

Edge Virtual Bridging: Some Thoughts on Discovery

Joe Pelissier

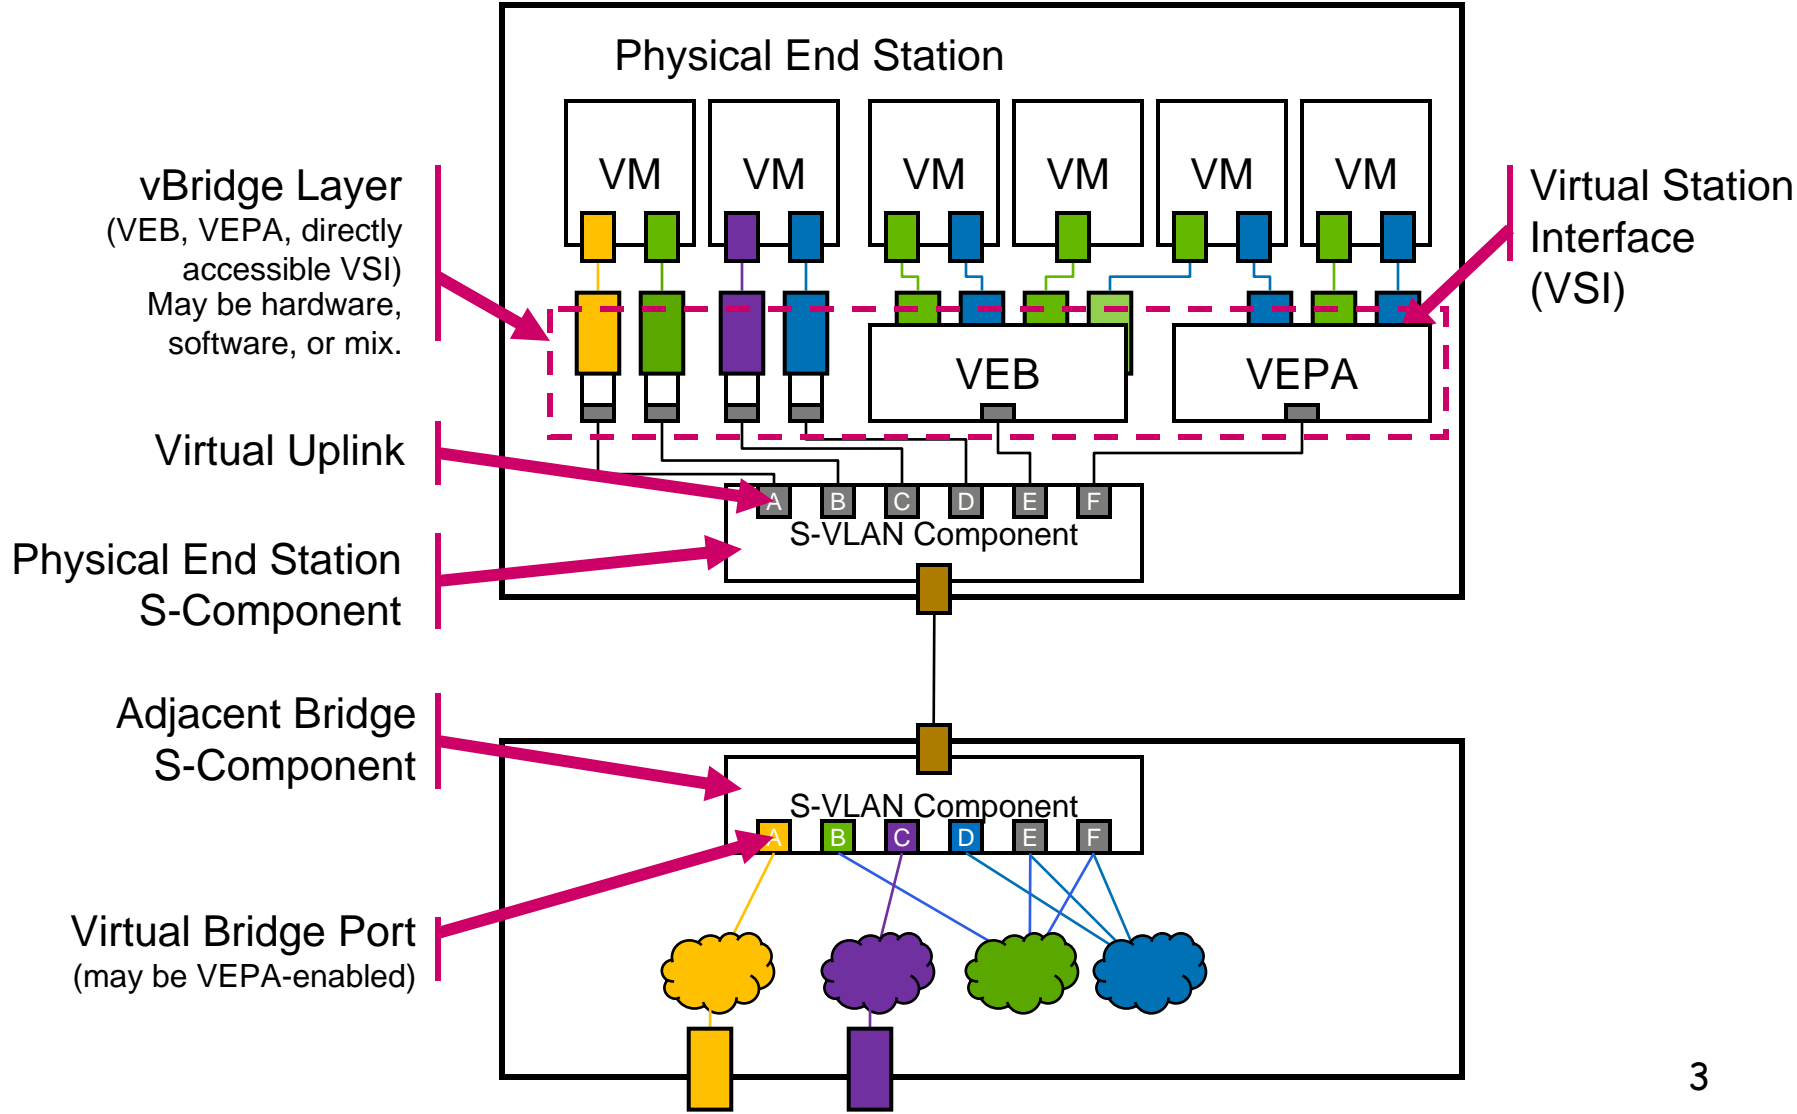
bg-pelissier-evb-discoverythoughts-0110

Introduction

- We have an active debate on whether EVB discovery should be done with LLDP only or if LLDP should be augmented with a new lightweight protocol
- Many of these presentations on either side of the debate over-simplify what needs to happen during discovery
 - Discovery is still quite simple, just not quite as simple as some have indicated
- This presentation hopes to illustrate a few of the required steps
 - What is presented is not the only possible solution, but it at least gets us to the ballpark

MultiChannel


New Anatomy and Terms



MultiChannel Discovery

Is there anybody out there?

