EVB Basic Architecture

V9

March 15, 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),
Renato Recio(IBM), Rakesh Sharma(IBM) , Pat Thaler(Broadcom)

EVB Overview

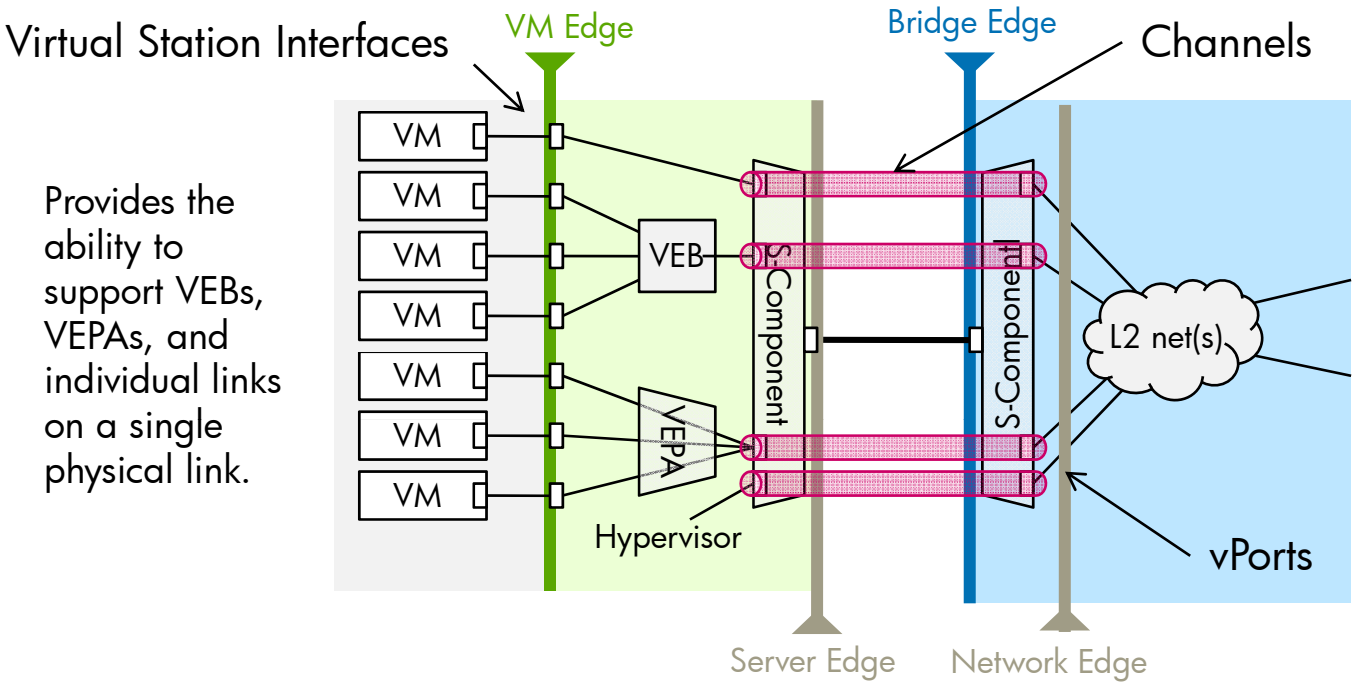

- **Channels** (or S-VLANs) are implemented using a reduced function S-VLAN aware Bridge relay.
- **Virtual Station Interfaces(VSI)** are the virtual end station attachment points for virtual machines
- **vPorts** (or internal Bridge Ports) are the termination points for S-VLANs (channels)
- **Virtual 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

