

November 15-18, 2010



Santa Clara Marriott
Santa Clara, CA

Virtual Ethernet Bridge CIM Model

John Parchem

Partner Architect Microsoft Corp.
Vice-Chair SVPC Virtual Networking

Disclaimer

- The information in this presentation represents a snapshot of work in progress within the DMTF.
- This information is subject to change. The Standard Specifications remain the normative reference for all information.
- For additional information, see the Distributed Management Task Force (DMTF) Web site.





Enhancements to virtual Ethernet bridge profiles and schema

- DSP1097 Virtual Ethernet Switch Profile and DSP1050 Ethernet Port Resource Virtualization are DMTF standards
- SVPC virtual networking group was formed to enhance the virtual network CIM model to support additional use cases supported by the IEEE 802.1Qbg work including Virtual Edge Bridge (VEB) and Virtual Ethernet Port Aggregator (VEPA) capabilities and support of PCI SRIOV embedded bridges) and to develop a schema for a Network Port Profile.



Enhancements to virtual Ethernet bridge profiles and schema

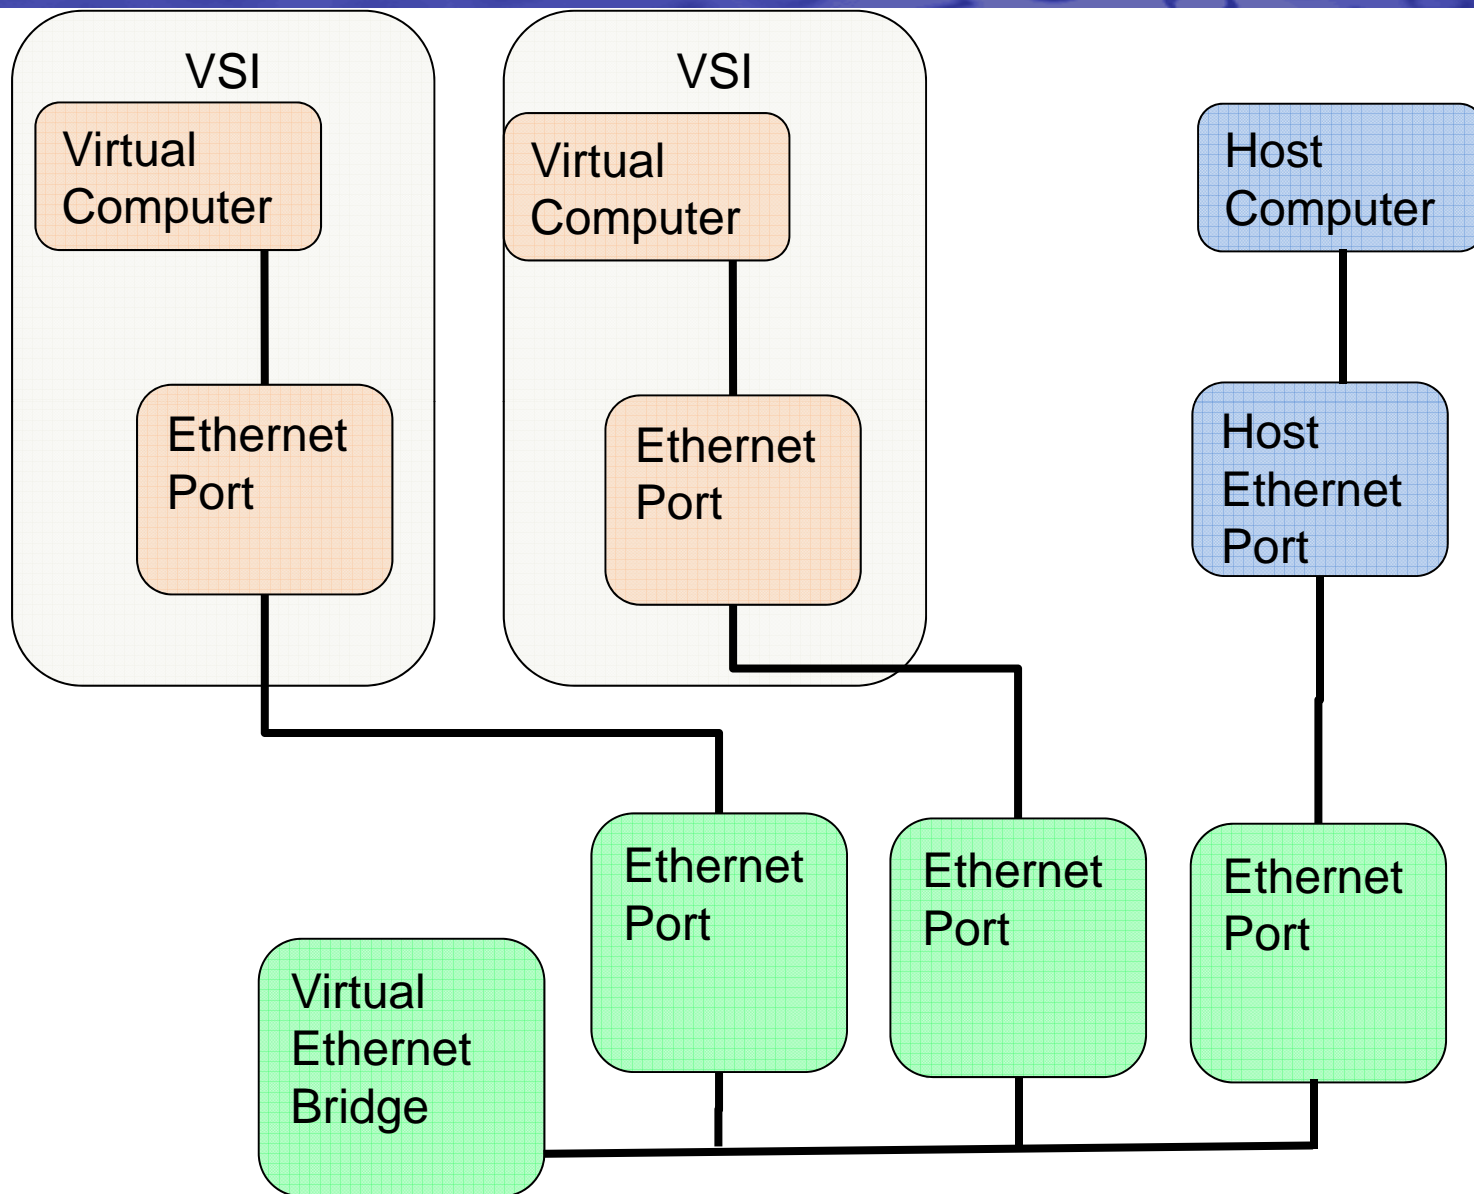
- CIM Schema 2.27 includes changes to CIM_VirtualEthernetSwitchSettingData and CIM_VirtualEthernetPortAllocationSettingData to support Edge Bridge (VEB) and VEPA capabilities and support of SRIOV embedded bridges.
- V1.1 of DSP1050 and DSP1097 are being developed to support the above described uses case and normalize the networking terms used in the document to those described in IEEE 802.1.



