

# Possible EVB-related TLVs

v 09

Based on multiple presentations and numerous discussions in the EVB conference calls. Does not yet reflect full consensus.


January 21, 2010

# Summary

## Proposed EVB-related TLVs

- Multichannel
  - Negotiate whether to use multichannel mode
  - Share enough information to setup the channels
- EVB Discovery
  - Negotiate aggregation mode (VEB, VEPA, etc.)
  - Negotiate whether to use VSI discovery protocol
  - Negotiate whether Hypervisor Authentication is supported
- VSI Discovery/Configuration
  - Announce the arrival status of a (Virtual) Station Interface and communicate information to allow the edge bridge to retrieve the appropriate configuration for the connection.

# Multi-Channel: Multiple isolated channels sharing a single physical link.

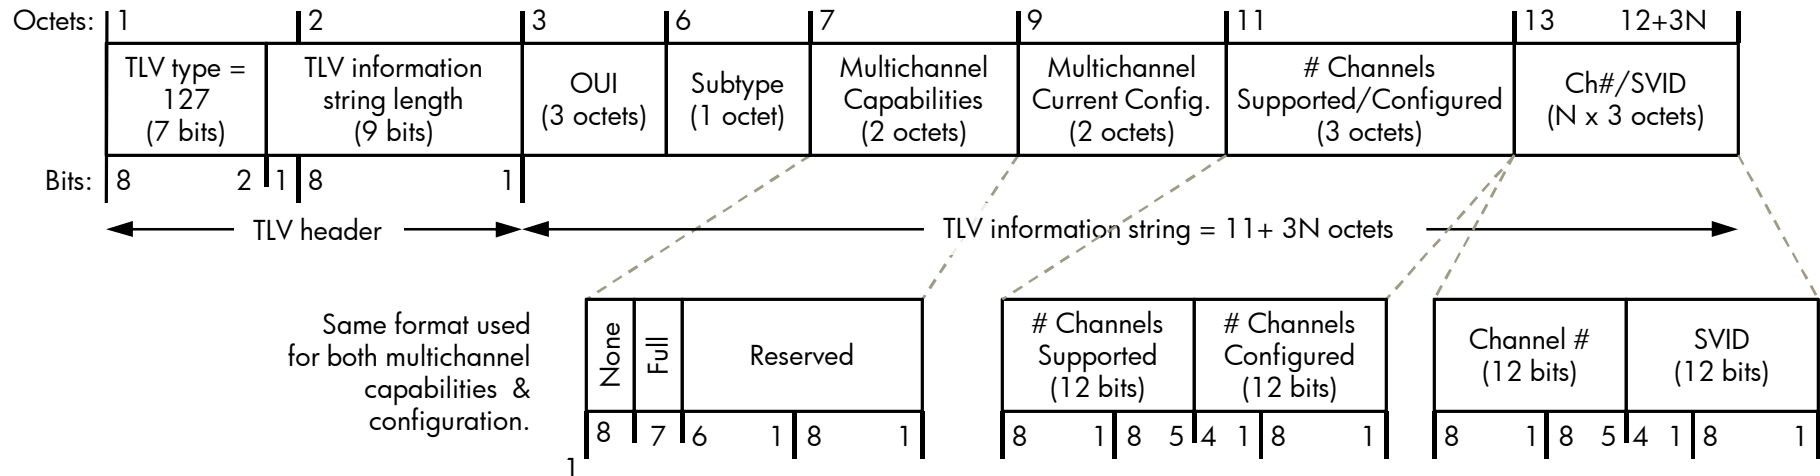


Multi-channel TLV establishes the basic channels.

The specific configuration of individual channels is handled by the EVB and VSI TLVs.