EVB VSI Co-ordination Overview

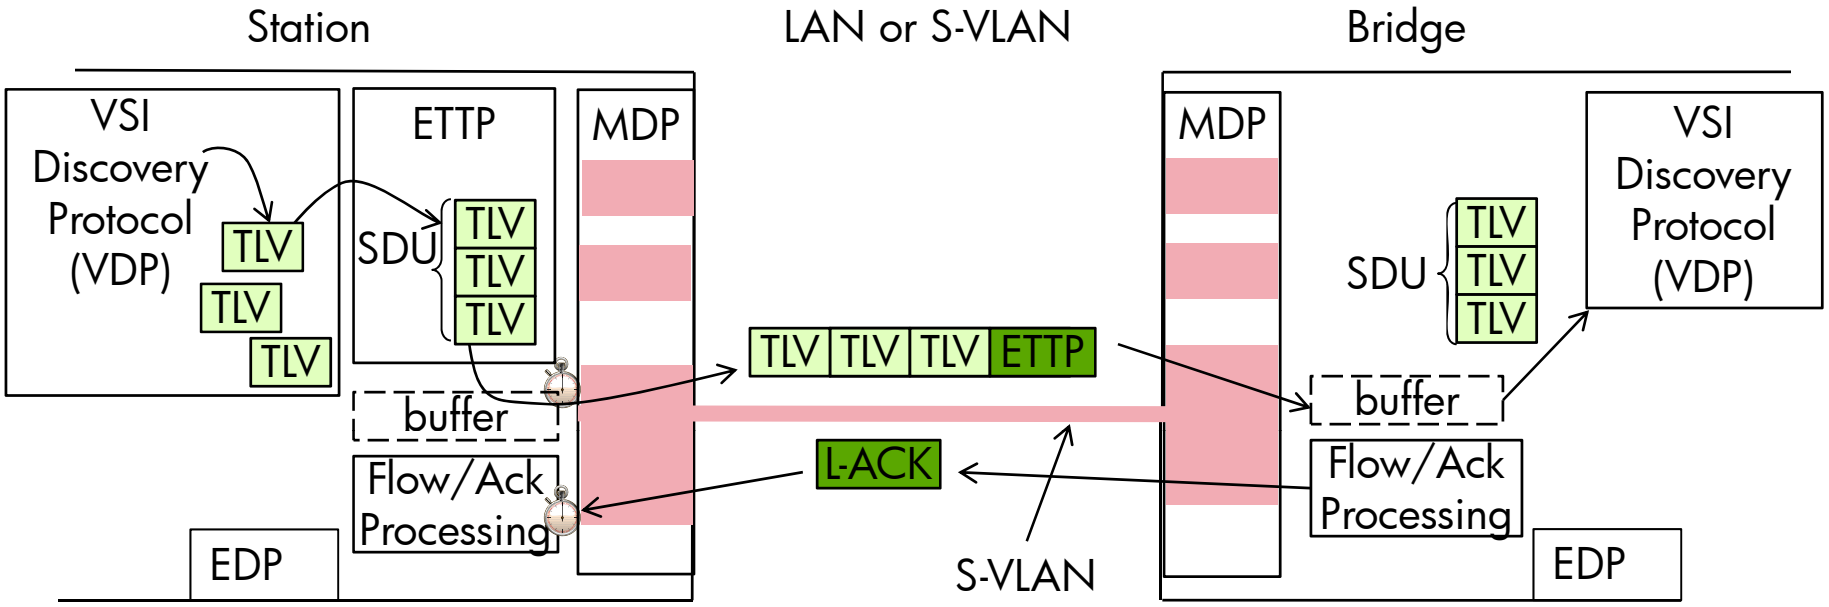

- *Virtual Station Interface(VSI) Discovery and Configuration Protocol (VDP)* co-ordinates network resources for Virtual Machines (VMs) and packs and unpacks VDP TLVs into PDUs which are handed to ETTP for delivery.
- *Edge TLV Transport Protocol (ETTP)* provides reliable delivery of SDUs between the station and bridge
- MultiChannel Discovery and Configuration Protocol (MDP)is an LLDP based protocol used to configure the S-VLANs.
- Edge Discovery and Configuration Protocol (EDP) is an LLDP based protocol used to discover and configure VEPA, ETTP and VDP

