EVB PAR Proposal

September 2009

Paul Congdon (HP) ptcongdon@ucdavis.edu

Agenda

- Edge Virtual Bridging Status
- Base EVB PAR Discussion

EVB Yahoo Group
http://tech.groups.yahoo.com/group/evb/

- Unofficial ad hoc group working to develop concepts and proposals related to Edge Virtual Bridging for consideration by the IEEE 802.1 working group.

- Membership
  - 100+ members have joined Yahoo group
  - Affiliated with 20+ companies (including server, switch, NIC, hypervisor & OS companies)
- Weekly Conference Calls
  - Tuesdays 1PM Central
  - Since February 20th 2009
  - 25-30 Attendees Weekly

Approaches

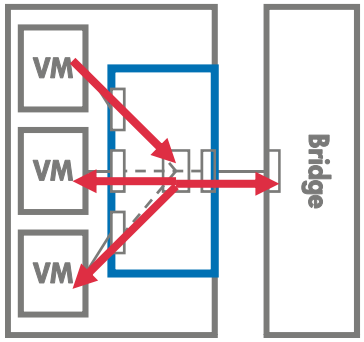

Virtual Ethernet Bridge (VEB)

MAC+VID to steer frames

- Emulates 802.1 Bridge
- Existing implementations (vSwitch, SR-IOV switch)
- Works with all existing bridges
- No changes to existing frame format.
- Limited bridge visibility
- Limited feature set
- Best local performance.
- Legacy, pervasive solution

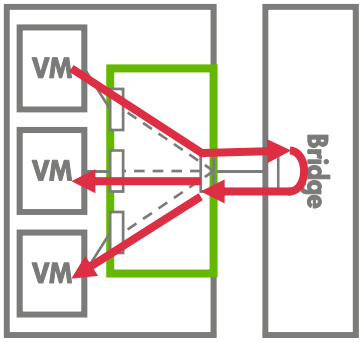

Virtual Ethernet Port Aggregation (VEPA)

MAC+VID to steer frames

- Exploits 802.1 Bridge
- Works with many existing bridges (hairpin)
- No changes to existing frame format.
- Full bridge visibility
- Access to bridge features
- Constrained performance
- Leverages VEB resources

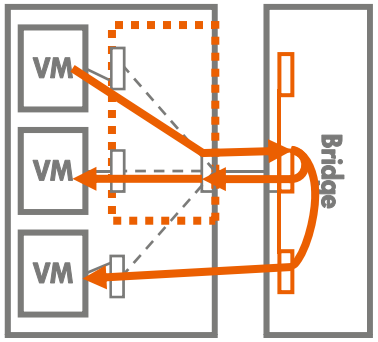

Multichannel

uses tag for remote ports

- Exploits Provider Bridge
- Similarities to Remote Service Interface
- Uses existing frame formats (S-tags).
- Creates bridge virtual ports
- Defines restricted S-Component
- Access to bridge features
- Adjacent bridge multicast replication (constrained performance)

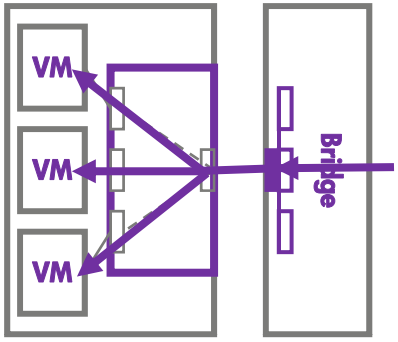

Remote Replication

uses tag to replicate packets

- Extends 802.1 Bridge
- Optimizes multicast delivery
- Defines new tag format
- Defines new name space

Discovering VEBs and VEPAs

VEB mode allows VM-to-VM traffic with limited policy enforcement.

VEPA mode forces all traffic to fully-capable edge for full policy enforcement.