# Proposed Multichannel TLV (previously discussed format)

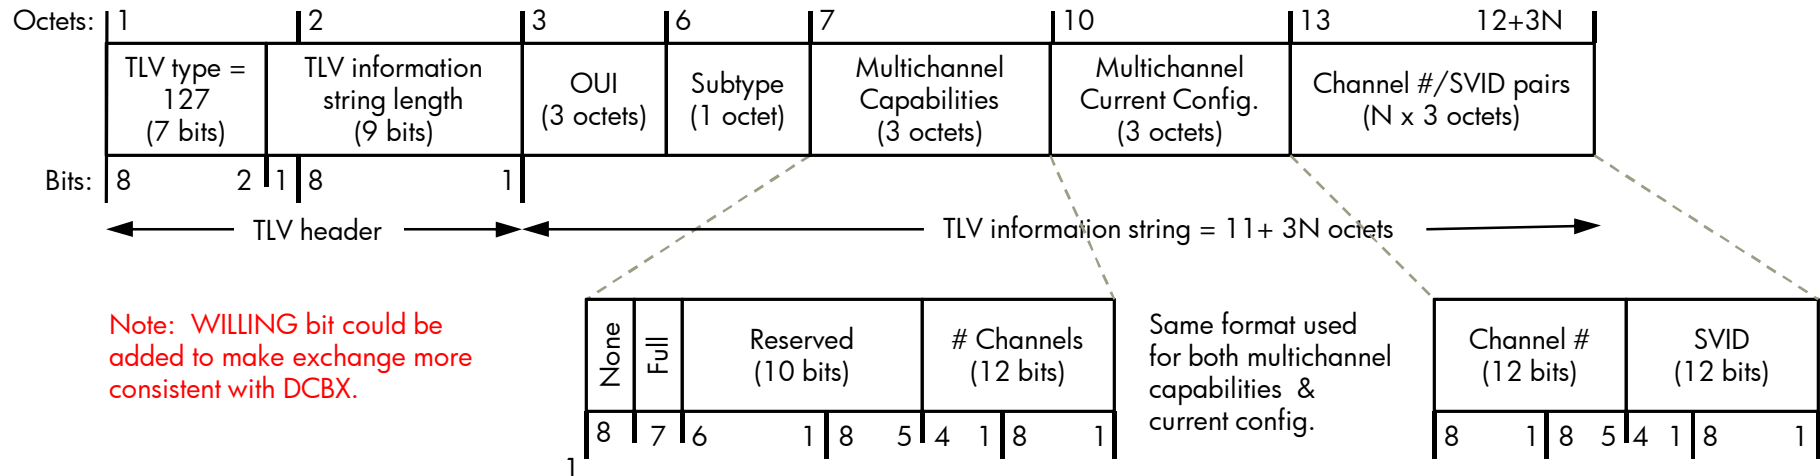
Transport: LLDP



- Multichannel Capabilities– Describes EVB multichannel capabilities that can be supported by the sender.
- Multichannel Current Configuration – Identifies the EVB multichannel capabilities that are currently enabled by the sender. Only one channel mode may be selected.
- Number Channels Supported – Identifies the number of SVID channels that are supported by the sender.
- Number Channels Currently Configured – Identifies the number of SVID channels that are configured/desired by the sender.
- Ch#/SVID Pairs
  - Channel # - indicates the index number of the channel. Allows insertion or deletion of specific channels while only listing the currently configured channels.
  - SVID – The S-Tag VLAN ID assigned to the channel. This is identified by the bridge. SVID of 0 means that no VLAN ID has been assigned.

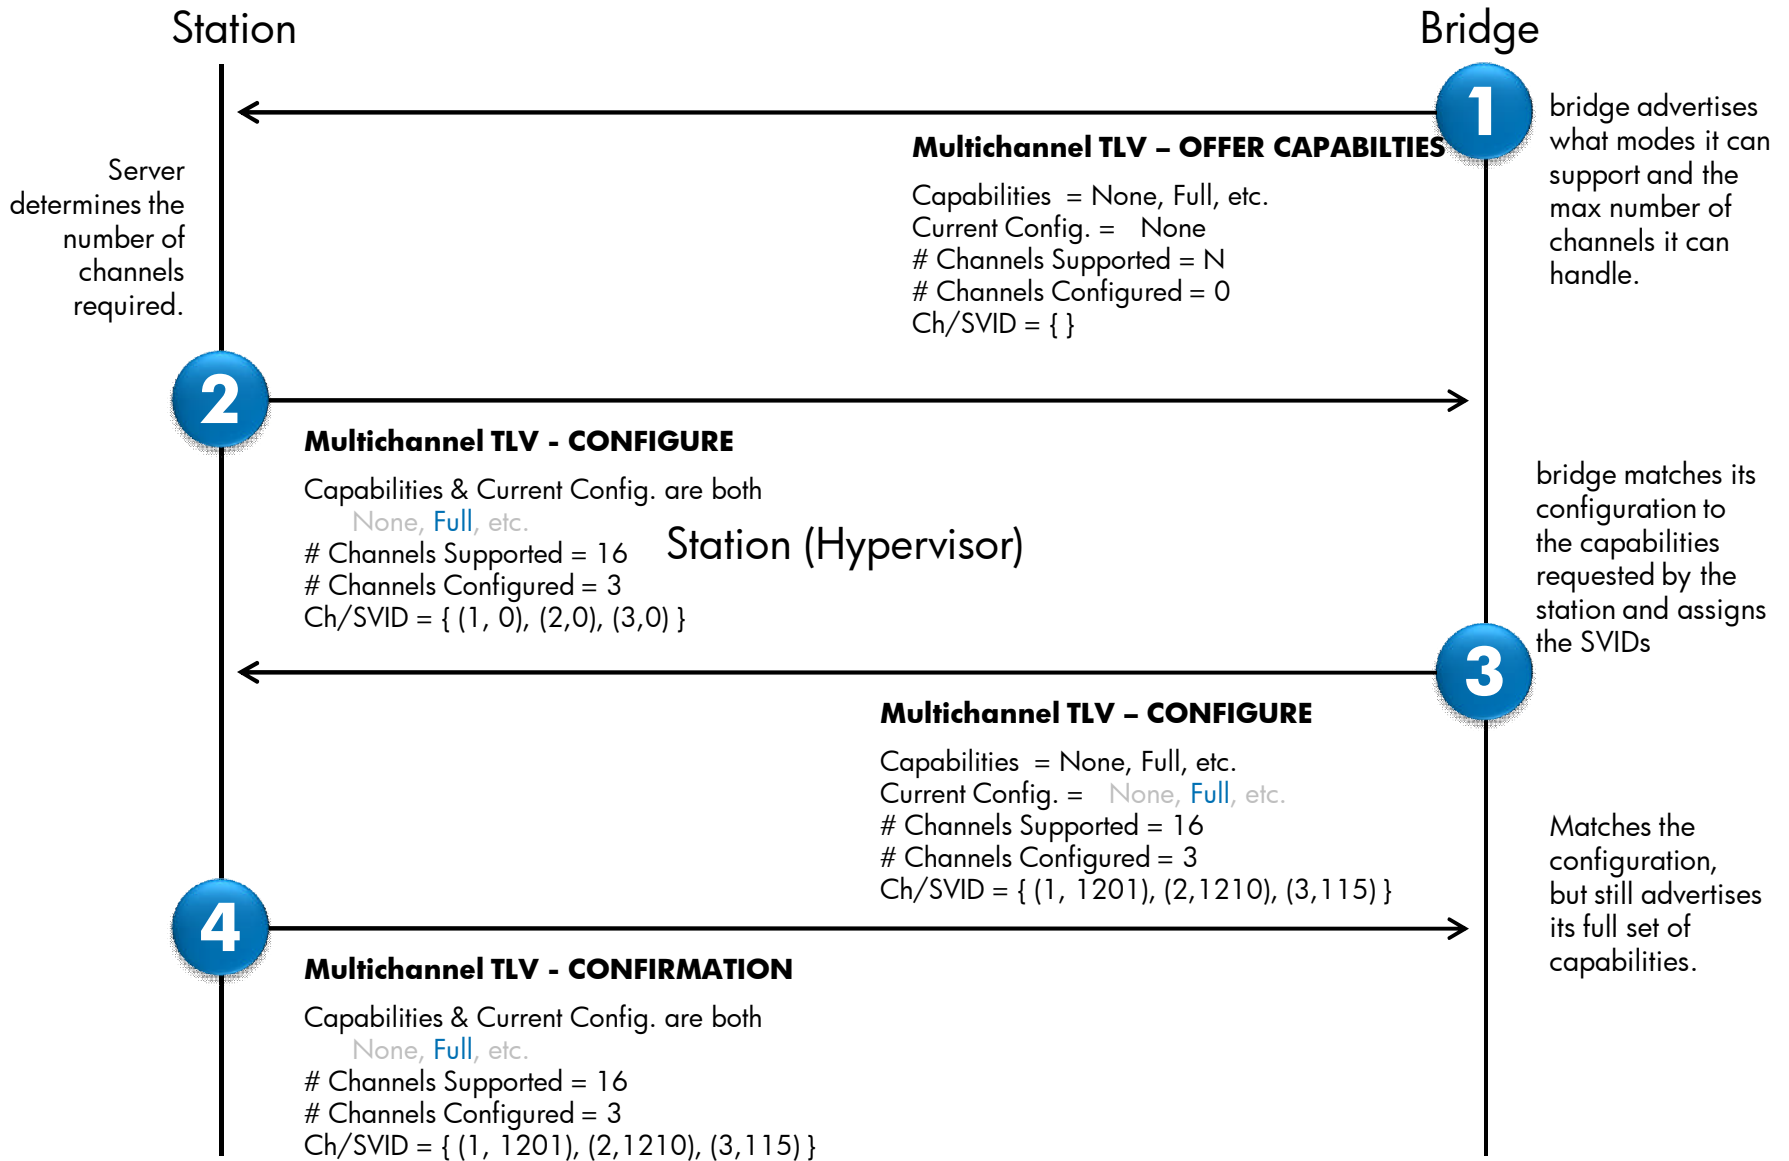
# Proposed Multichannel TLV (alternative format)

