


VID on VDP Reply

Bob Sultan (bsultan@huawei.com)

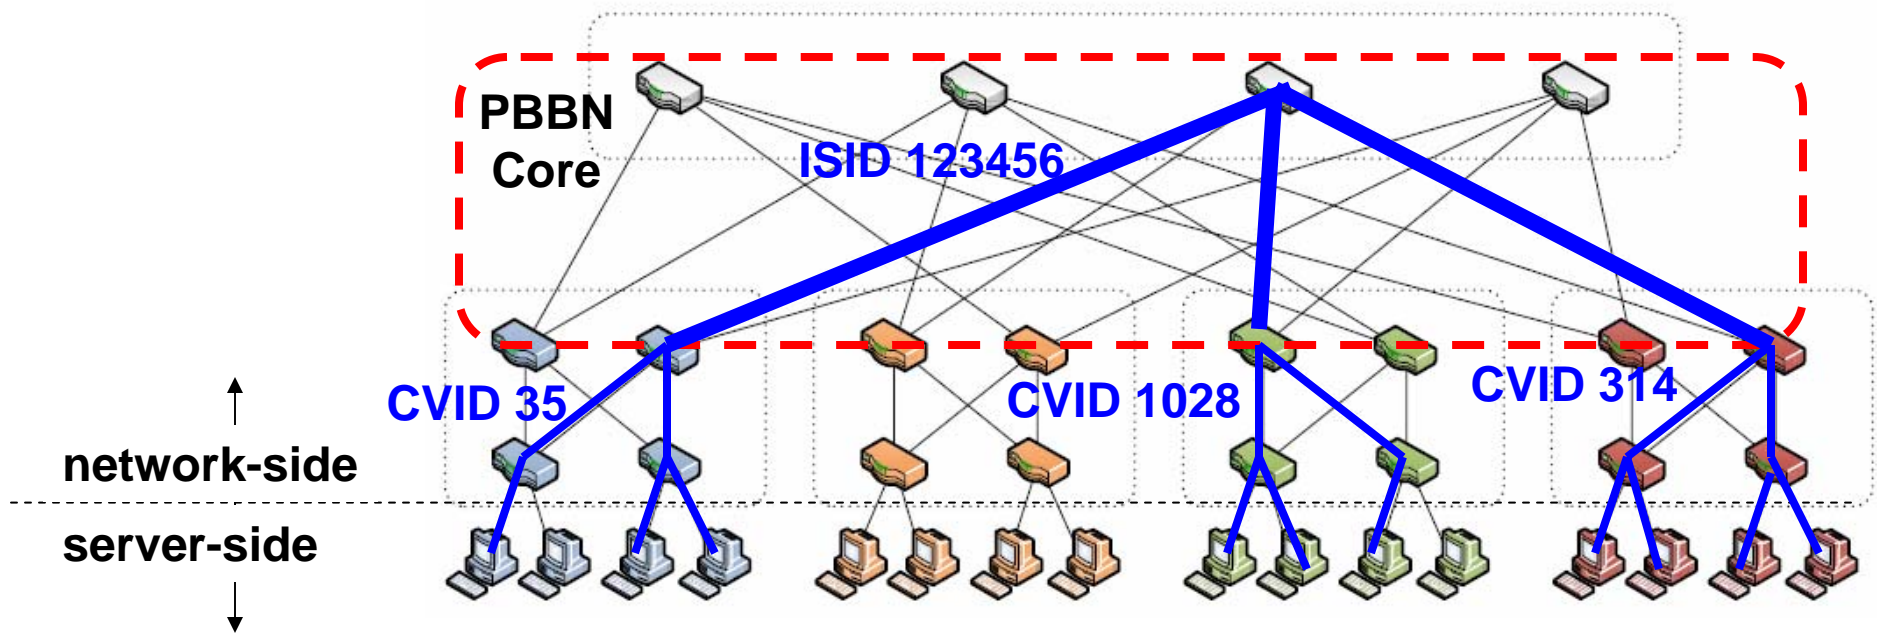
Li YiZhou (liyizhou@huawei.com)

Why is VID a *network-side* parameter?