Example VSI Discovery Implementation

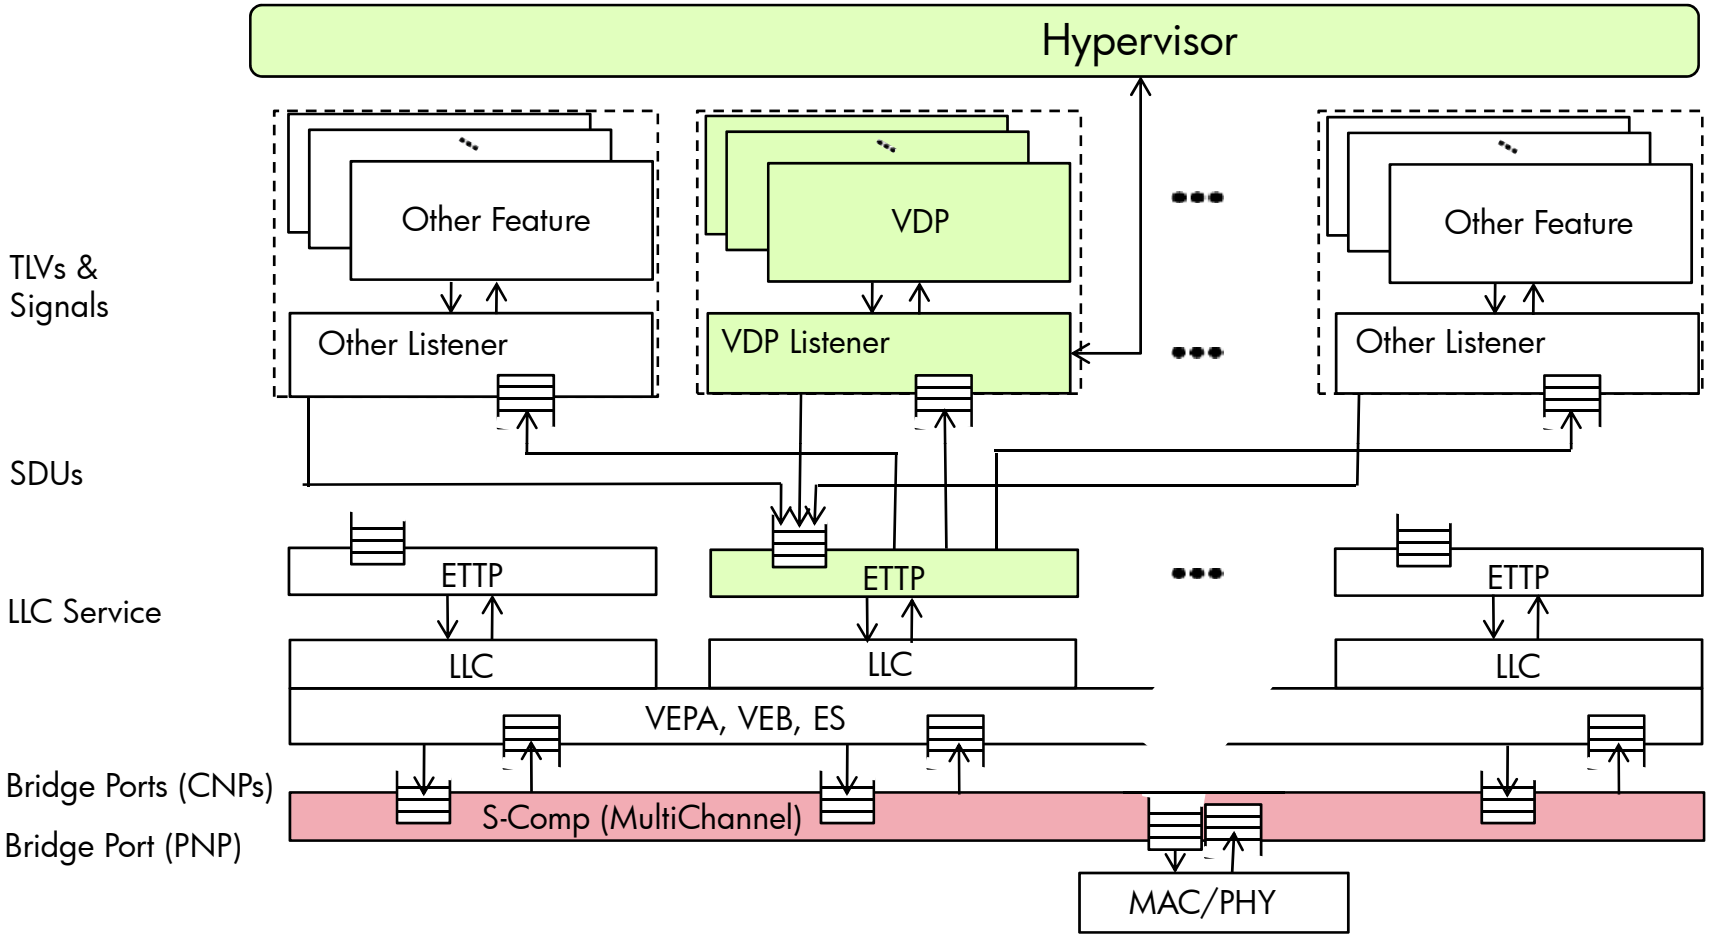

- There is one VDP per VSI. The VDP handles all VSI state transitions. It sends and exchanges TLVs and Signals from listener.
- There is one listener per ULP. The listener handles packing and unpacking TLVs into PDUs.
- There is one ETTP per LAN. It provides reliable SDU delivery for SDUs which are placed in its transmit queue by the listeners. It also delivers SDUs to the appropriate listener.

EVB Stack Architecture

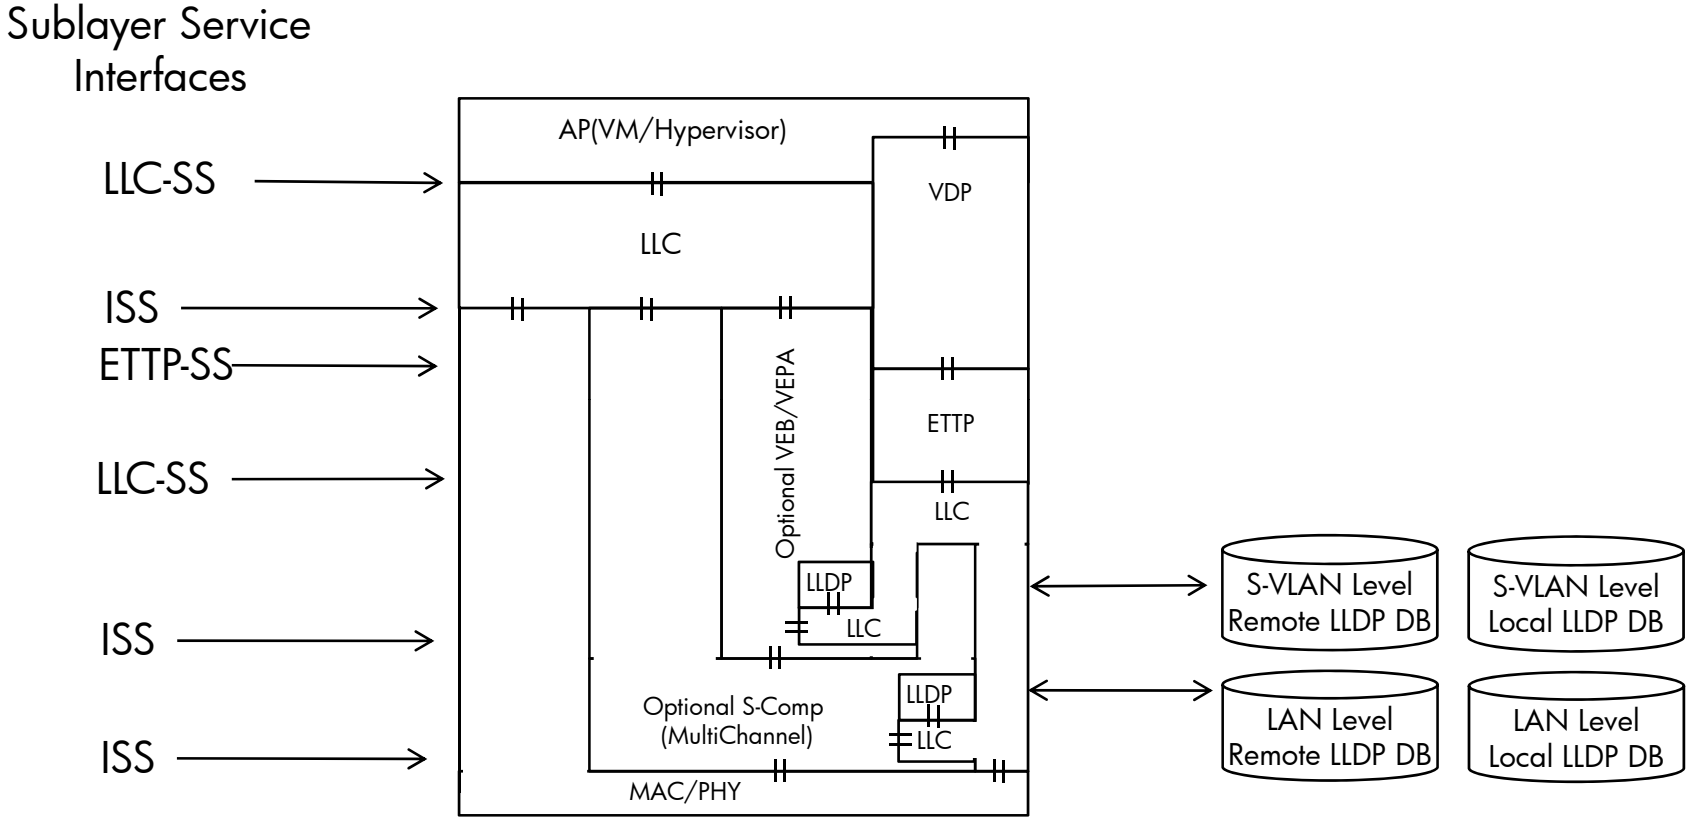

- Virtual Station Interfaces (VSIs) are created/deleted using the Virtual Station Interface Discovery and Configuration Protocol (VDP). VDP interfaces to the Hypervisor to co-ordinate VM creation and ETTP to communicate with the Bridge.
- Edge TLV Transport (ETTP) may be connected to a LAN through a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