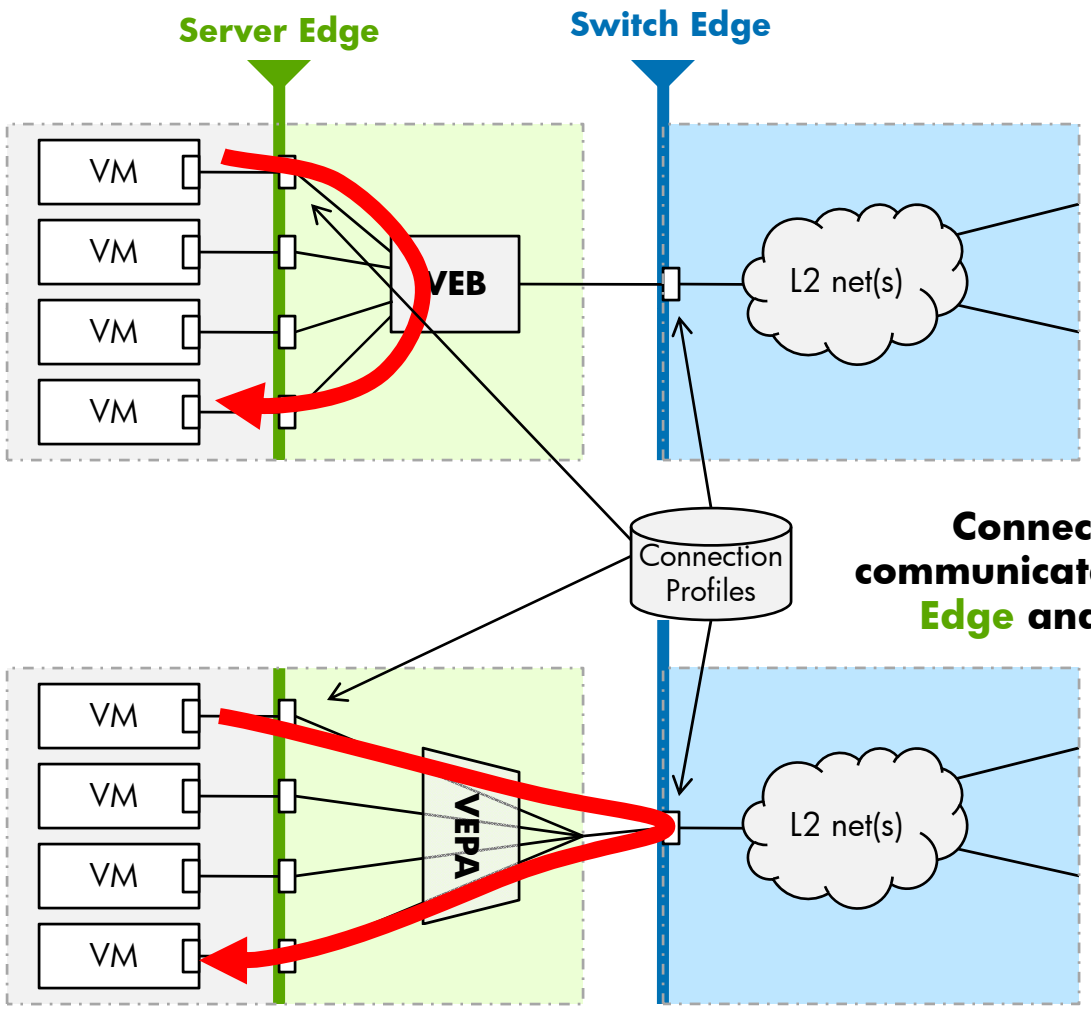

Connection settings communicated to both Server Edge and Switch Edge.

Multi-Channel: Allows VEB, VEPA, & dedicated links on the same switch port

Provides the ability to support a vSwitch and VEPA on the same switch port (with a single NIC)