Transport: LLDP



- Multichannel Capabilities– Describes EVB multichannel capabilities that can be supported by the sender.
  - Number Channels Supported – Identifies the number of SVID channels that are supported by the sender.
- Multichannel Current Configuration – Identifies the EVB multichannel capabilities that are currently enabled by the sender. Only one channel mode may be selected.
  - Number Channels Configured – Identifies the number of SVID channels that are configured/desired by the sender.
- Channel #/SVID Pairs
  - Channel # - indicates the index number of the channel. Allows insertion or deletion of specific channels while only listing the currently configured channels.
  - SVID – The S-Tag VLAN ID assigned to the channel. This is identified by the bridge. SVID of 0 means that no VLAN ID has been assigned.

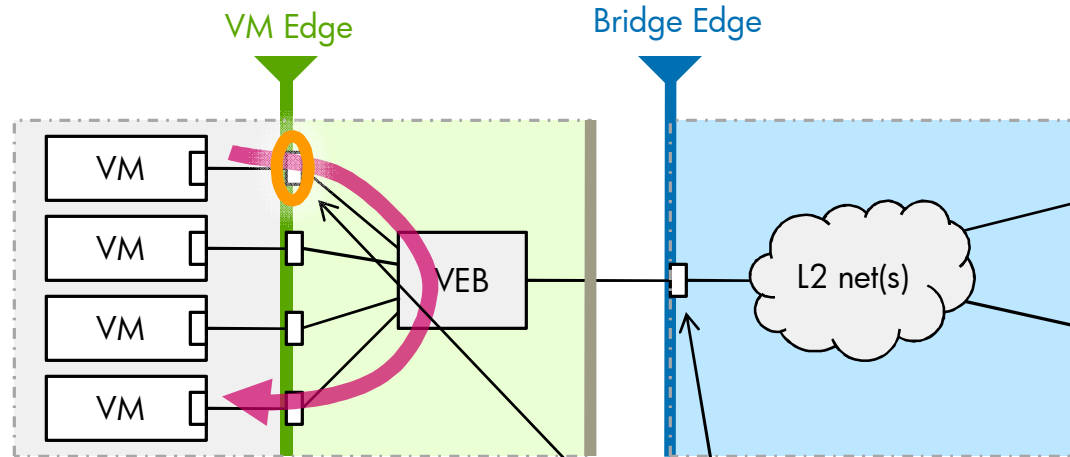
# Multichannel TLV Usage




# EVB Discovery

## VEPA and Reflective Relay

Regular vSwitch mode allows VM-to-VM traffic with limited policy enforcement.




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




Config. Connection settings communicated to both VM Edge and Bridge Edge.

# Proposed EVB TLV



Note: WILLING bit could be added to make exchange more consistent with DCBX.

Same format used for both multichannel capabilities & configuration.



Separate capabilities bit to indicate support for the TLV transport?

- EVB Capabilities - Describes EVB capabilities that can be supported by the sender. Includes port forwarding modes and other capabilities.
- EVB Current Configuration - Identifies the EVB capabilities that are currently enabled by the sender. (Only one port forwarding mode may be selected)
  - Forwarding Mode
    - Standard 802.1Q forwarding
    - Reflective Relay (as required for VEPA operation)
  - Capabilities Bits
    - VSI Discovery Protocol
    - 802.1X Authentication Required
- Number VSI's Supported - Identifies the number of VSI's that are supported by the sender. It may indicate limits imposed by buffering, ACL rules, etc.
- Number VSI's Configured - Identifies the number of VSI's currently in use.
  - From station, it indicates the number of resources that should be reserved on the bridge.
  - From the bridge, it indicates the number of active Virtual Station Interfaces (VSI) discovered/configured.