EVB Stack Architecture – Simplified View

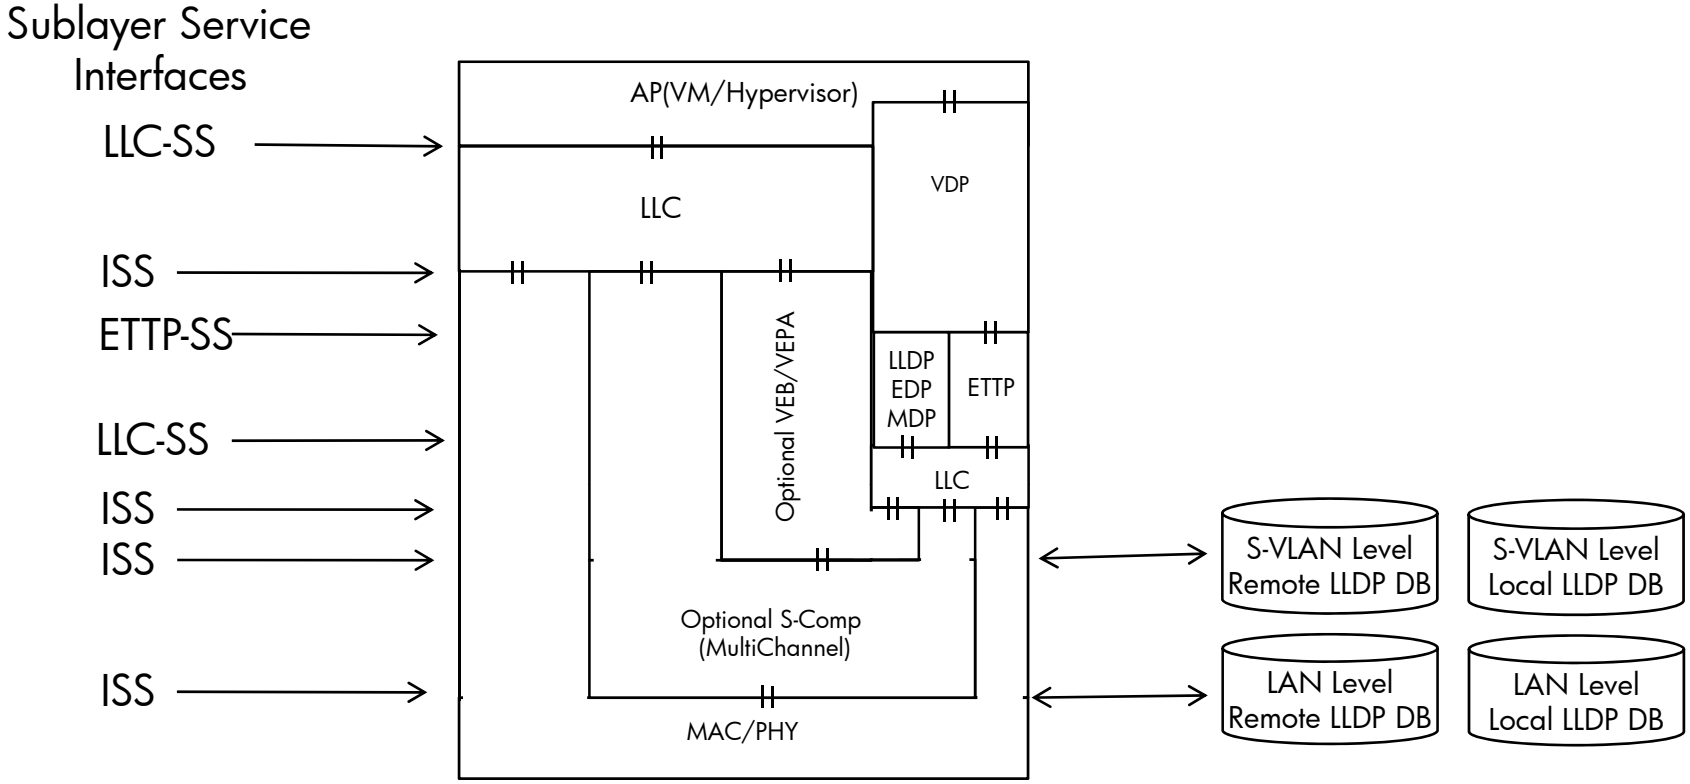

- MAC: Media Access Control 802 2001 subclause 6.2.3 and 802.1Q Rev 2010 subclause 6.1
- ISS: Internal Sublayer Service 802.1Q Rev 2010 subclause 6.6
- LLC: Link Layer Control Protocol see 802 2001 subclause 6.2.2 and 802.2 (note: see 802.1AB 2009 subclause 6.7)
- LLC-SS: Link Layer Control Protocol Sublayer Service 802.1AB-Rev 2009 subclause 6.7
- LLDP: Link Layer Discovery Protocol 802.1AB Rev 2009
- MDP: MDP MultiChannel Discovery and Configuration Protocol is an LLDP based MultiChannel 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

