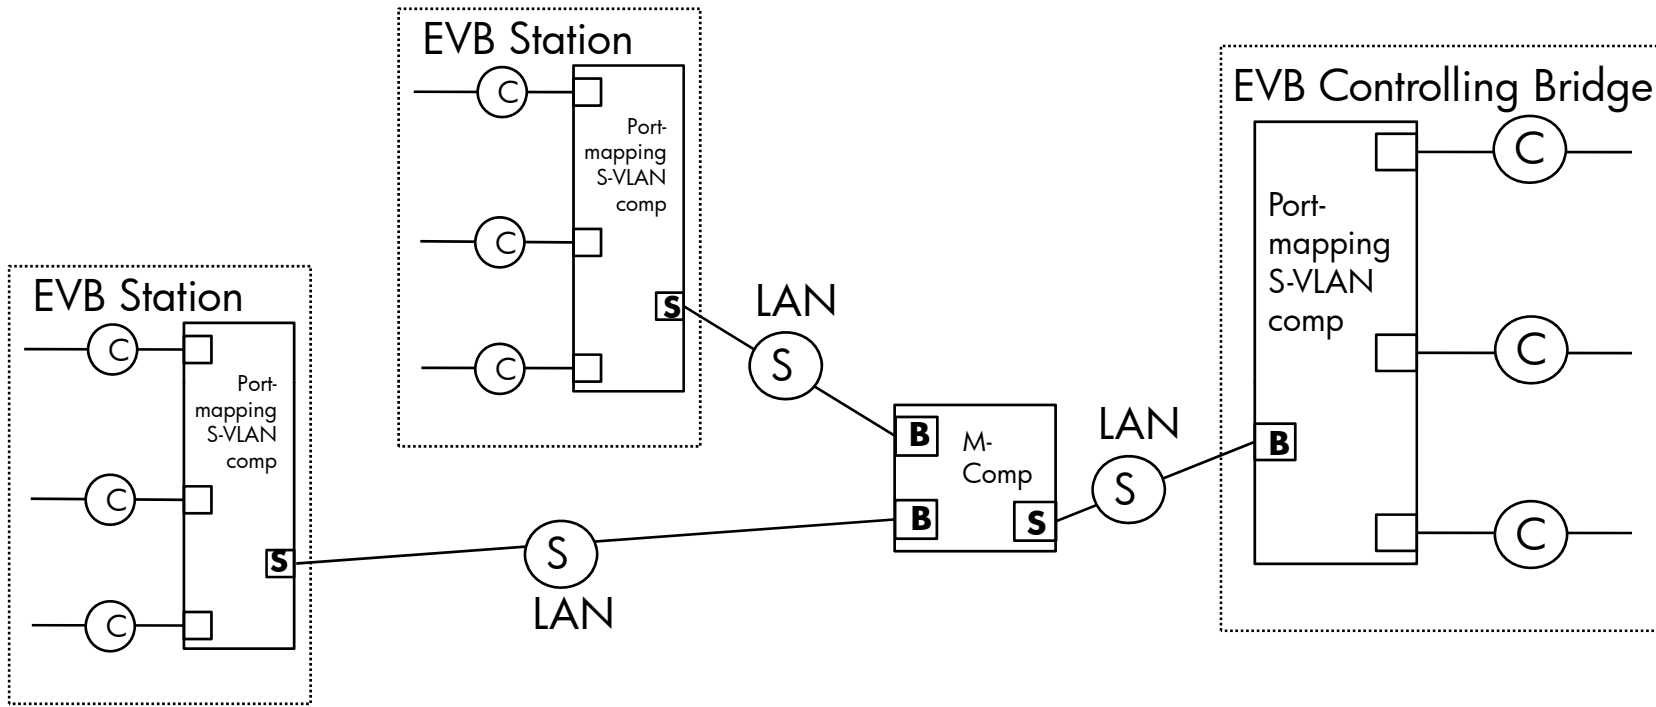
- **LAN** – the link between station and bridge
- **S-channel** – a pt-pt S-VLAN implemented between a station and Bridge over Port-mapping S-VLAN components
- **S-TAG** – the tag used to implement S-channels
- **Port-mapping S-VLAN component** – a bridge component used to implement S-channels
- **S-channel Access Port (CAP)** – a CNP (or Qbc PAP) used to terminate a S-channel in a station or bridge
- **Station Bridge Access Port (SBAP)** – a PNP (or Qbc RCAP) used between a station and bridge to implement S-channels

S-channel Components



- CAP and SBAP are same definitions as PAP and RCAP
- S-TAG remains unchanged and with same name
- S-channel is a point-to-point S-VLAN terminating at a CAP

Cascade of Port-mapping S-VLAN Comps



- LLDP from S-B hop-by-hop through cascade.