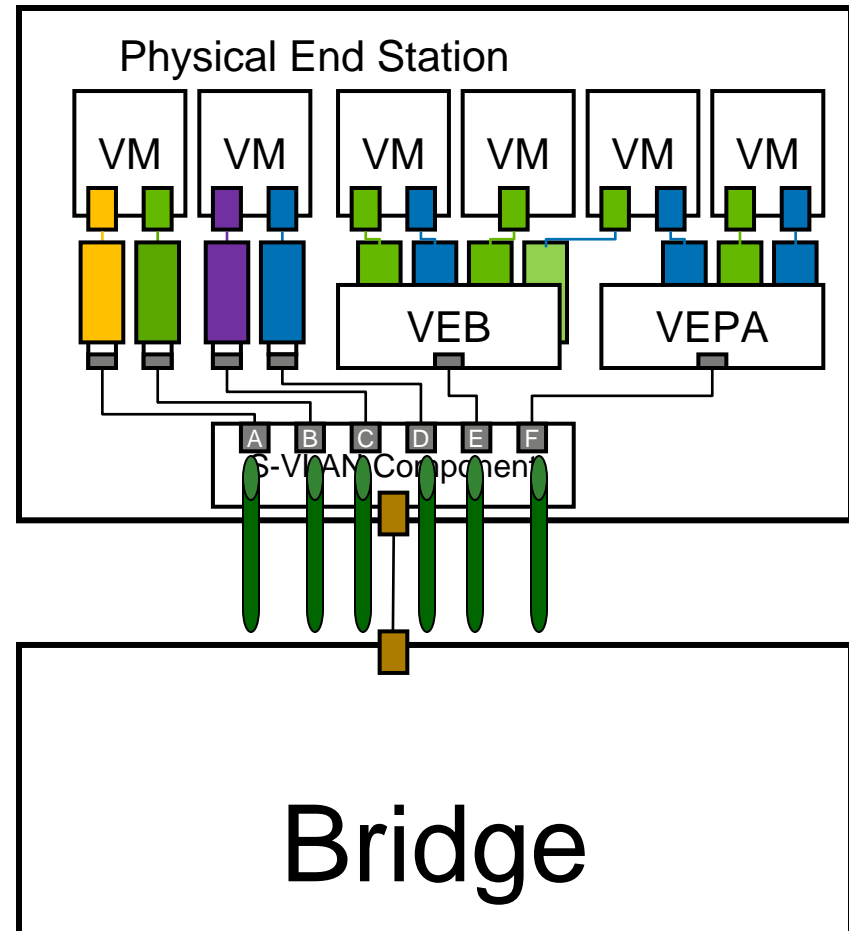
- LLDP/DCBX runs between Bridge and S_VLAN Mapping Component (SMC)
 - Bridge discovers the SMC
 - SMC discovers bridge
 - Link parameters (ETS, PFC, etc.) established



MultiChannel Discovery

Creating Channels

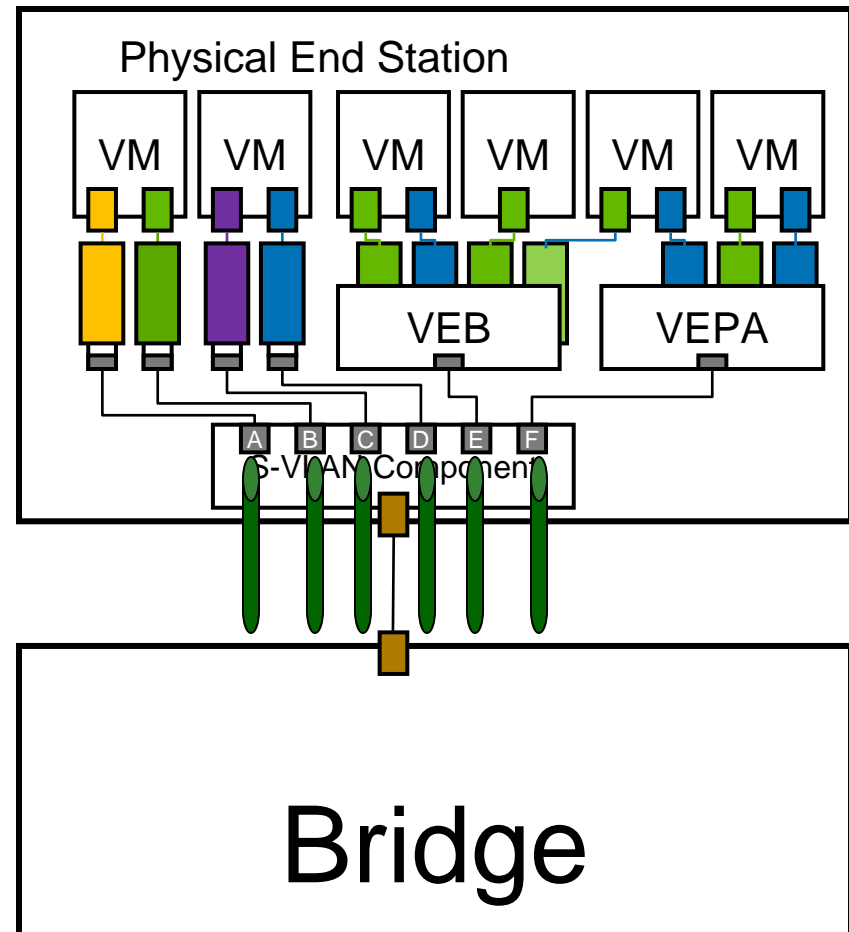
- We need six channels between the bridge and the SMC
 - Maybe a 7th channel to operate as a control channel between the SMC and Bridge?
- Each channel needs to be associated with a port (A-F) on the SMC
 - This setup is done on the control channel?