Bridge "Baggy Pants" Model

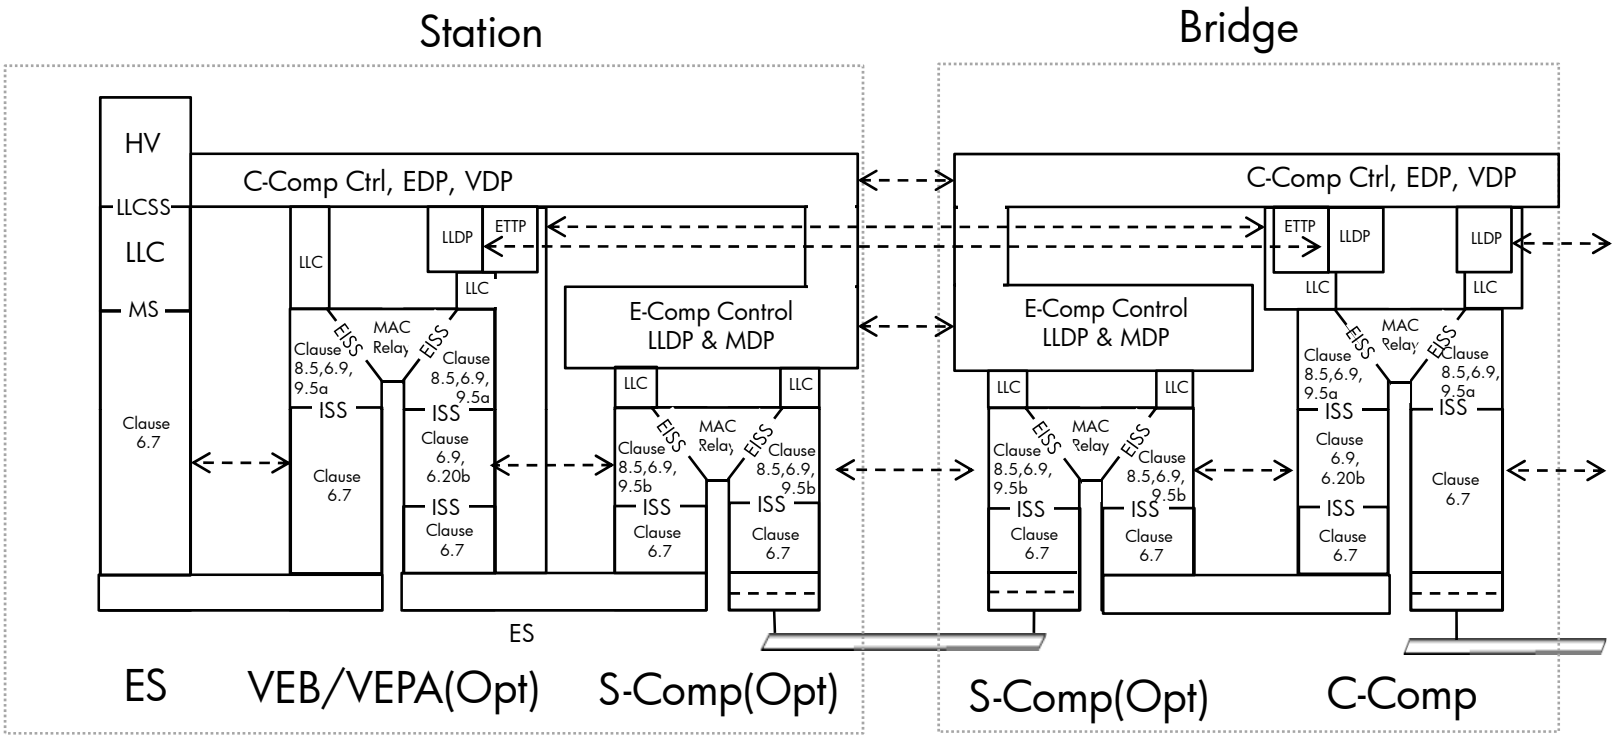

- Station uses a dual relay. The outside relay is an S-VLAN aware Component. The inside relay is a VEB or VEPA. Stations also have an ES layer providing LLC service to the Hypervisor and VMs.
- Both the MultiChannel (S-Comp) and VEB/VEPA relays are optional. The S-Comp can terminate in a ES rather than a VEB/VEPA.
- Multichannel is implemented using the peered S-Comps. An LLDP database exists on each exterior facing leg of the S-Comp
- Each VEB or VEPA has an LLDP database on it's exterior facing legs