Profiles Under Development

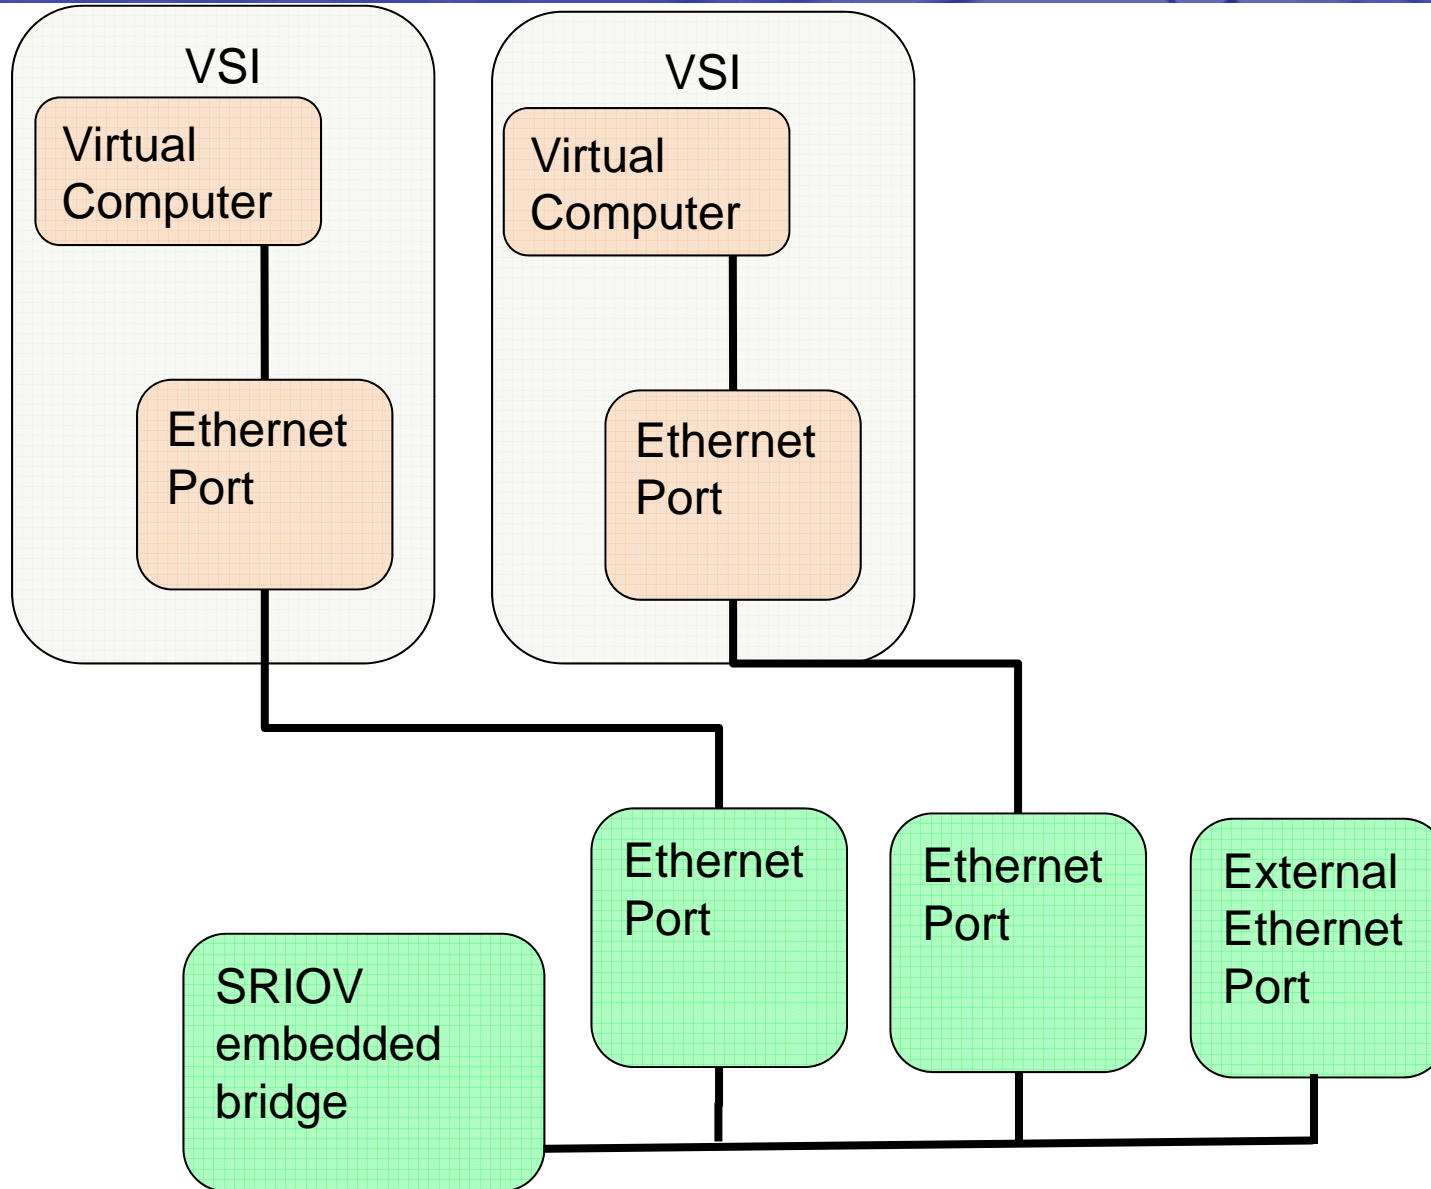
- DSP 1097 V1.1 – Virtual Ethernet Switch Profile
 - specializes the autonomous Virtual System Profile. This profile defines the minimum top-level object model needed to define a virtualization system's internal Ethernet bridge.
- DSP 1050 V1.1 – Ethernet Port Resource Virtualization Profile
 - specializes the abstract Resource Allocation Profile and the Allocation Capabilities Profile to specify the allocation and management of a host computers Ethernet network access, in support of virtual computer systems ability to access external and internal network.