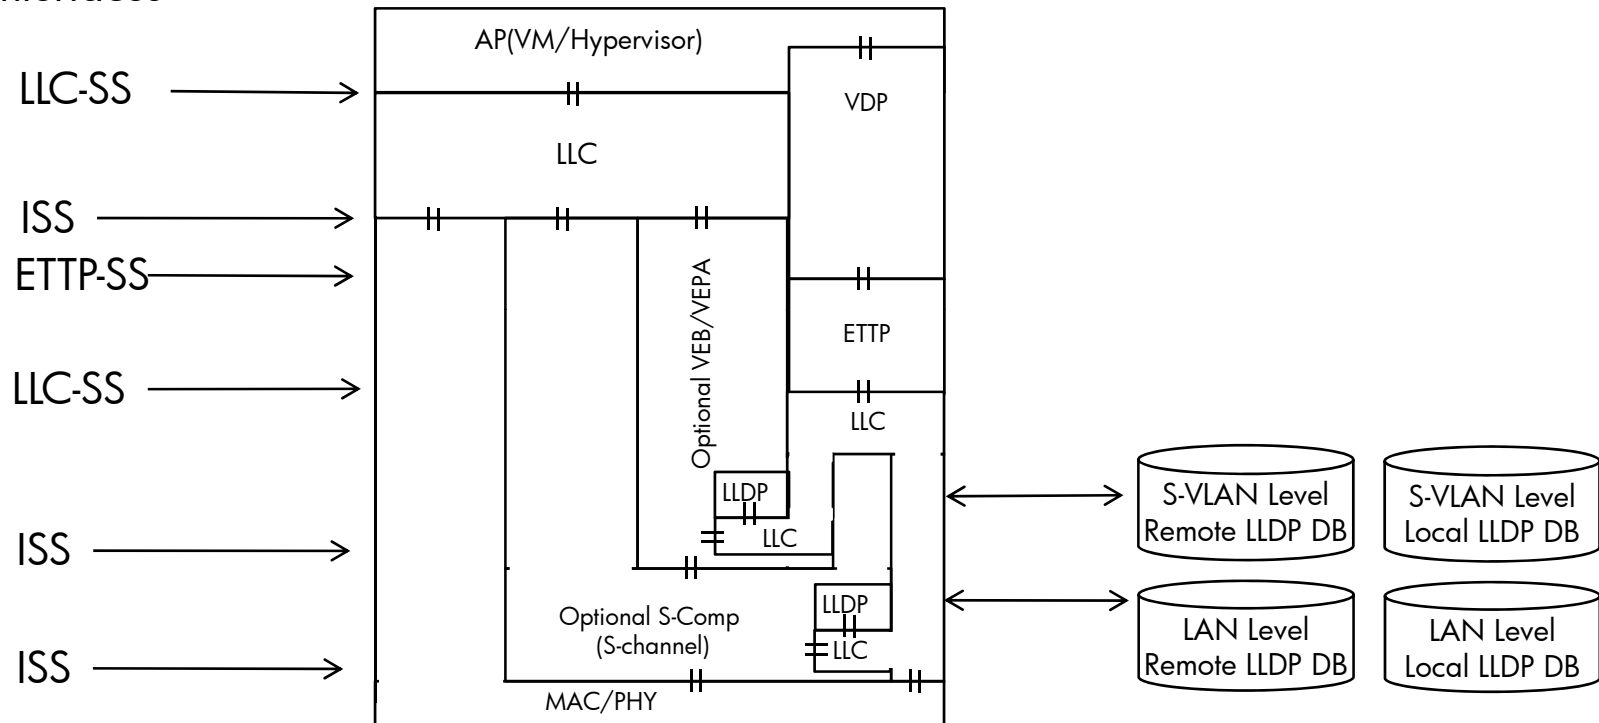
Summary

- EVB Layer Model
 - S-Comp profile for S-channels
 - VEB/VEPA bridges
 - ETTP client of LLC (like LLDP)
 - SDU service interface between ETTP and ULPs
- S-channel naming
 - Each S-channel is a point-to-point S-VLAN of 802.1Qbc
 - A S-channel terminates at a CAP (renamed from PAP)
 - The LAN between Station and Bridge terminates at a SBAP (or renamed RCAP)

BACKUP SLIDES

EVB Stack Architecture

Sublayer Service Interfaces



- Virtual Station Interfaces (VSIs) are created/deleted using the Virtual Station Interface Discovery and Configuration Protocol (VDP). VDP interfaces to the Hypervisor to co-ordinate VM creation and ETPP to communicate with the Bridge.
- Edge TLV Transport (ETPP) may be connected to a LAN through a MAC, to a VEB/VEPA internal bridge port, or ES
- If S-channels are present then two levels of LLDP databases exist, one at the LAN and one at the S-channel. If S-channels don't exist then only a single LLDP database exists.