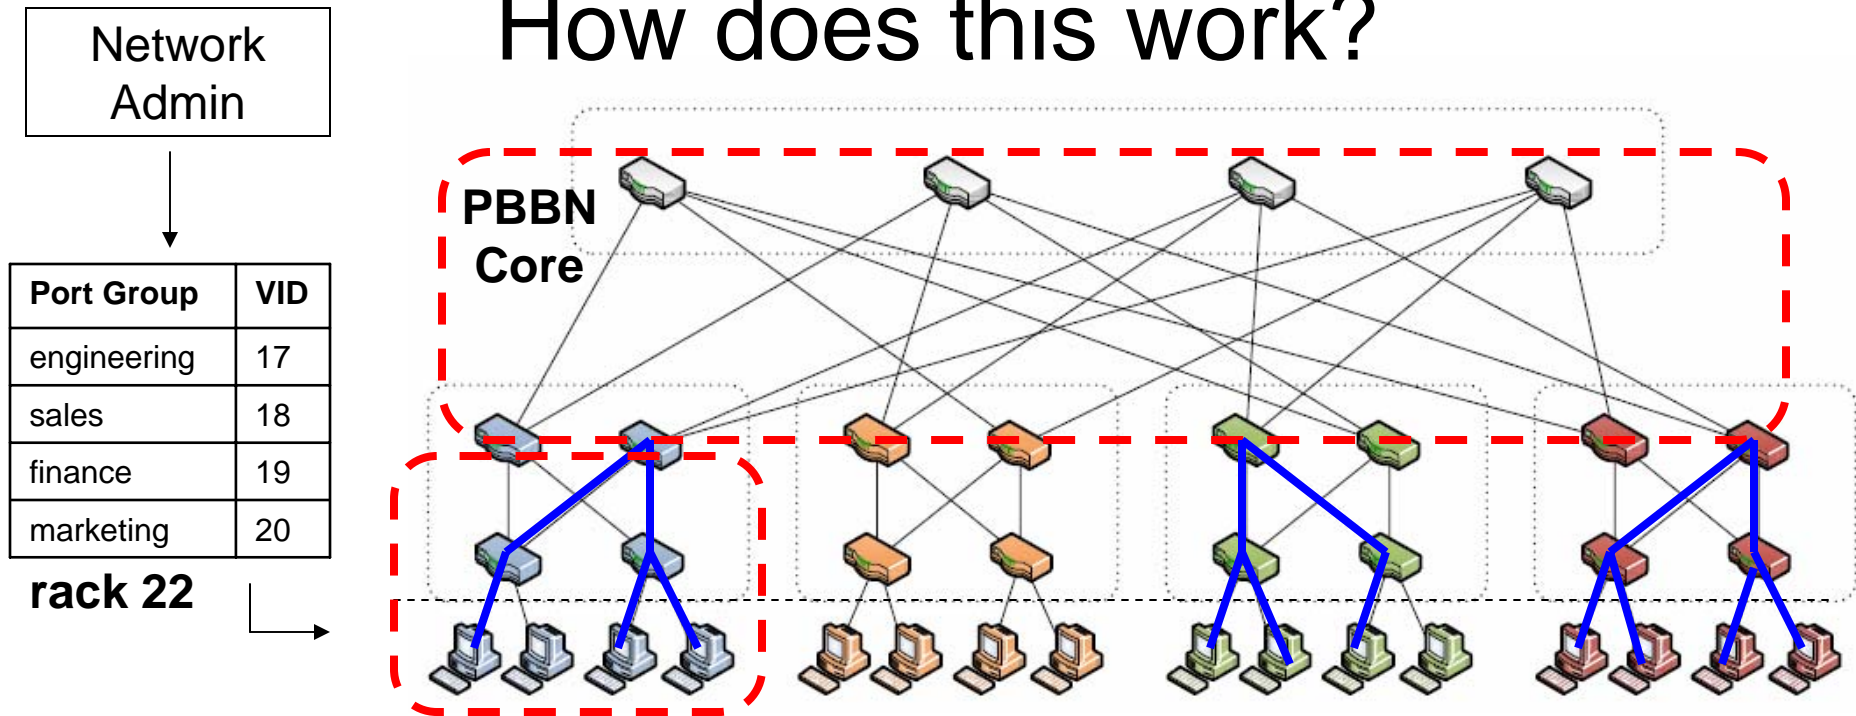
- While VLAN-awareness may extend to the server, the establishment of the VLAN (member-set for each bridge port) is a *network-side* activity;
- Server-side *admin* need not, and probably should not, know about VID;
 - *hypervisor* needs to know VID in order to apply VLAN tag;
 - VMware (for example) supports VID in VM configuration because there is currently no way for bridge-side to communicate this VID to server-side;
 - not what we want going forward;