Some Adapter Types

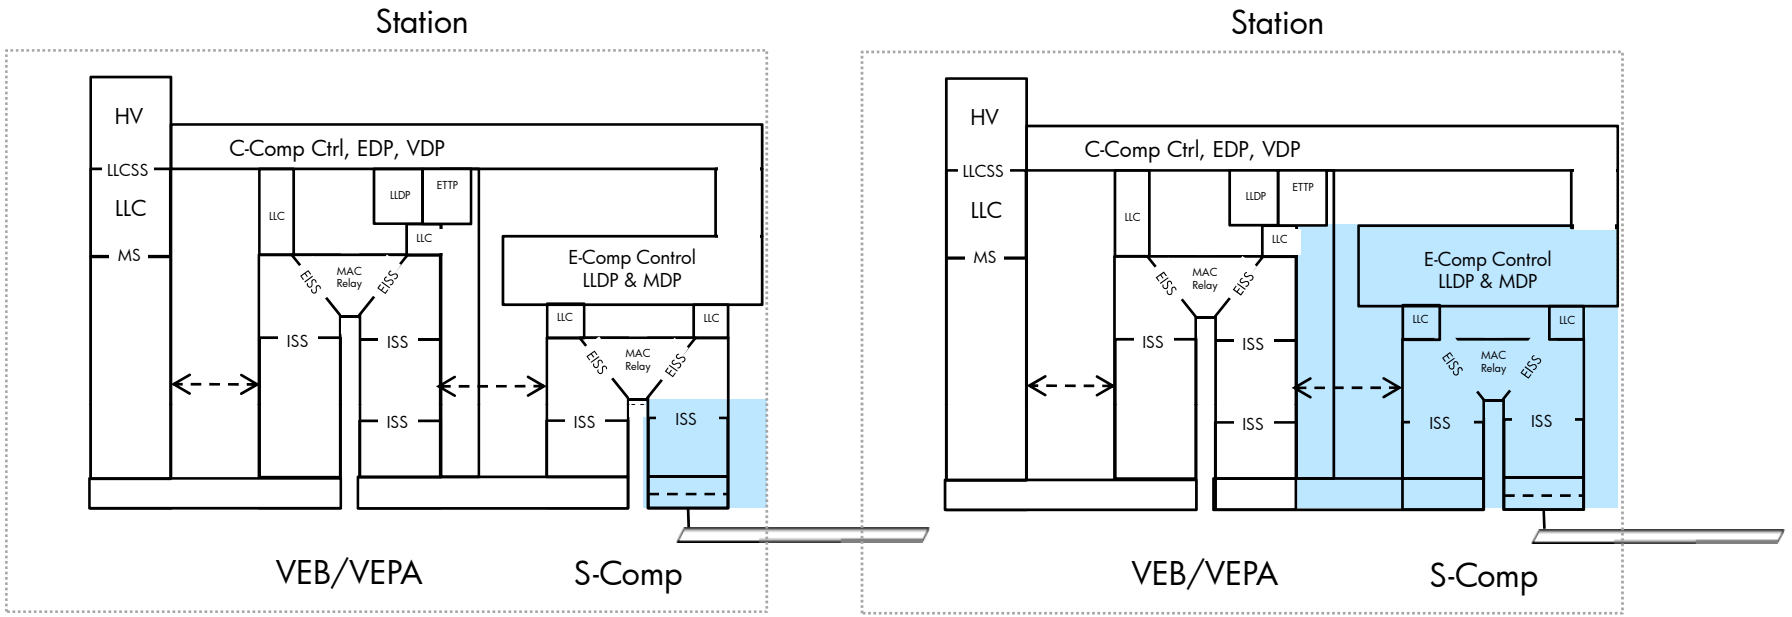

- Link Adapter Provides
  - Ethernet MAC
  - Frame IO PF/VF

- Multi-Channel Adapter Provides
  - Ethernet MAC
  - Frame IO PF/VF
  - S-Comp (Simplified S-VLAN Relay)
  - LAN Link Layer Discovery (LAN-LLDP)
  - MultiChannel Discovery (MDP)

Some Adapter Types Continued

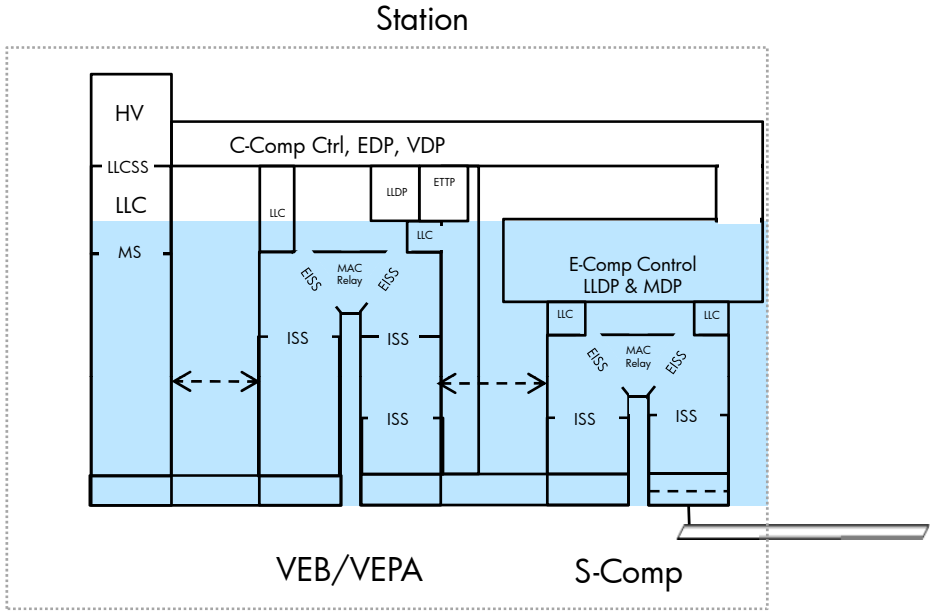

- MultiChannel/VEB Adapter Provides
  - Ethernet MAC
  - Frame IO PF/VF
  - S-Comp (Simplified S-VLAN Relay)
- Continued
  - LAN Link Layer Discovery (LAN-LLDP)
  - MultiChannel Discovery (MDP)
  - VEB (or VEPA)

New Service Interface for ETTP

- Parameters:
  - List of TLVs (7 bit ULP type, 9 bit length, tlv-list)
    - TLV list contains TLVs from a single ULP
  - Signals: None
- Primatives (Unconfirmed Datagram Service)
  - ETTP_UNITDATA.request (ulptype, ulptlv-list)
  - ETTP_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-Comp (MultiChannel) Stack Architecture