Virtual Ethernet Bridge V1

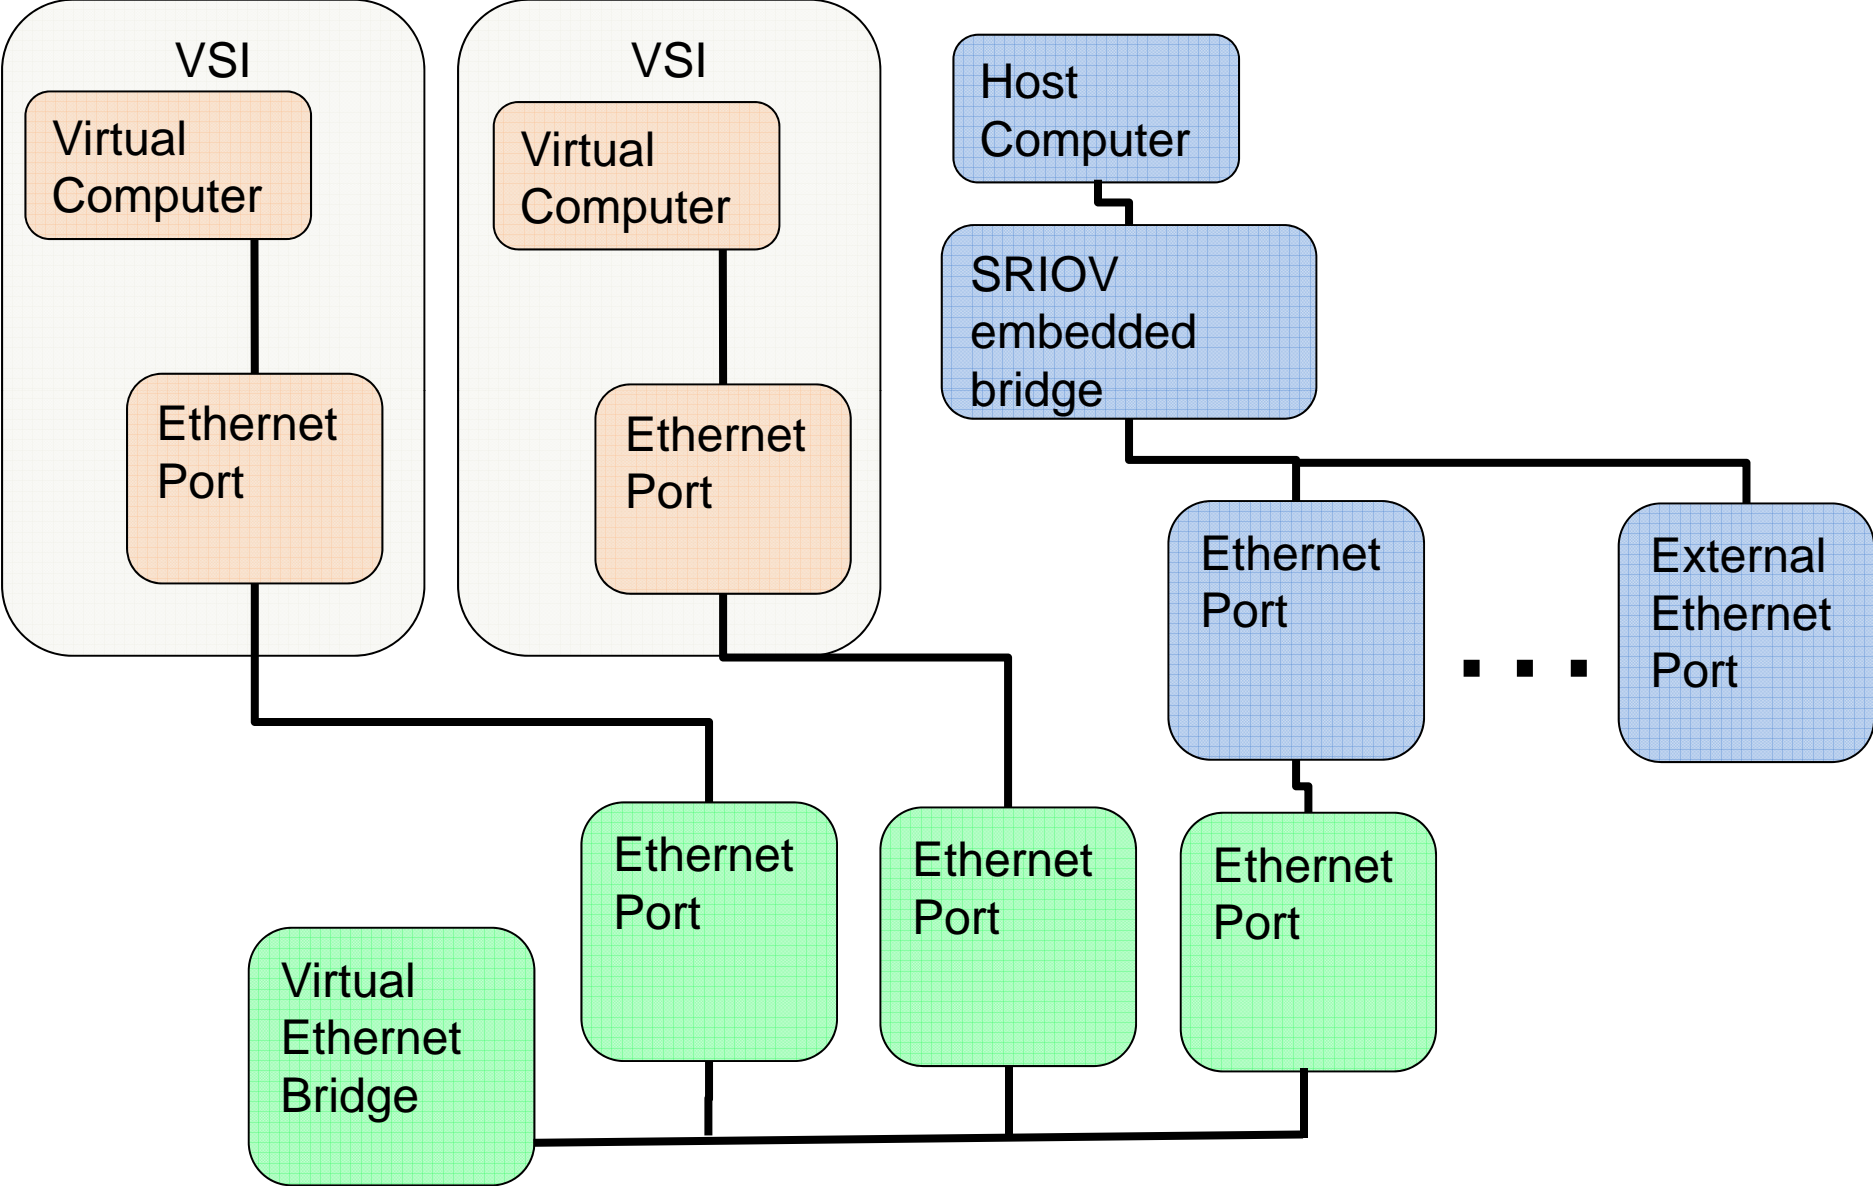


Virtual Ethernet Bridge V1.1a

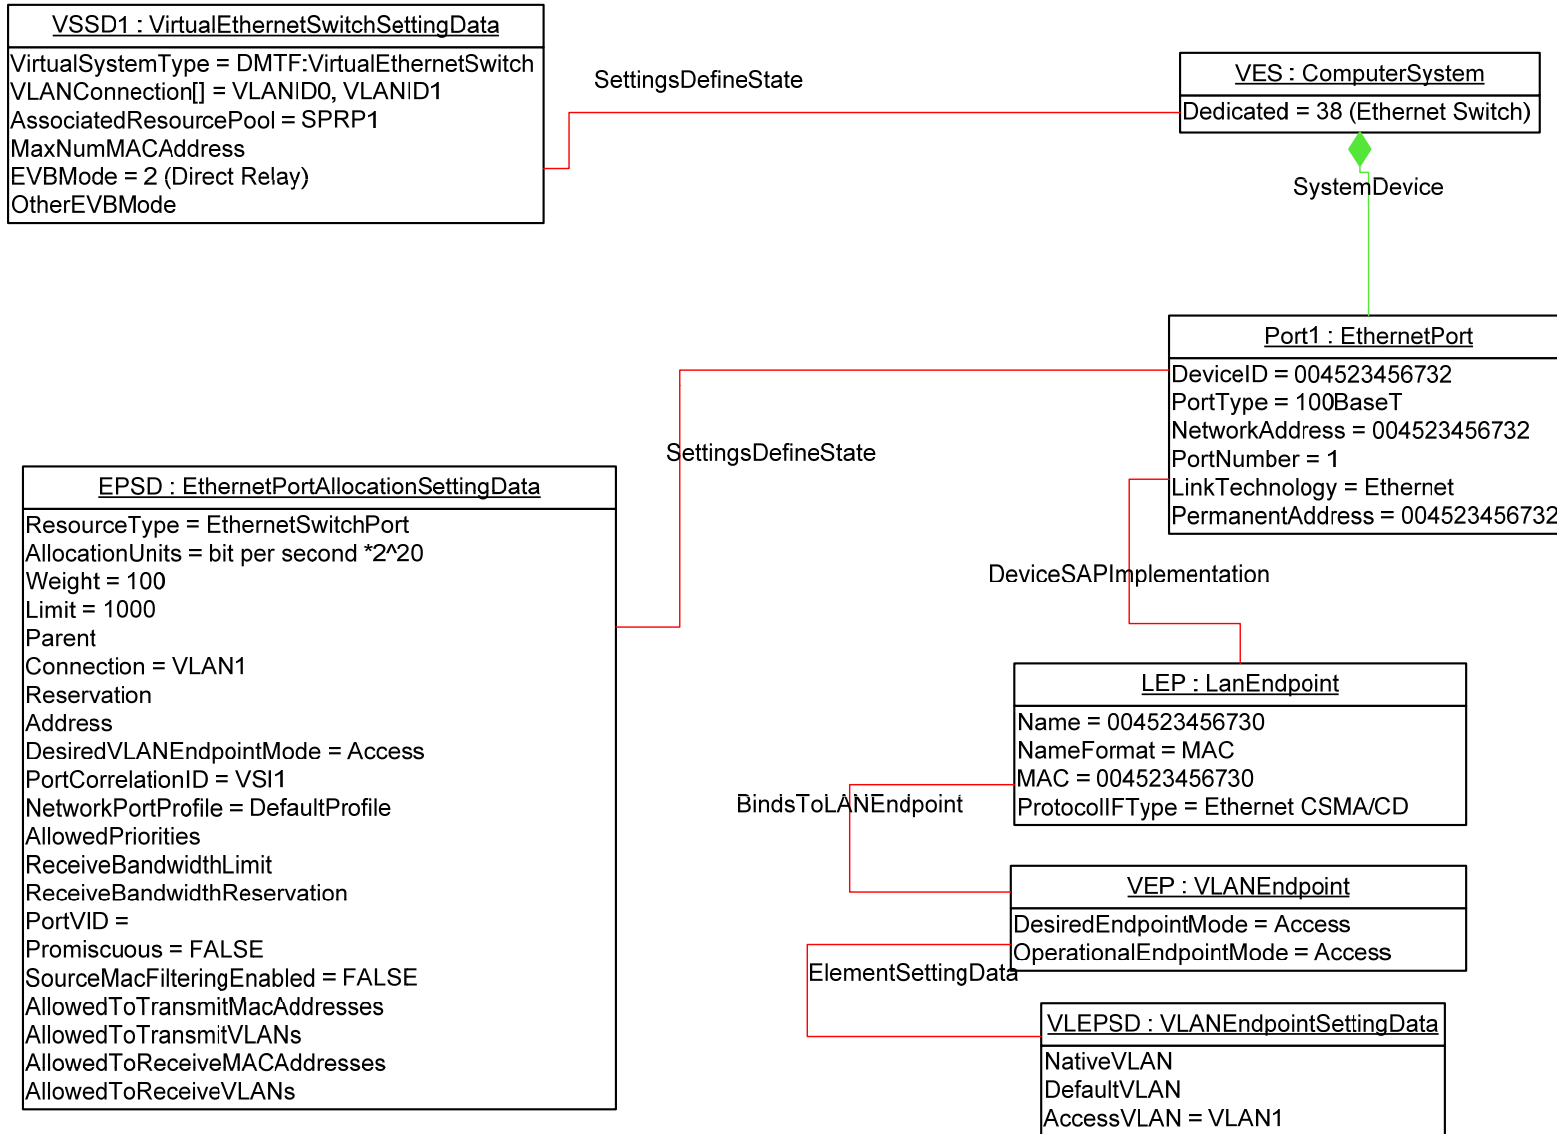
SRIOV Embedded Bridge



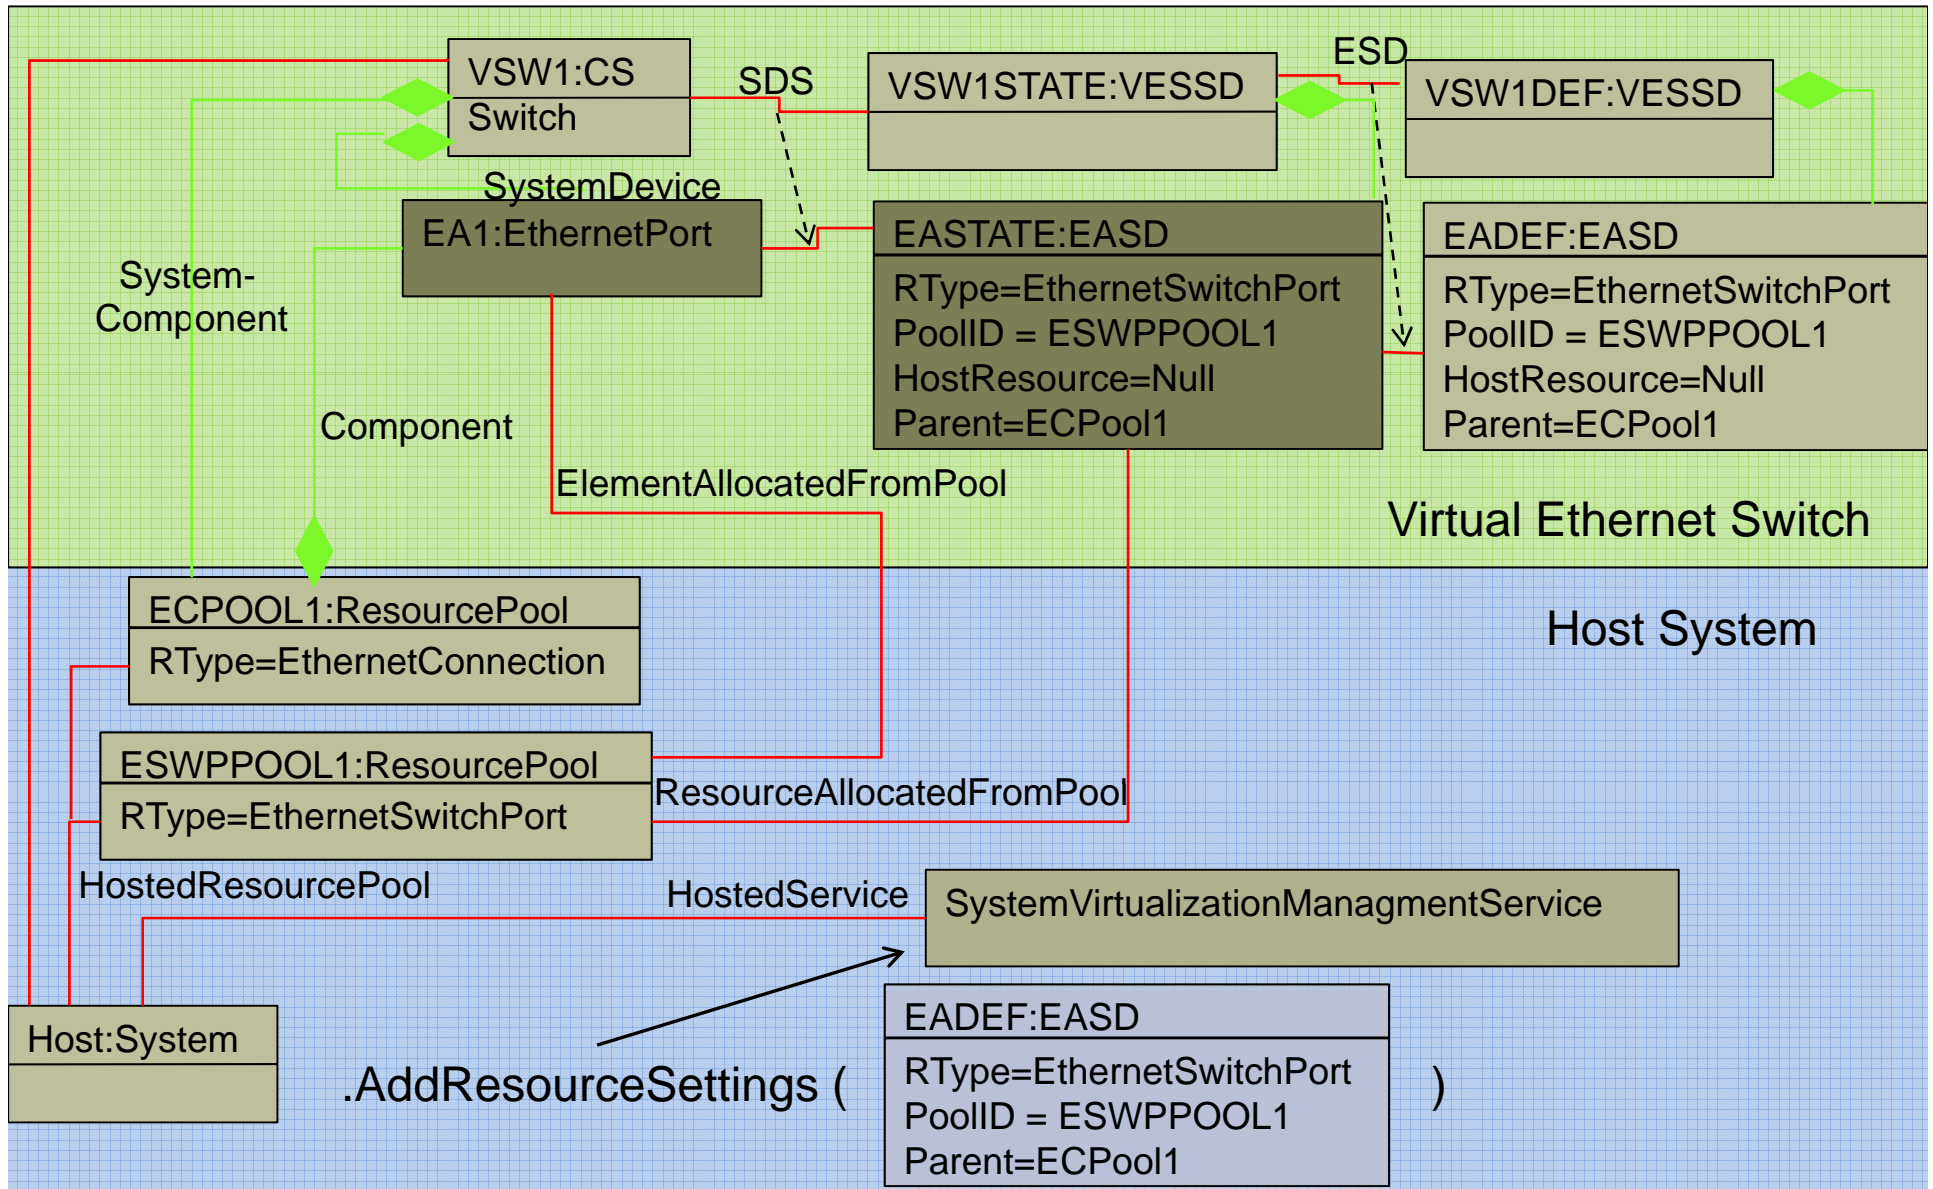
Virtual Ethernet Bridge V1b



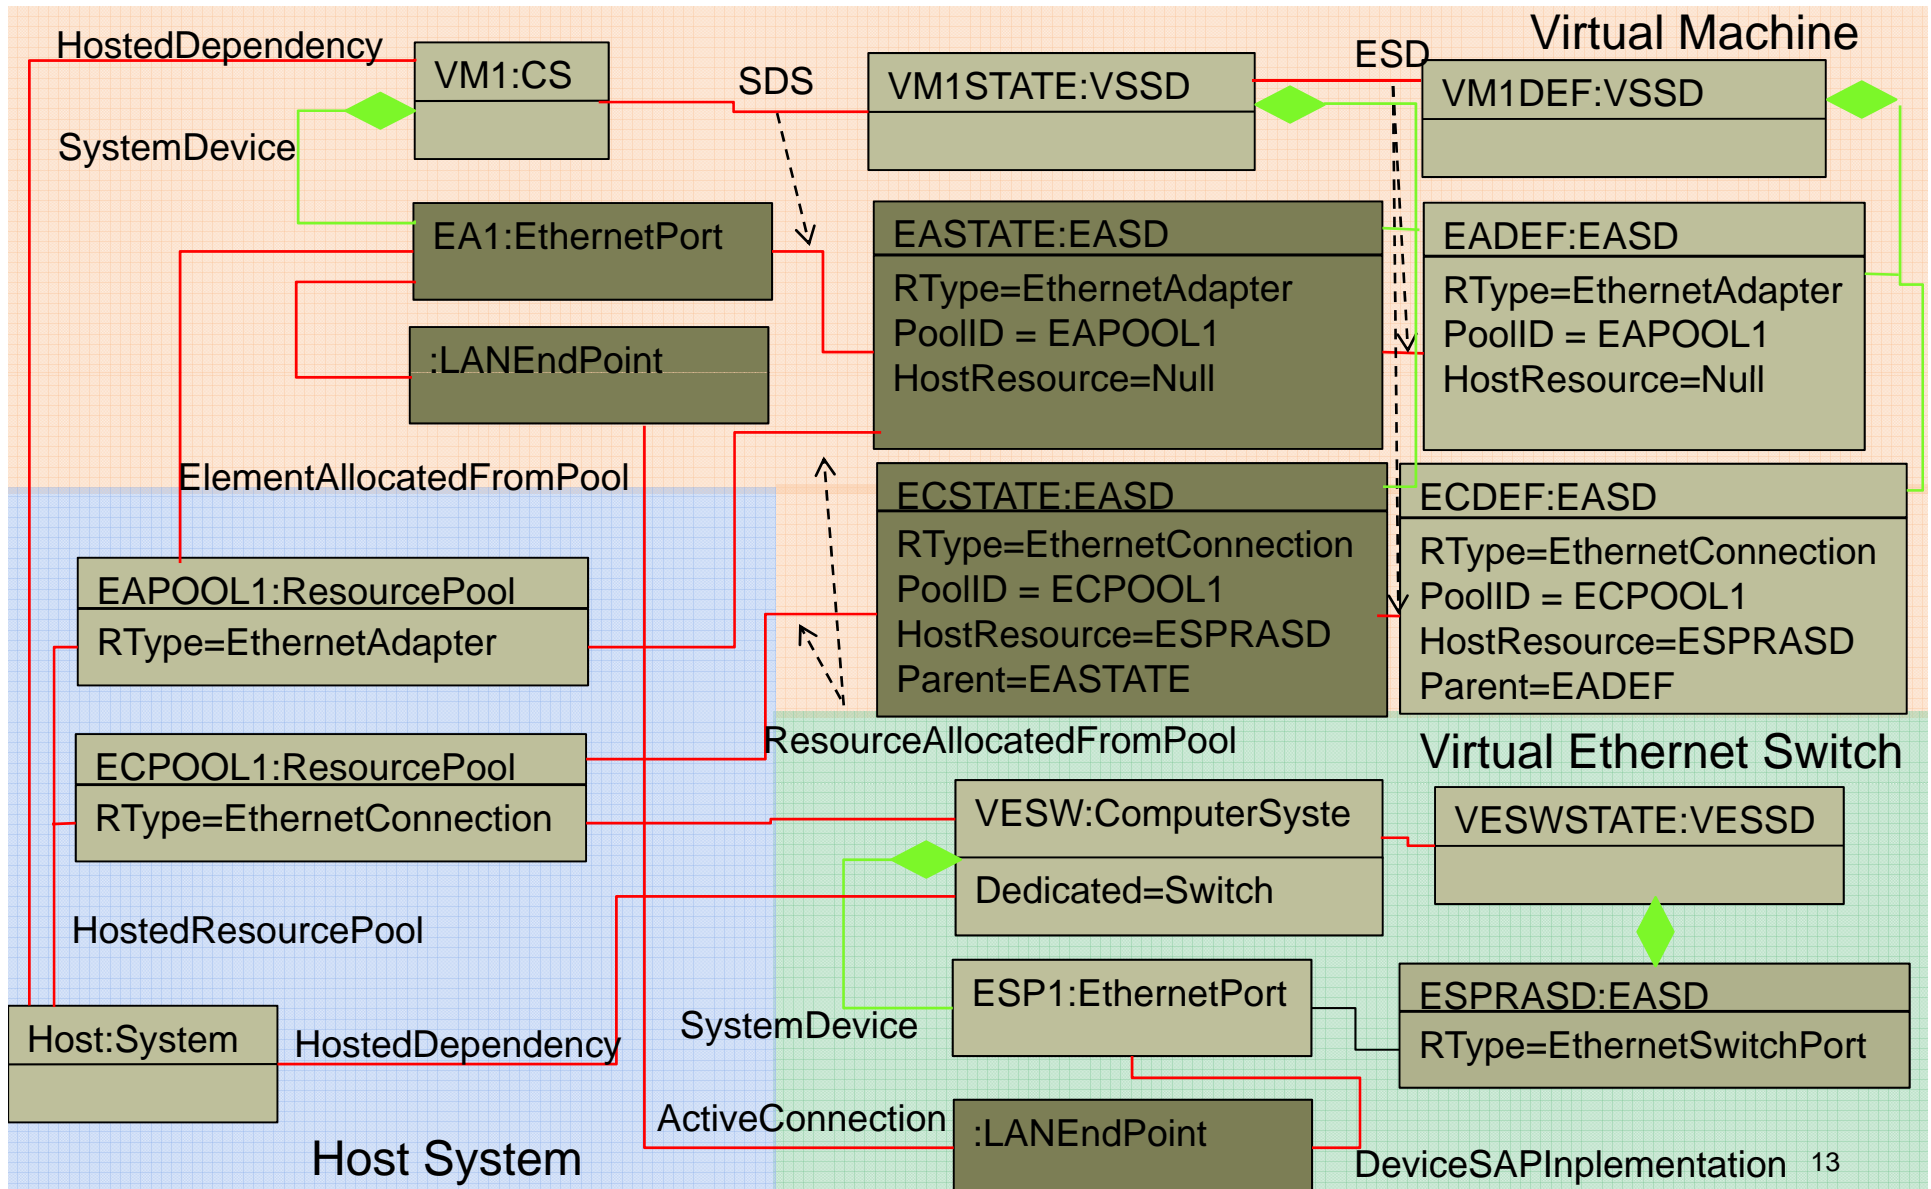
Ethernet Bridge Virtualization



Allocation of Ethernet Switch Port to Virtual Switch



Allocation of Ethernet Adapter and Ethernet Connection



VESSD: VirtualEthernetSwitchSettingData

- VirtualSystemType= “DMTF:Ethernet Switch”