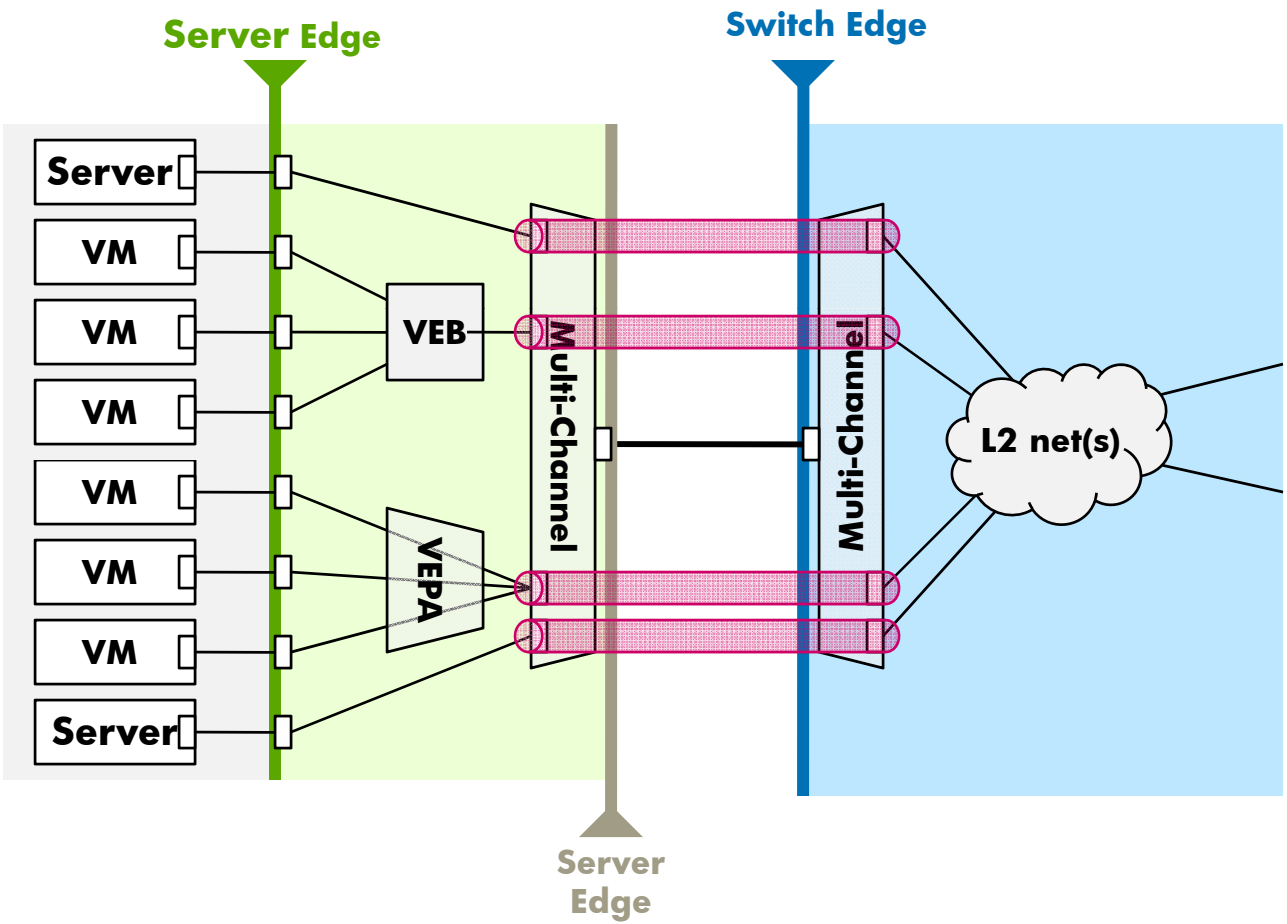

The type of link (VEB, VEPA, or direct) could be specified as part of the connection profile.