Need network-side VID assignment



- Data Center providing cloud services likely to have >4K service instances;
- Any one rack supports up to 4K;
- Implies that ISID will be mapped to different VIDs at BEB to avoid conflict;
- Only network-side has knowledge to avoid these conflicts;

How does this work?

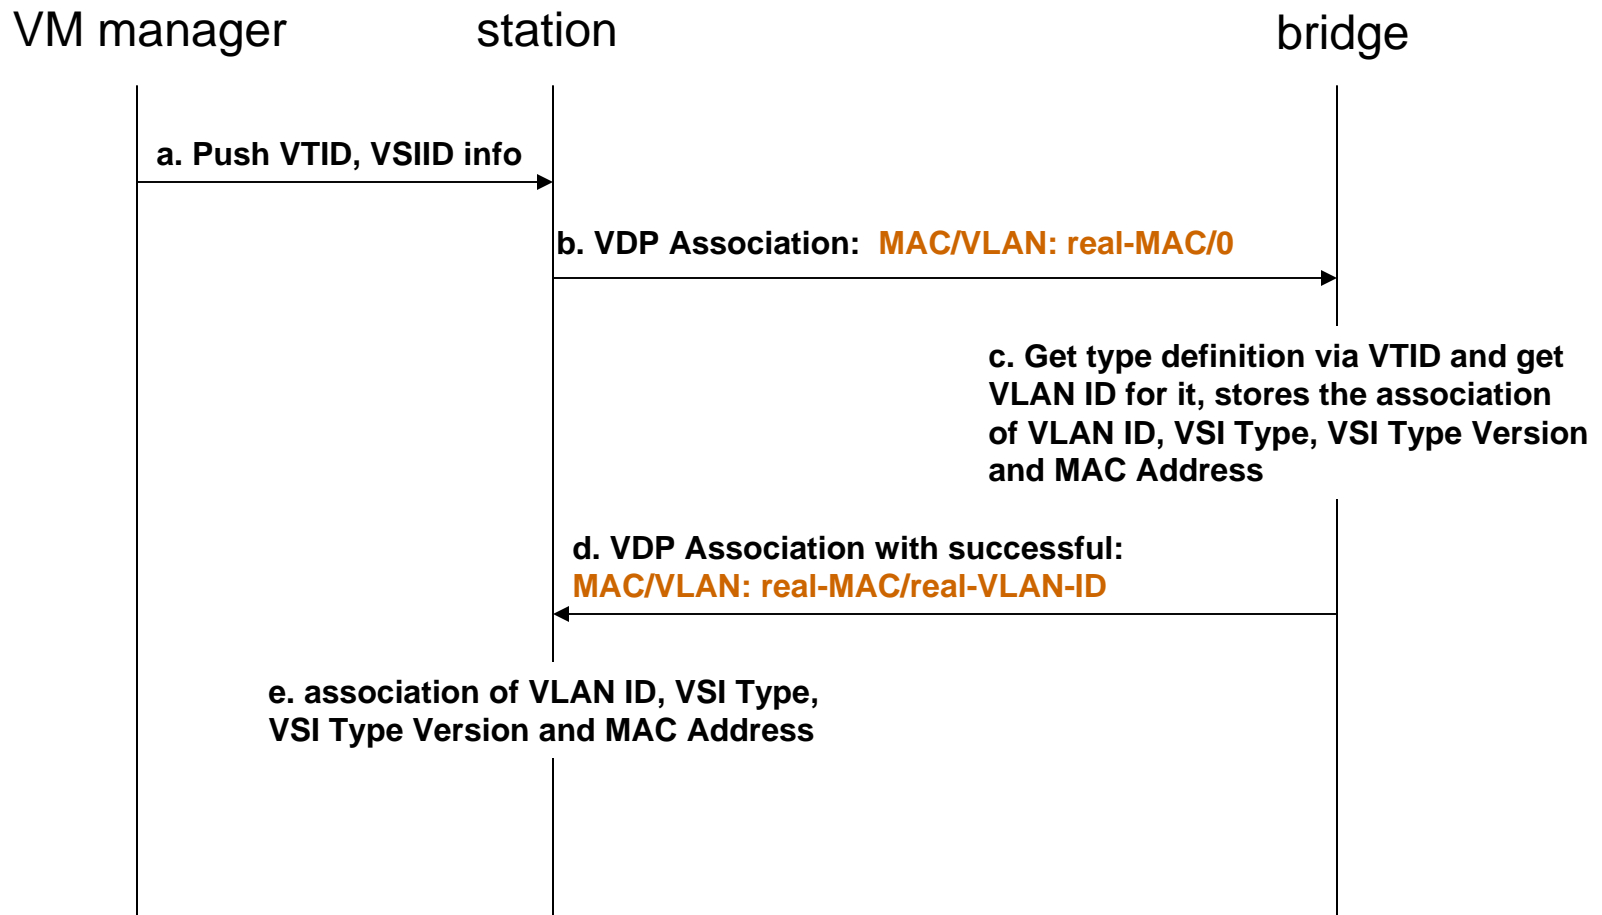


- VDP request identifies VSI-type;
- VSI-type contains Port Group(s);
- If TOR is BEB, for example, then all Bridges within rack use same table to map Port Group to VID;
 - table is different, in general from rack to rack;
 - table supplied by network admin;