- EVBMode - Describes the Ethernet virtual bridge mode that this port is operating in. Direct Relay and Reflective Relay are as defined in IEEE 802.1Qbg (VEB mode and VEPA mode). Direct is used to describe an outbound port that is directly connected to or is part of a hardware Ethernet NIC
- AssociatedPoolID[] :Array of Ethernet Connection pools associated with the switch
- VLANConnection[]: A list of VLANs this switch can access (results in a set of allocation capabilities on the switch)
- MacNumMACAddress: Number of MAC addresses that this switch can learn



CIM_EthernetPortAllocationSettingData: Ethernet Switch Port

Page 1

- ResourceType = 30 (Ethernet Switch Port)
- AllocationUnits: PUnit bit per second
- Reservation,Limit: bandwidth in AllocationUnits
- HostResource: URI to an existing host Ethernet port for host access to network
- DesiredEndPointMode: the VLAN endpoint mode on the Ethernet Switch Port
- Connection: depending on the VLANEndpoint's DesiredEndpointMode either Access VLANID or array of Trunked VLANIDs
- PoolID: specifies desired pool for allocation of the Ethernet Switch Port
- Address: MAC address

- PortCorrelationID - typically a VSI Instance ID
- NetworkPortProfile – used to map a port to an existing network port profile
- AllowedPriorities[] - allowed values of 802.1Q PCP bits that this port is allowed to transmit
- ReceiveBandwidthLimit - Maximum receive bandwidth this port is allowed to use
- ReceiveBandwidthReserve – Minimum receive bandwidth allocated to this port
- PortVID – VLAN ID that is used as the “DefaultVLAN” (used to mark traffic when no VLAN ID is specified)
- Promiscuous – True if the port passes all traffic without destination MAC address filtering