# EVB Discovery TLV Usage

This exchange use the Hypervisor's MAC Address as the Source MAC Address

Station (e.g., Hypervisor)

Bridge

Server configures itself from the available capabilities according to local policy.

2

## EVB TLV - CONFIGURE

Capabilities & Current Config.  
Forwarding: RR  
Other: VSI, Auth, etc.  
# VSIs Supported = J  
# VSIs Configured = K

1

## EVB TLV - OFFER CAPABILITIES

Capabilities  
Forwarding: Std, RR  
Other: VSI, Auth, etc.  
Current Config.(Std, None)  
# VSIs Supported = J  
# VSIs Configured = 0

Bridge advertises what modes it can support and the max number of VSI's it can handle.

Bridge matches its configuration to the limited capabilities advertised by the station.

3

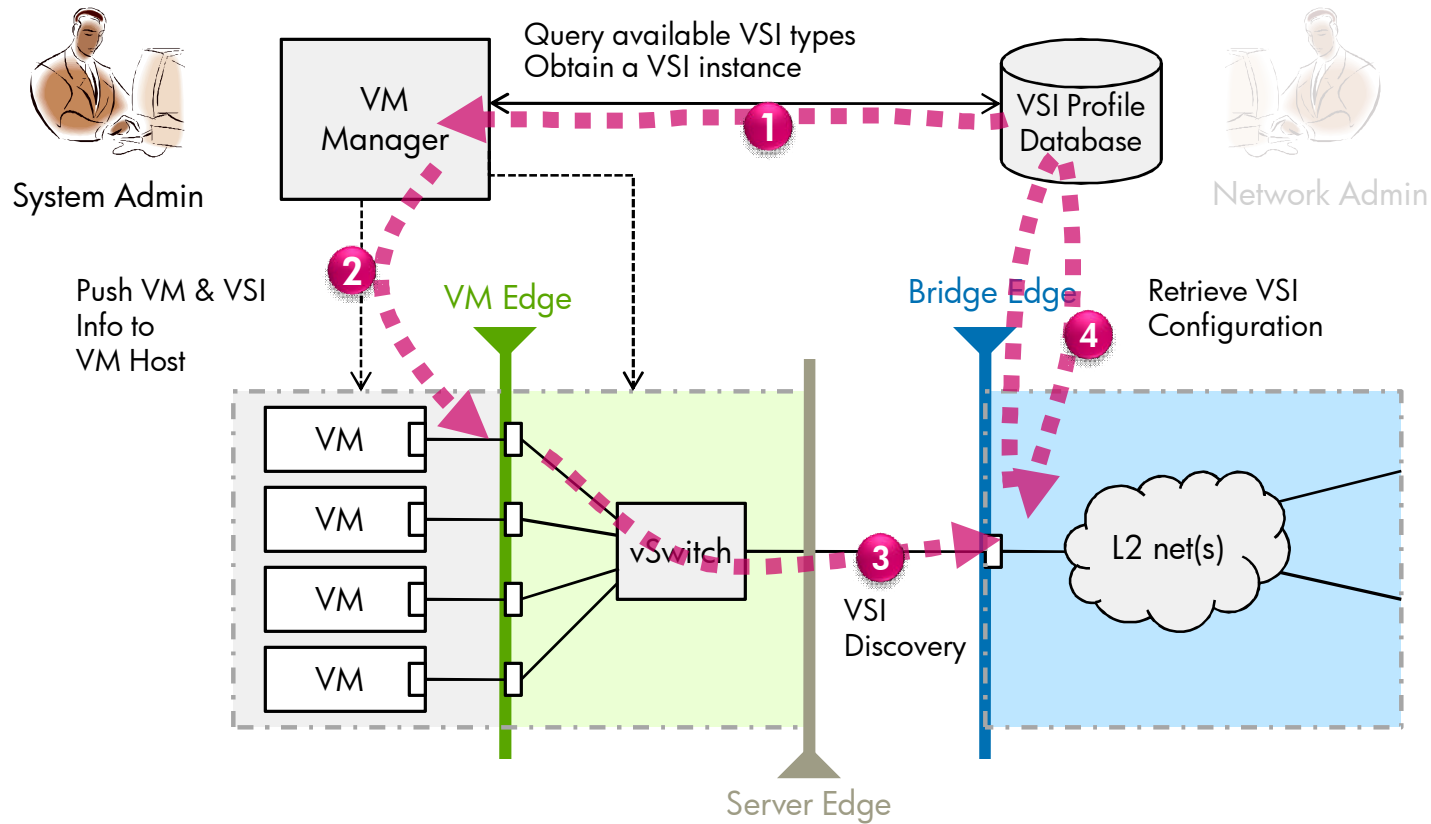
## EVB TLV - CONFIRMATION

Capabilities  
Forwarding: Std, RR  
Other: VSI, Auth, etc.  
Current Config.  
Forwarding: RR  
Other: VSI, Auth, etc.  
# VSIs Supported = J  
# VSIs Configured = K