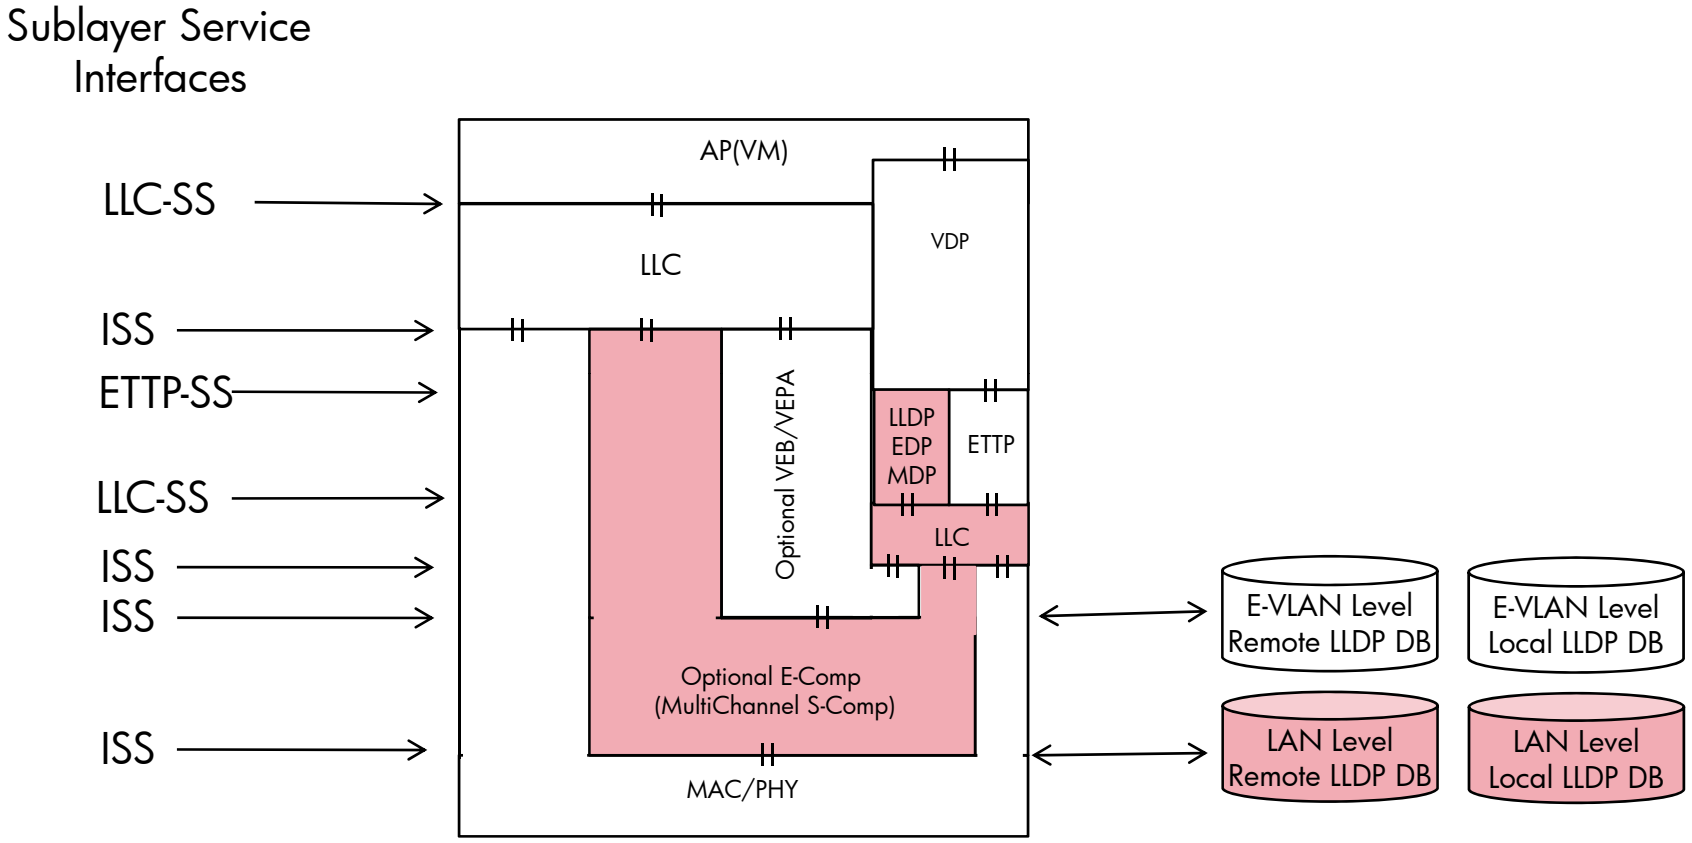

- S-Comp uses 802.1Qbc Port-mapping S-VLAN component and LAN Level LLDP to run MDP

Multi-Channel Bridge Component

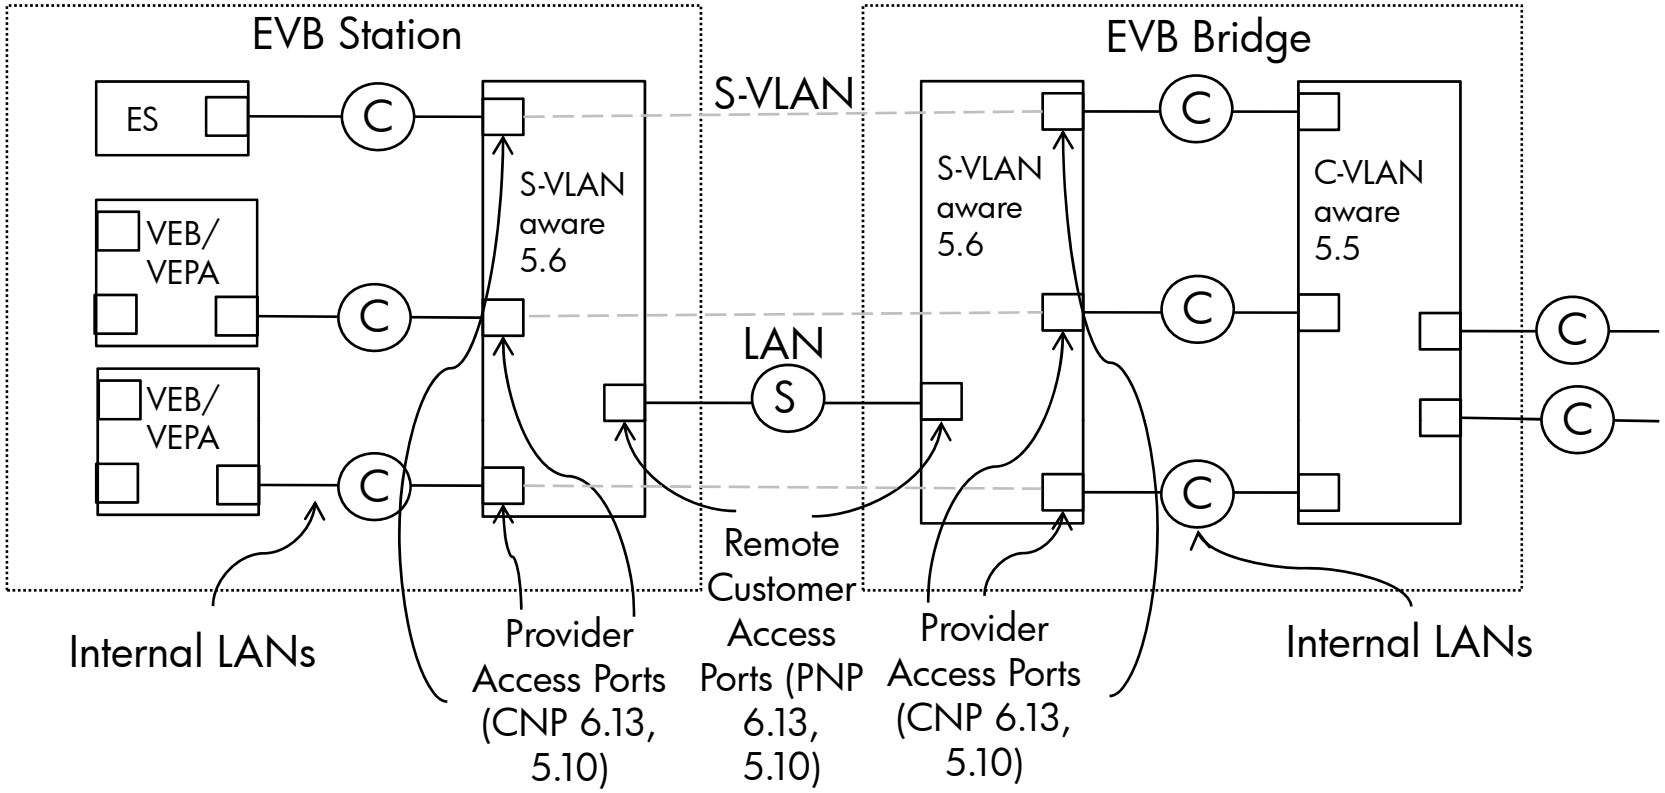

- A Port-mapping S-VLAN component (802.1Qbc) is used to implement MultiChannel
- Component disables spanning tree and MAC learning
- 1-1 relationship between CNP to C-Comp Bridge Ports of the EVB Bridge
- 1-1 relationship between CNP to VEB/VEPA Bridge Ports or End-Station ports of the EVB Station

Thoughts on Naming?