- SourceMacFilteringEnabled – If enabled the filter is used to match either the ports MAC address or one of the pairing of the property value in the indexed arrays AllowedToTransmitMACAddresses and AllowedToTransmitVLANs
- AllowedToTransmitMACAddresses - indexed array of MAC addresses this port is allowed to transmit
- AllowedToTransmitMACAddresses - indexed array of MAC addresses this port is allowed to transmit
- AllowedToReceiveVLANs and AllowedToReceiveMACAddresses - indexed arrays of MAC addresses and VIDs to support filtering on MAC,VID pairs.

Further Reference

- Further information: DMTF System Virtualization, Partitioning and Clustering WG
- For questions:
 - johnp@microsoft.com John Parchem Microsoft
 - tm-redundancy@dmtof.org SVPC workgroup
 - svpc-iov@dmtof.org SVPC Virtual Networking
- Referenced Profiles:
 - DMTF DSP1014, *Ethernet Port Profile 1.0*
 - DMTF DSP1041, *Resource Allocation Profile 1.1*
 - DMTF DSP1042, *System Virtualization Profile 1.0*
 - DMTF DSP1043, *Allocation Capabilities Profile 1.0*
 - DMTF DSP1057, *Virtual System Profile 1.0*
 - DMTF DSP1050, *Ethernet Port Virtualization Profile 1.0*
 - DMTF DSP1097, *Virtual Ethernet Switch Profile 1.0*