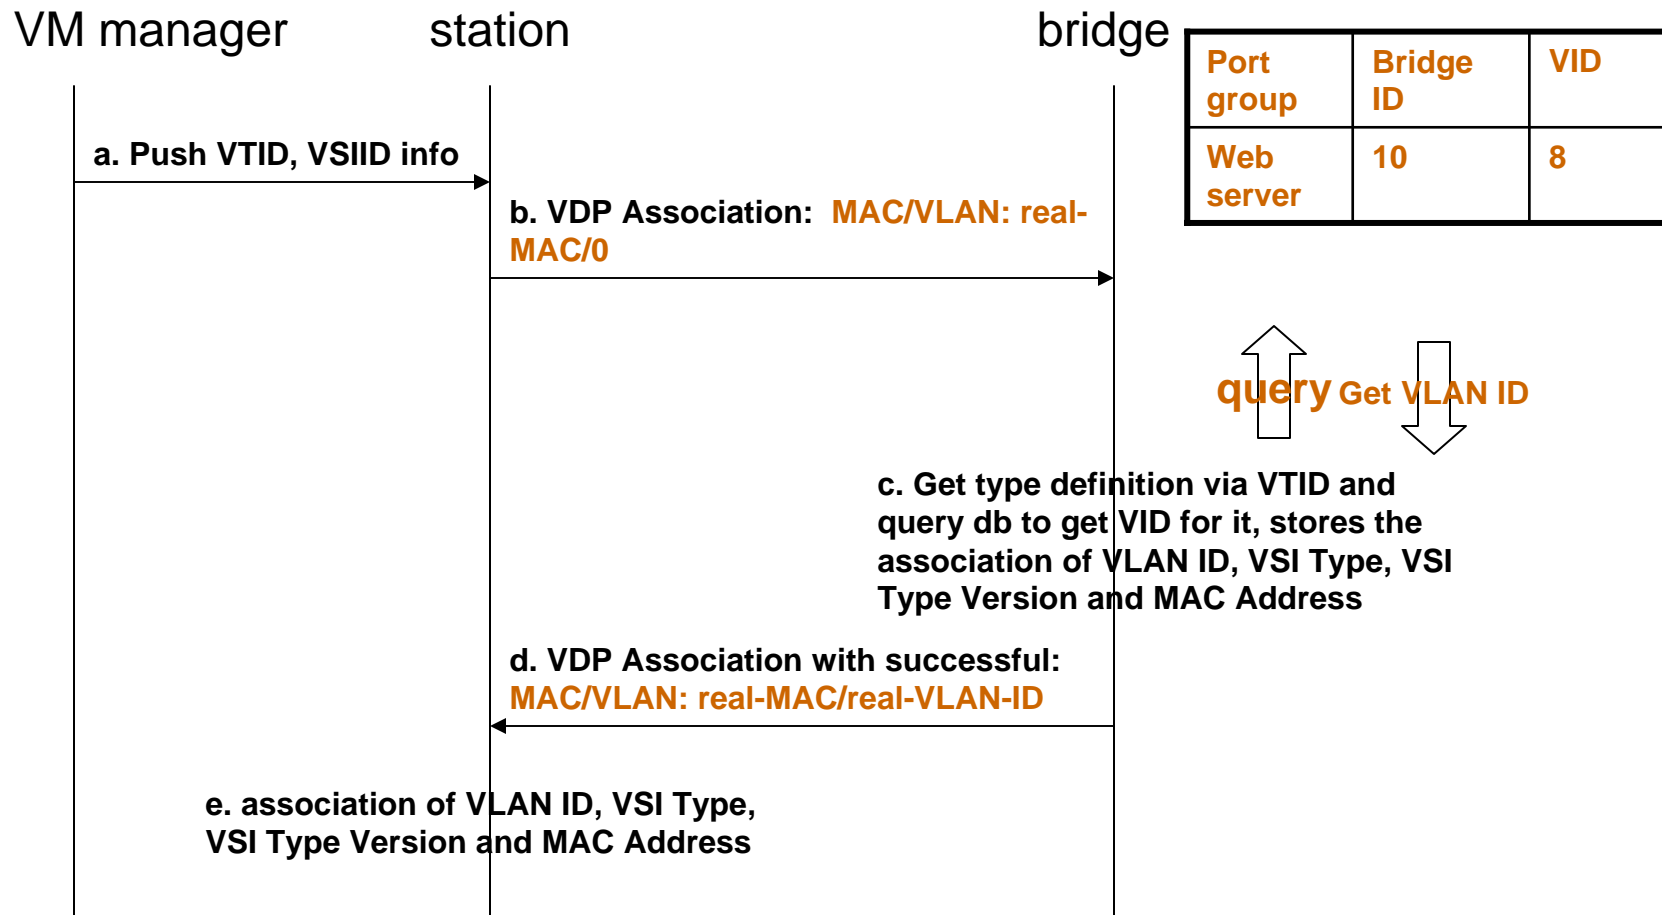
Recap VID assignment by bridge

- **Rationale:** VLAN ID is a network relevant parameter which may be under control of network admin. System admin may not care about the exact value though VSIs need to be provisioned with correct VID.
- System admin only needs to do the server planning based on logical partition or service, say port group named “web server” or “IT department”. It is an abstraction of VID which has certain logical meaning to system admin.
- Network admin knows the exact number for a particular VLAN, say port group “web server” uses VID 10. And he/she may not want to put VID in VSI type definition. Port group name/index or service name/index would be defined in VSI type definition instead.
 - Network Admin does not want to disclose the VID info to VM manager directly in order to demarcate the management domain clearly
 - Network Admin may want to map same port group to different VID under different bridge as VID number space is limited

Use case for getting VLAN ID from bridge (1)



Use case for getting VLAN ID from bridge (2)




Changes to current VDP

- Allow station to send VDP Associate without VID and to receive the exact VID from the bridge via VDP Associate Successful
 - VDP Associate request: VID = 0
 - VDP Associate Successful TLV to get real VID
- Correlation of VDP request and reply

Correlation of VDP request and reply

- Current approach
 - no explicit sequence number
 - Implicitly using **full TLV correlation**: correlation based on the whole VDP TLV fields except successful/failure octet
- Problems with bridge-assigned-VID: MAC/VLAN fields are not exactly same in VDP request and reply. Therefore full TLV correlation does not work.
- **Proposal: use partial TLV correlation.** i.e. use fields excluding MAC/VID as the key for correlation



Summary

- Must be able to communicate VID from Bridge to Server
 - necessary because service instance can be associated with a different VID in each rack;
- Partial correlation (ie., excluding list of MAC/VID from match) doesn't appear to pose a problem;
 - but open to looking at counterexamples;