Discovering EVB Environment

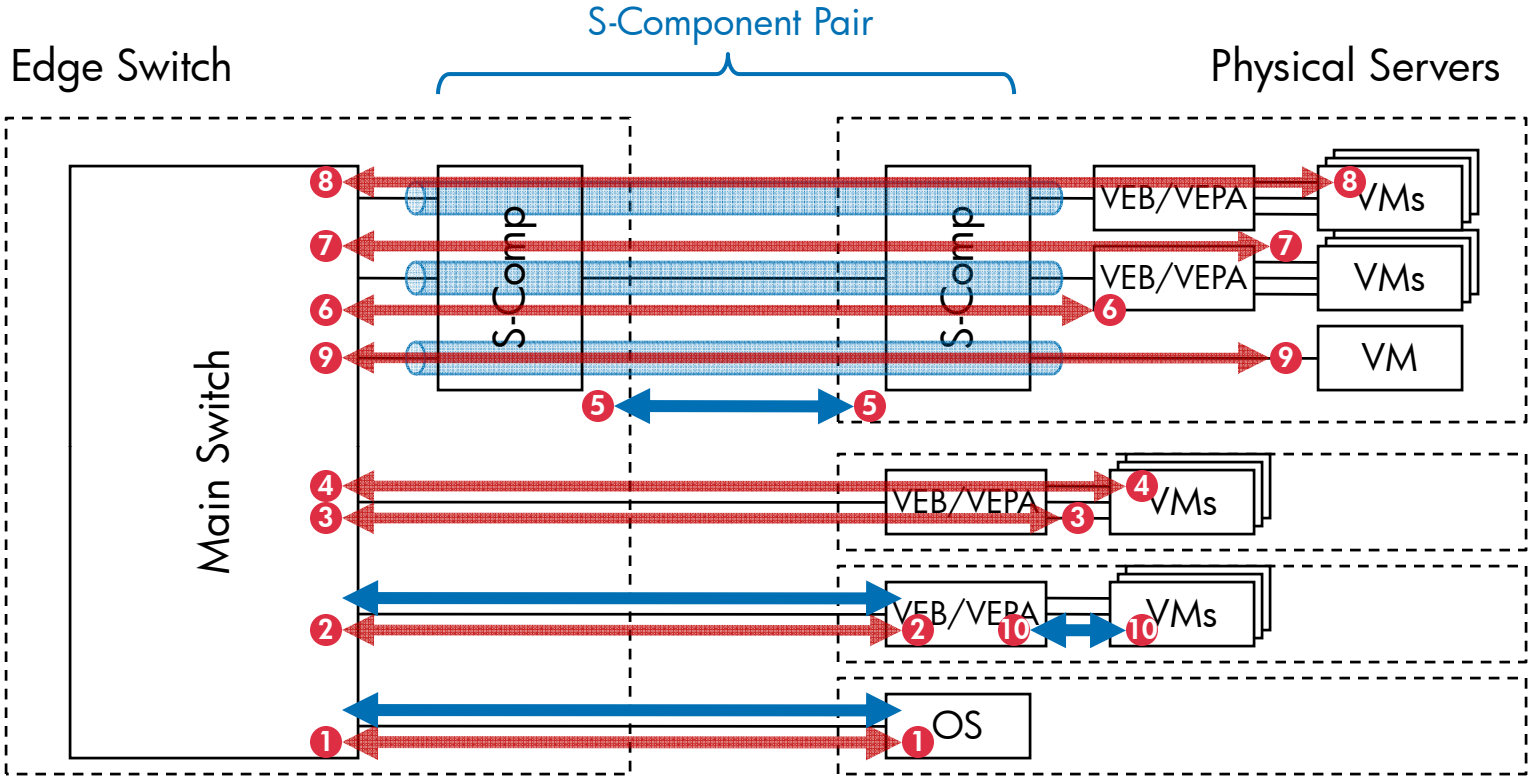

Trace	LLDP Address	Appropriate Items
1,2,5,10	Nearest bridge	LLDP, DCBX, Multi-Channel
All	Nearest customer bridge	LLDP, AMPP, VEB/VEPA