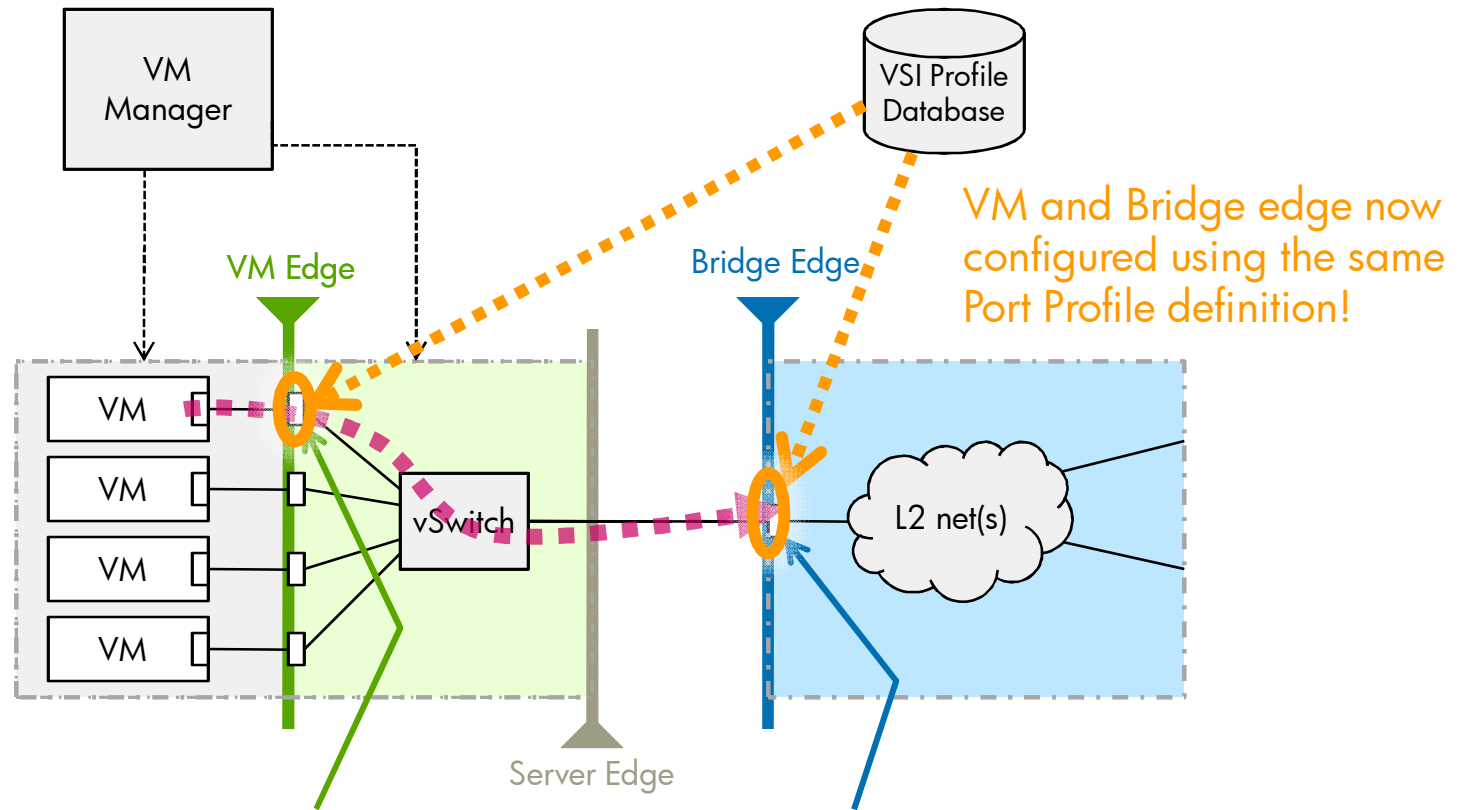
But still advertises its full set of capabilities.

# (Virtual) Station Interface Discovery and Configuration

# One Scenario for Configuring Edge Connections (VSIs)



# Result: Dynamic & Coordinated Configuration of Edge Connections



VM and Bridge edge now configured using the same Port Profile definition!

*VM Edge provisions and enforces some connection settings (VLAN, BW, etc.)*

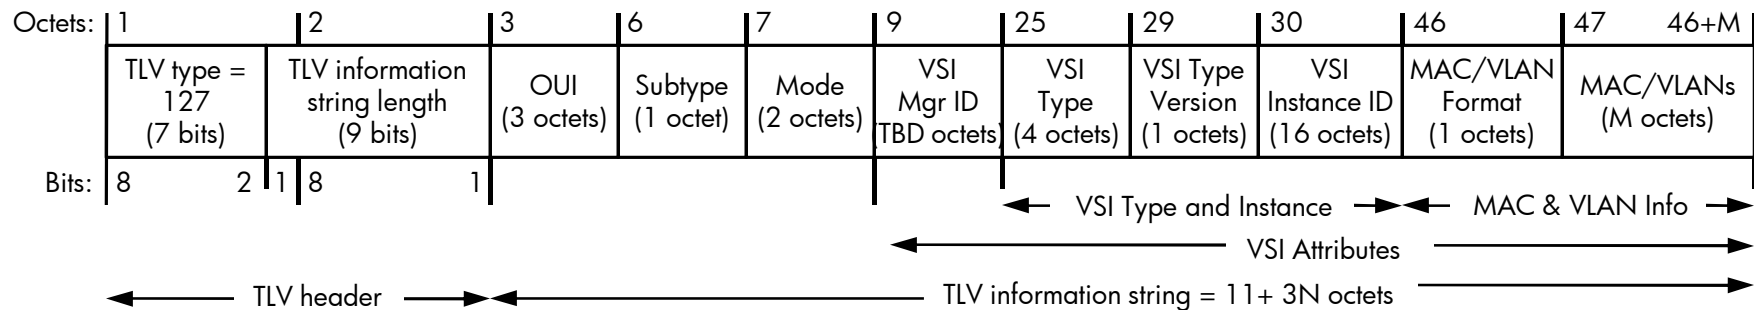
*Bridge edge provisions and polices coordinated set of server connection settings.*

# VSI Discovery TLV Supports Multiple Usage Models

- Database-driven, Instance Aware
  - Driven primarily by Port Profile Domain & Profile Instance
  - Allows database to track instance location, etc.
- Profile-driven
  - Driven primarily by Port Profile ID and Version (Instance ID ignored)
  - Allows for smaller, simplified port profile databases
  - Port Profile definitions may be fully cached on edge bridge
- SMB or Lab Config
  - Directly uses VLAN ID contained within vPort discovery LLDP frame to configure VLANs on the edge bridges