MultiChannel Discovery

Discovering the Ports, VEBs and VEPA

- Discovery of ports, VEBs, and VEPA done on a per-channel basis
 - Separate LLDP instance per channel
- In the case of ports and VEB, as single “VPort” is created in the Bridge for each channel
- In the case of VEPA, multiple VPorts are created per channel
- Same basic principle for Port Extension, just repeat as necessary



MultiChannel Discovery

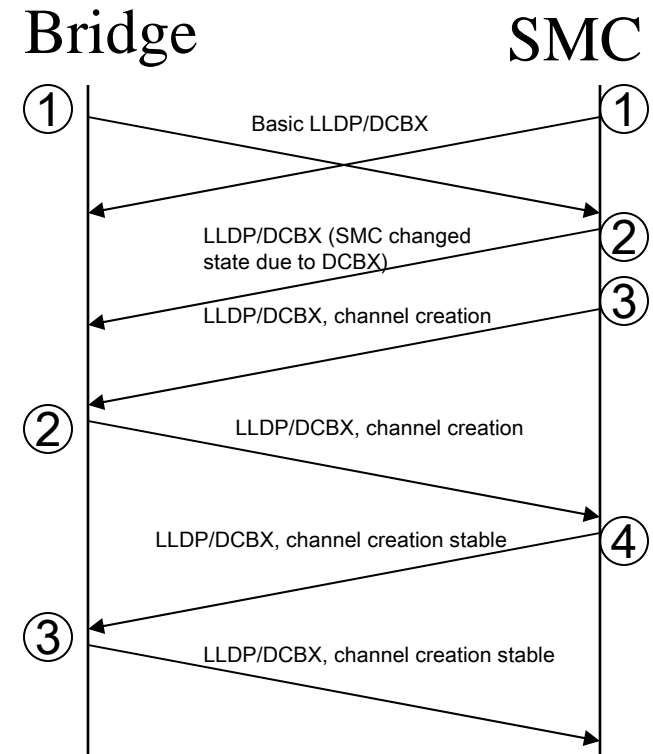
Discovering the Ports, VEBs and VEPAs

- By default, a port starts of with and can have a maximum of five credits
 - This may be set as high as ten through management
- Every time an LLDP PDU is sent, a credit is consumed
 - An LLDP PDU cannot be sent without a credit
 - Not even if some local has changed
- A new credit is bestowed once per second
 - Not to exceed the maximum mentioned above

MultiChannel Discovery

Discovering the Ports, VEBs and VEPAs

- Illustrated is a nearly best-case bring up
 - In this well-behaved scenario, the SMC generates four LLDP PDUs nearly immediately
 - If VMs, VEPAs, and VEBs come up unsynchronized, additional PDUs are generated
 - If some of these things happen to reset during the bring up process, even more PDUs are generated
- We also need to consider forwarding tables and other things yet to be invented that need to be communicated
 - Some of these are likely to need to occur serially
- We also need to consider other things that need to be communicated
 - VPort setup with VEPA
 - Forwarding tables in PEs
 - ACLs, security, and proprietary features
- We also need to consider that we are at the genesis of this new technology
 - Do we really want to be this close to a scalability limit this soon?
- In effect, we only get one PDU per second
 - IMHO, this is really scary



Thank you!