Specification Needs for EVB Operation

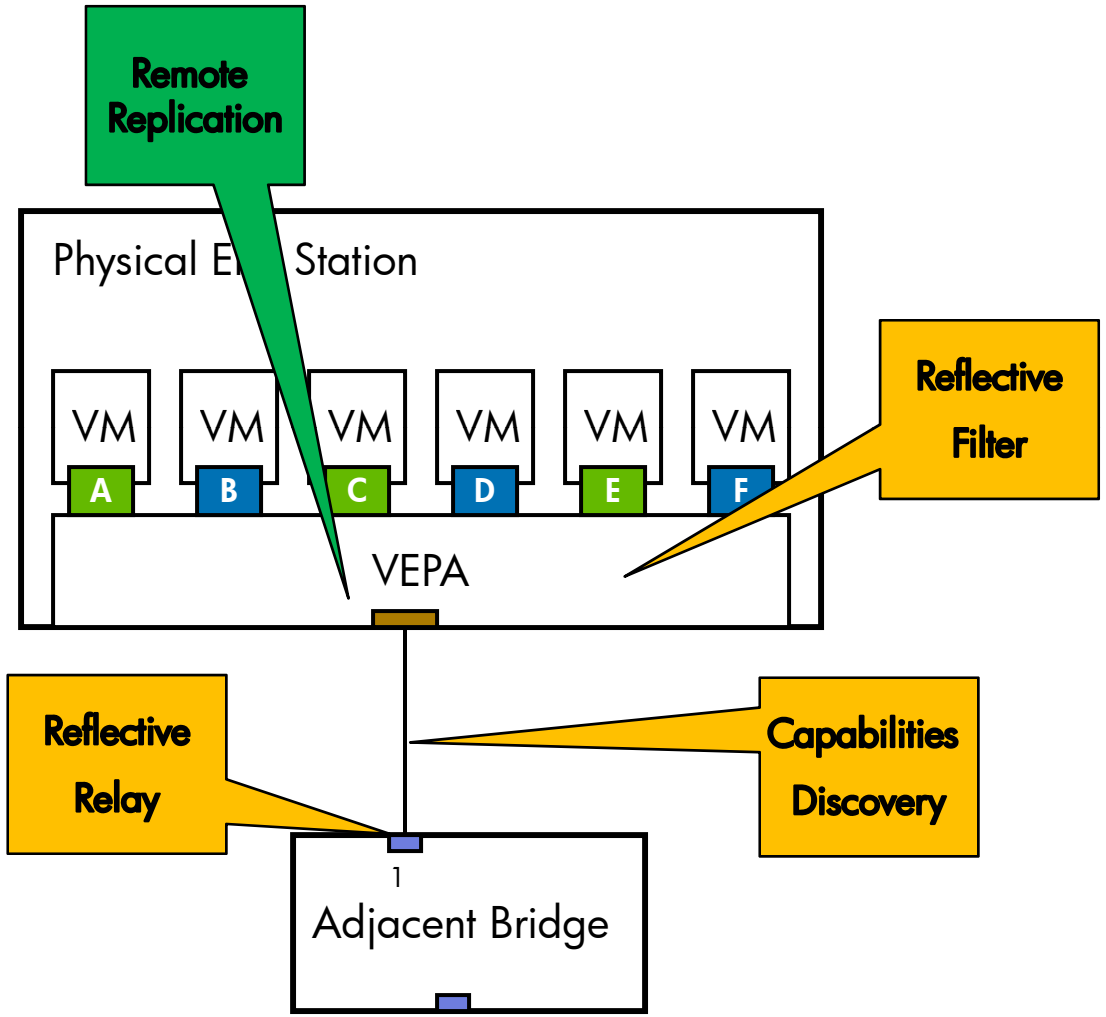

Static VEPA Address Table

DST MAC	VLAN	Copy To (ABCDEF)
A	1	100000
B	2	010000
C	1	001000
D	2	000100
E	1	000010
F	2	000001
Bcast	1	101010
Bcast	2	010101
MulticastC	1	101010
Unk Mcast	1	100010
Unk Mcast	2	010101
Unk Ucast	1	000000
Unk Ucast	2	000000

Base EVB PAR Proposal

- Current text located at: http://www.ieee802.org/1/files/public/docs2009/new-evb-congdon-evbPar5C-0709-v01.pdf

- Wording worked at July plenary PM meeting

Base EVB PAR - Scope

5.2 Scope: This standard specifies protocols, procedures, and managed objects that:

- Provides for the discovery, configuration, and control of a pair of direct-attached, limited-function S-components to extend the services of a customer bridge to remote ports and enable coexistence of multiple services on station-resident ports.(e.g. port aggregation services, embedded bridging)
- Provides for discovery, configuration, and control of a Reflective Relay Service for a bridge port when it is connected to a Port Aggregation Service.
- Defines the requirements for, and operation of, a Port Aggregation Service required to allow the loop free operation of the Reflective Relay Service.
- Provides for discovery of, and coordinated configuration of, station embedded Port Aggregators and station embedded bridging.

Base EVB PAR – Purpose and Need

5.4 Purpose:

- The purpose of this standard is to allow multiple virtual stations to share a common bridge port to obtain the services of bridge relay. The standard will enable coordinated configuration and management of bridge services for virtual stations.

5.5 Need for the Project: Station (desktop and server) virtualization

- Introduces a proliferation of virtual stations that share access to a network through an embedded bridge. The embedded bridge in a virtual station host and bridges in the adjacent network may be under different management domains.
- Network administrators may desire an extension of the adjacent bridge capabilities to the virtual station ports. A reflective relay service is needed so that inter-virtual-station traffic can be exposed to the relay in the adjacent bridge as well as its associated services (e.g., security, statistics, etc.).
- A pairing of limited-functionality S-Components is needed to allow for multiple virtual links between an adjacent bridge and a virtual station host so that the host can support multiple services (port aggregation, embedded bridging, dedicated bridge link).