# Proposed VSI Discover/Configuration TLV

Transport: TBD



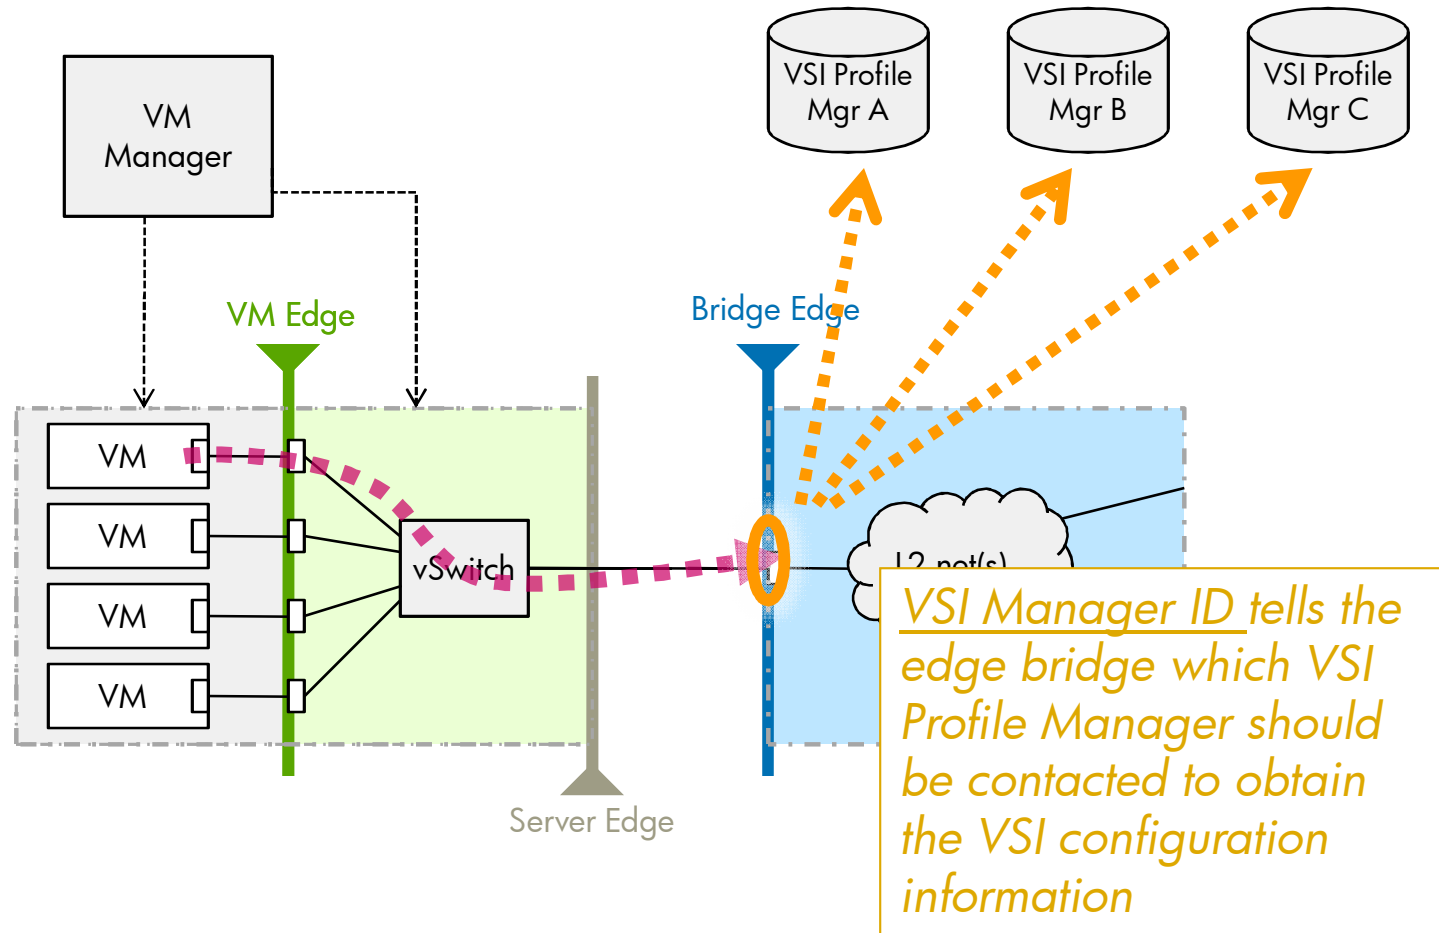
- Mode – Indicates VSI TLV Mode
  - First octet identifies a pre-associate, associate, de-associate, or the corresponding confirmation or rejection for each.
  - Second octet is used during a rejection to indicate the reason for the pre-assoc or assoc rejection.
- Port Manager ID – Identifies the Port Manager with the Database that holds the detailed port/VSI type and or instance definitions. May be the IP address of the management server.
- Port Type ID (PTID)– The integer identifier of the port/VSI profile type.
- Port Type ID Version – The integer identifier designating the expected/desired version of the PTID.
- VSI Instance ID – A globally unique ID for the connection instance. The ID shall be done consistent with IETF RFC 4122.
- Format – identifies the format of the MAC and VLAN information that follows in the TLV.
- MAC/VLANs – Listing of the MAC/VLANs associated with the Port Instance (VSI).  
If Format = 1, this would be a simple listing of MAC/VLAN pairs as shown below.