- **LAN** – the link between station and bridge
- **E-VLAN** – a channel, implemented as an S-VLAN
- **S-TAG** – the tag used to implement multichannel
- **E-Comp** – a S-VLAN aware bridge component from 802.1Qbc used to implement DCB S-VLANs
- **Station-Bridge Access Port (SBAP)** – a PNP used between a station and bridge to implement E-VLANs
- **Network Access Port (NAP)** – a CNP used to terminate an E-VLAN in a station or bridge.

Multi-Channel E-Bridge Components

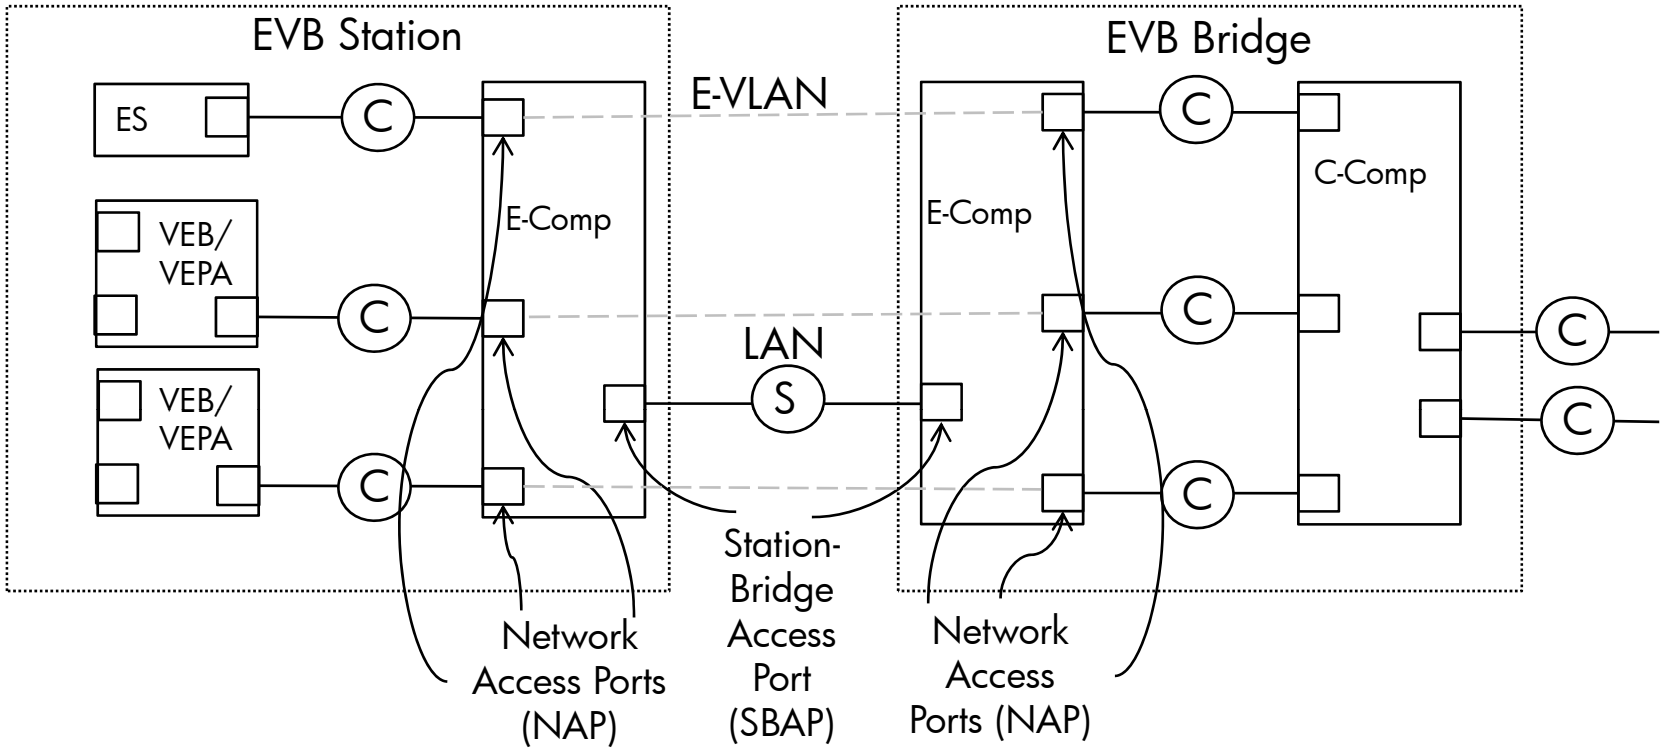

- NAP and SBAP are same definitions as CNP and RCAP
- S-TAG remains unchanged and with same name
- E-VLAN is an S-VLAN used as a channel
- E-Comp is a new S-VLAN aware component profile (new 5.x clause)

Summary

- EVB Layer Model
  - S-Comp profile for Multichannel
  - VEB/VEPA bridges
  - ETTP client of LLC (like LLDP)
  - SDU service interface between ETTP and ULPs
- EVB S-VLAN(MultiChannel) implementation naming
  - Each channel is a S-VLAN (or renamed E-VLAN)
  - A channel terminates at a CNP(PAP or renamed NAP - was vPort)
  - The LAN between Station and Bridge terminates at a PNP (or renamed SBAP)

BACKUP SLIDES