Backup

802.1X in the EVB Environment

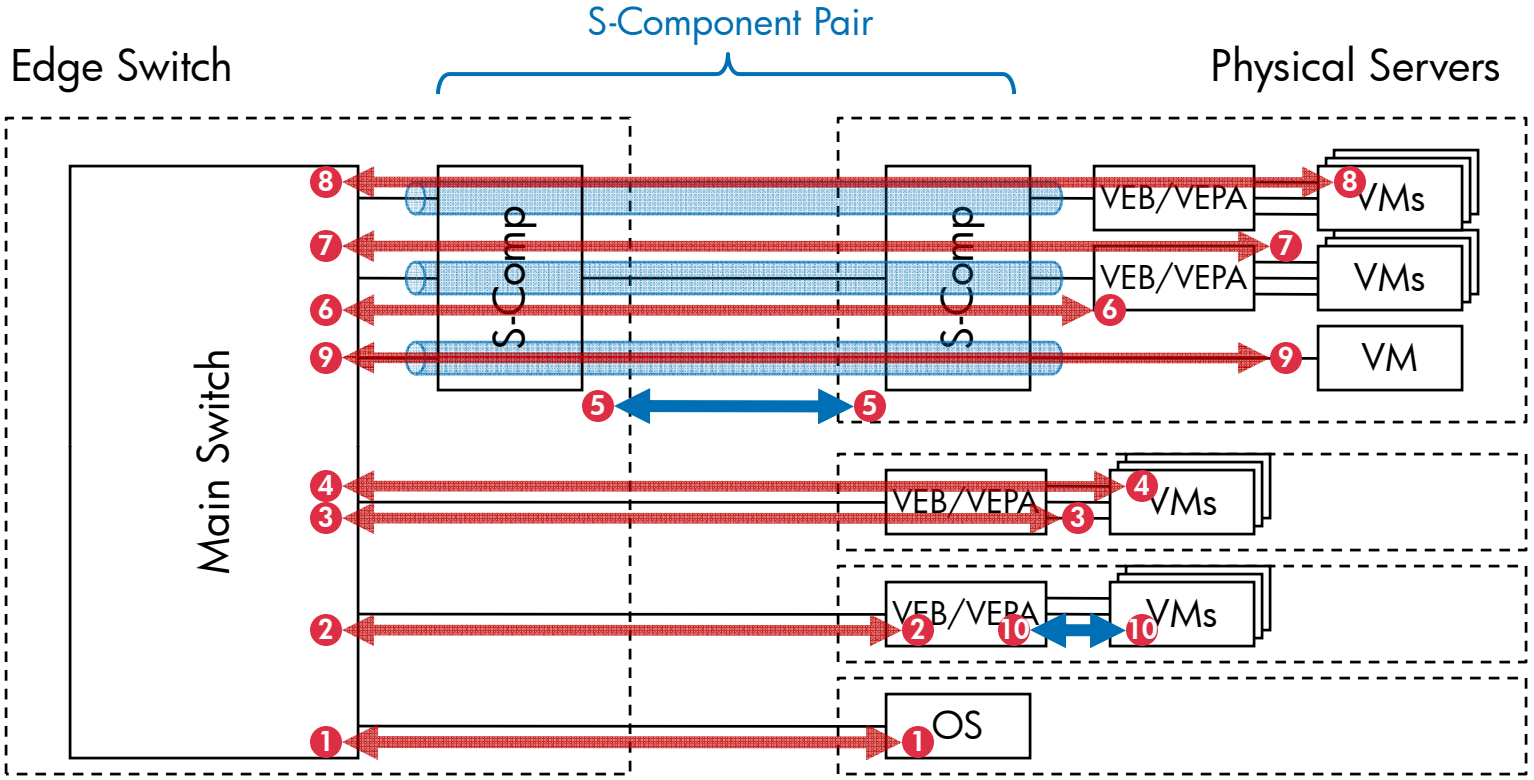

Trace	LLDP Address	Appropriate Items
5	Nearest bridge	authentication for shared MACSec
10	Nearest bridge	VEB/VEPA authenticator
7, 9, 3	Nearest customer bridge	proxy supplicant model
8, 4, 1	Nearest customer bridge	existing VM authentication model
6,2	Nearest customer bridge	VEB/VEPA supplicant