Default Port VLAN ID (2 octets)	# Entries (2 octets)	MAC (6 octets)	VLAN ID (2 octets)
---------------------------------	----------------------	----------------	--------------------

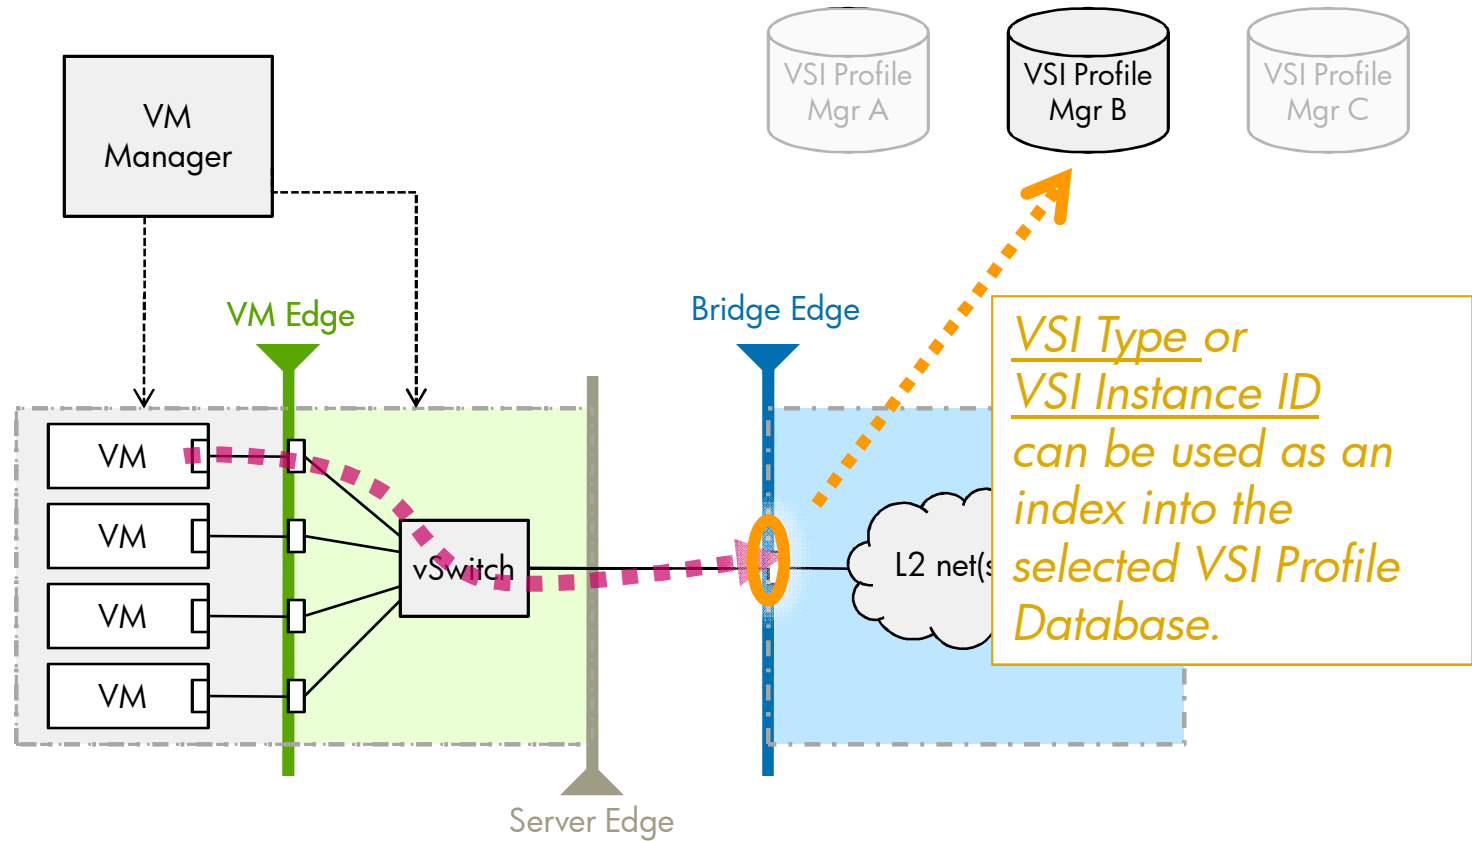
} x # Entries

NOTE: The station and bridge environments and their common understanding of the meaning of a port profile ID is outside the scope of this effort.