EVB Stack Architecture – Simplified View

Sublayer Service

Interfaces

LLC-SS

ISS

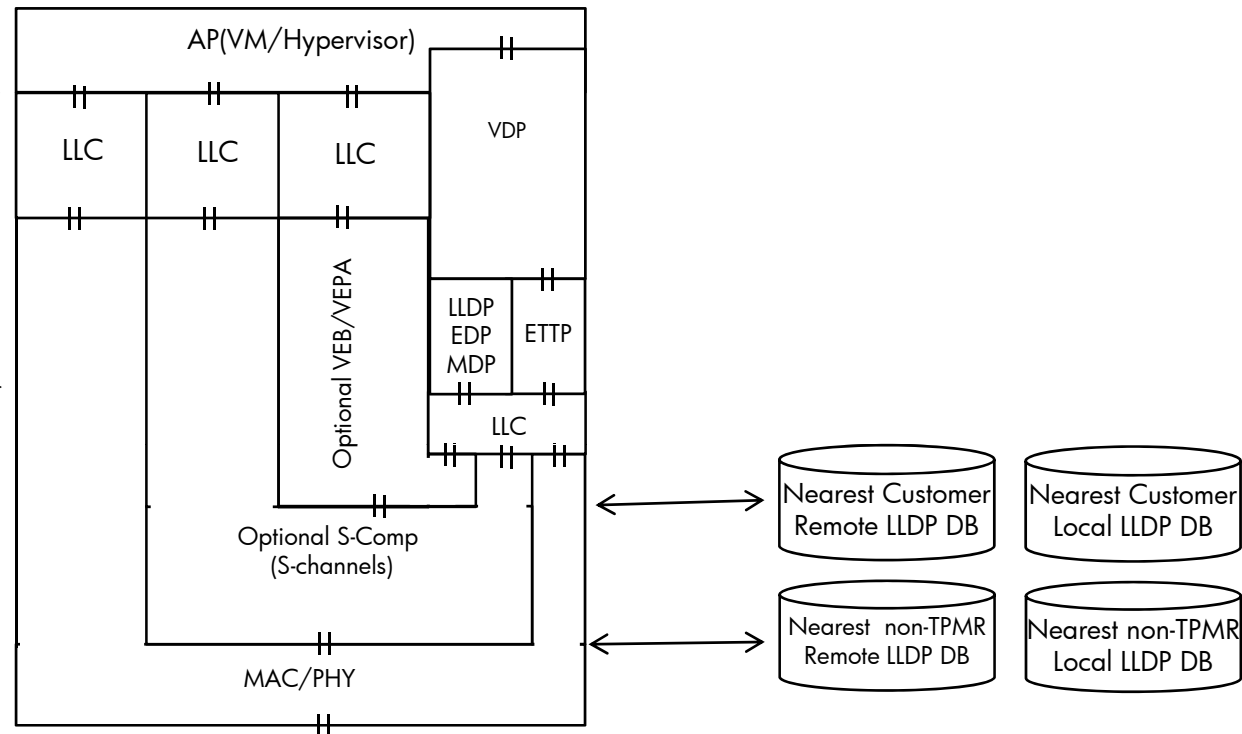
ETTP-SS

LLC-SS

ISS

ISS

ISS



- MAC: Media Access Control 802.2001 subclause 6.2.3 and 802.1Q Rev 2010 subclause 6.1
- ISS: Internal Sublayer Service 802.1Q Rev 2010 subclause 6.6
- LLC: Link Layer Control Protocol see 802.2001 subclause 6.2.2 and 802.2 (note: see 802.1AB 2009 subclause 6.7)
- LLC-SS: Link Layer Control Protocol Sublayer Service 802.1AB-Rev 2009 subclause 6.7
- LLDP: Link Layer Discovery Protocol 802.1AB Rev 2009
- CDCP: S-channel Discovery and Configuration Protocol is an LLDP based discovery protocol
- EDP: Edge Virtual Bridge Discovery and Configuration Protocol is an LLDP based EVB discovery protocol
- ETTP: Edge TLV Transport Protocol new link layer protocol
- ETTP-SS: Edge TLV Transport Protocol Sublayer Service new service interface for ETTP to ULP
- S-Comp: Draft 802.1Qbc Port mapping S-VLAN component subclause 5.10 & S-VLAN component 802.1Q Rev 2010 subclause 5.6
- VEB/VEPA: C-VLAN component 802.1Q Rev 2010 subclause 5.5