# VSI Manager ID



# VSI Type and Instance ID

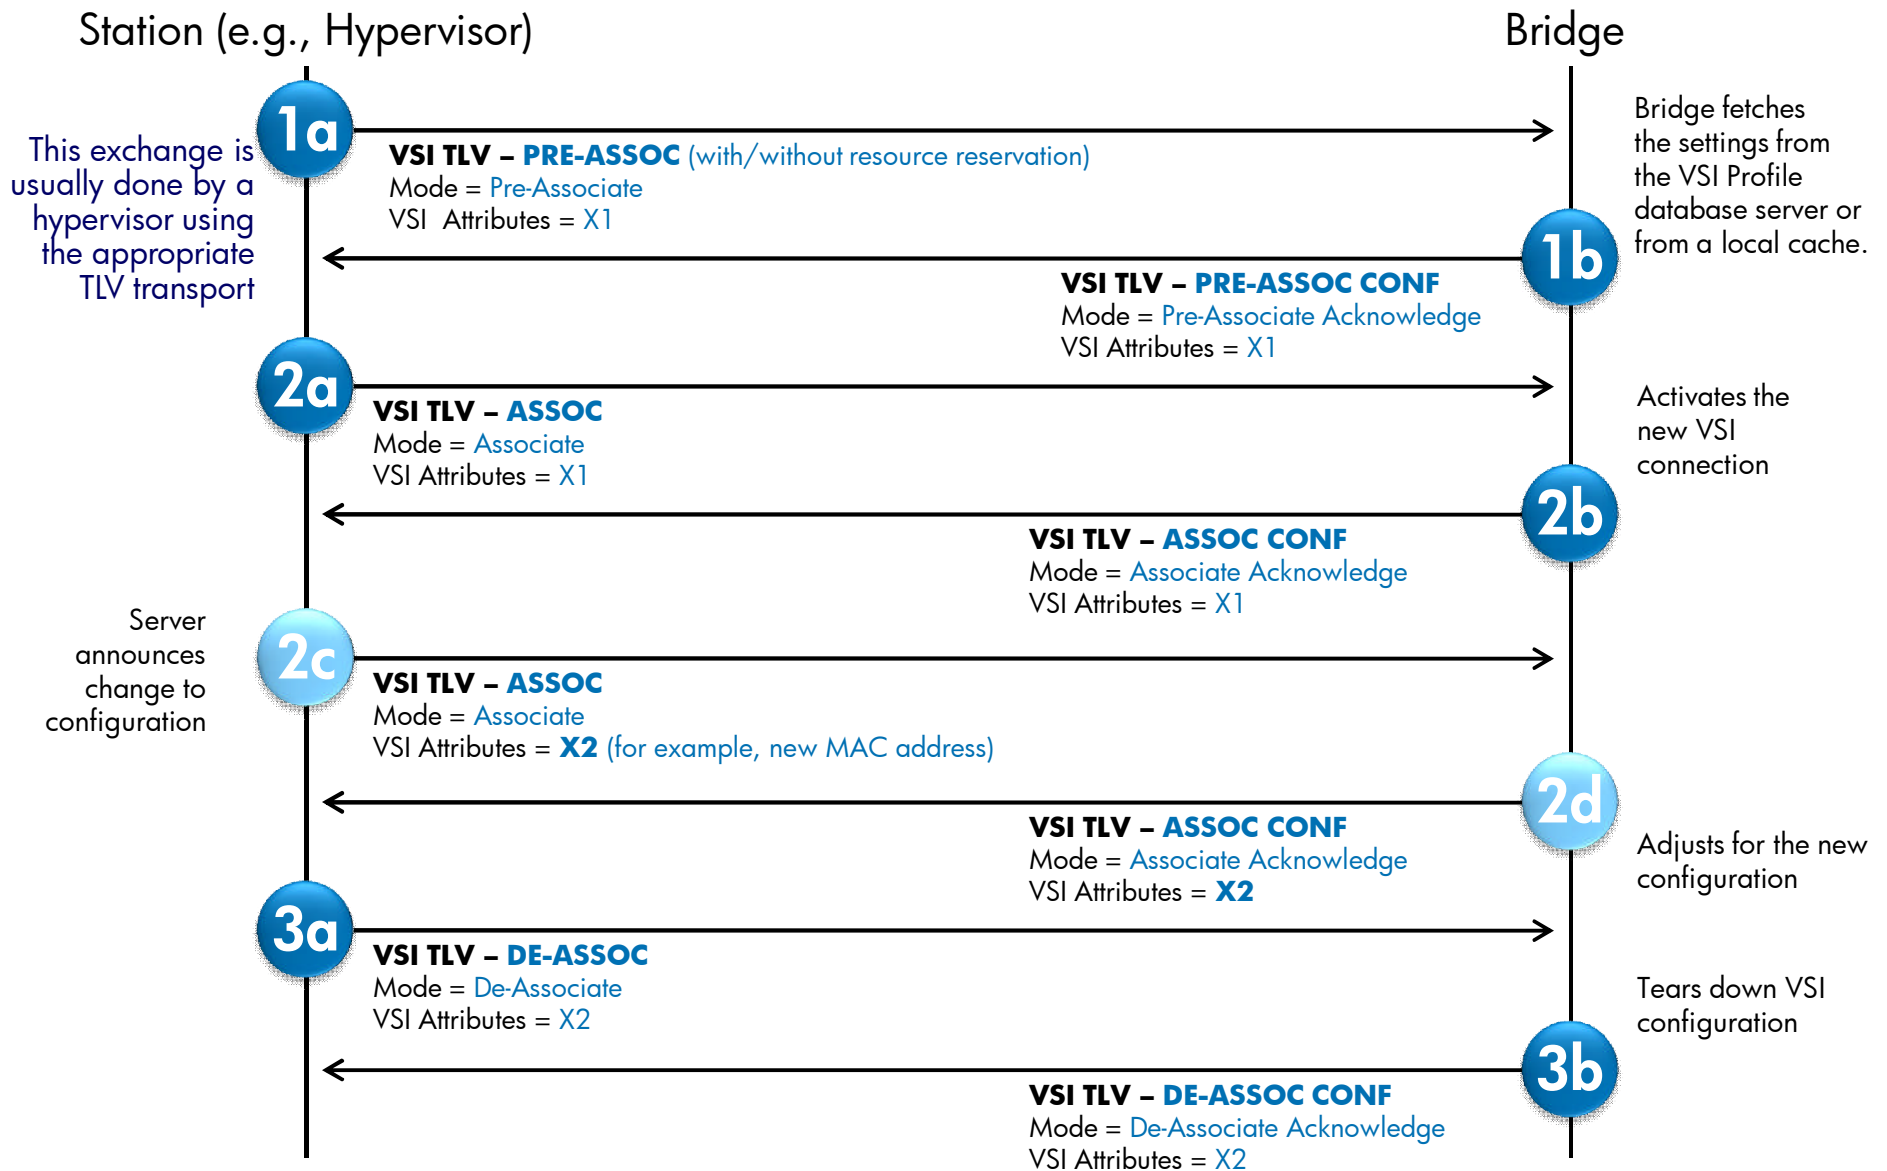




# Binding

- Specific MACs/VLANs are associated with a specific VSI instance and VSI type
  - MACs & VLANs may be represented in multiple formats
- Used by edge switch to associate a VSI profile (configuration parameters) to specific Ethernet traffic.

# VSI Discovery/Config. TLV Example



# Questions?