# (Port/Transport) Mux
# - Application Scenario in Ethernet Service Provider Domain

**Linda Dunbar (ldunbar@huawei.com)**
**Ben MackCrane (tmackcrane@huawei.com)**
**Bob Sultan (bsultan@huawei.com)**

# Problem Statement #1:
## Mux Feature Desired by Ethernet Service Providers

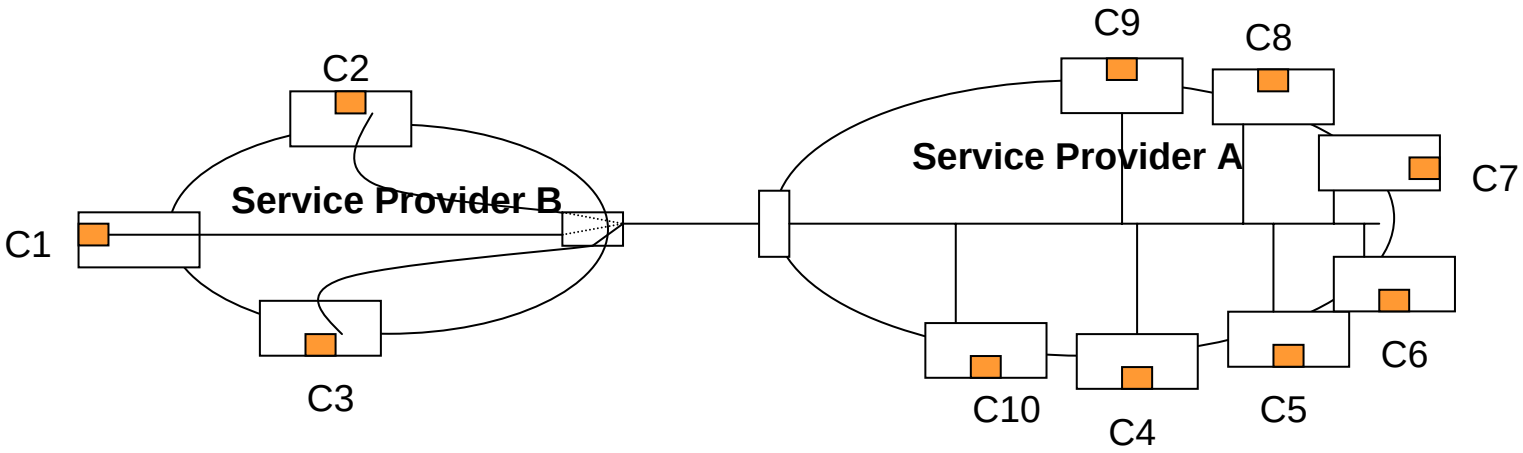


- **Application Scenario:**
  - Service Provider “A” leases three UNI ports from Service Provider B to offer a E-LAN service among client C (C1~C10 as shown above). There is an ENNI between two providers.
- **Desire from Ethernet Service Providers (from MEF)**
  - Service Provider A wants to have the full control over the traffic among their client ports, like enforcing certain types of policies among their clients.
  - Therefore, Service Provider A needs Service Provider B to backhaul all traffic from C1, C2 and C3 to “A”’s domain for layer 2 switch. i.e. Traffic from C1/C2/C3 are not switched within B domain

# Problem Statement #2:

## Data exchange among VMs are switched on external bridge

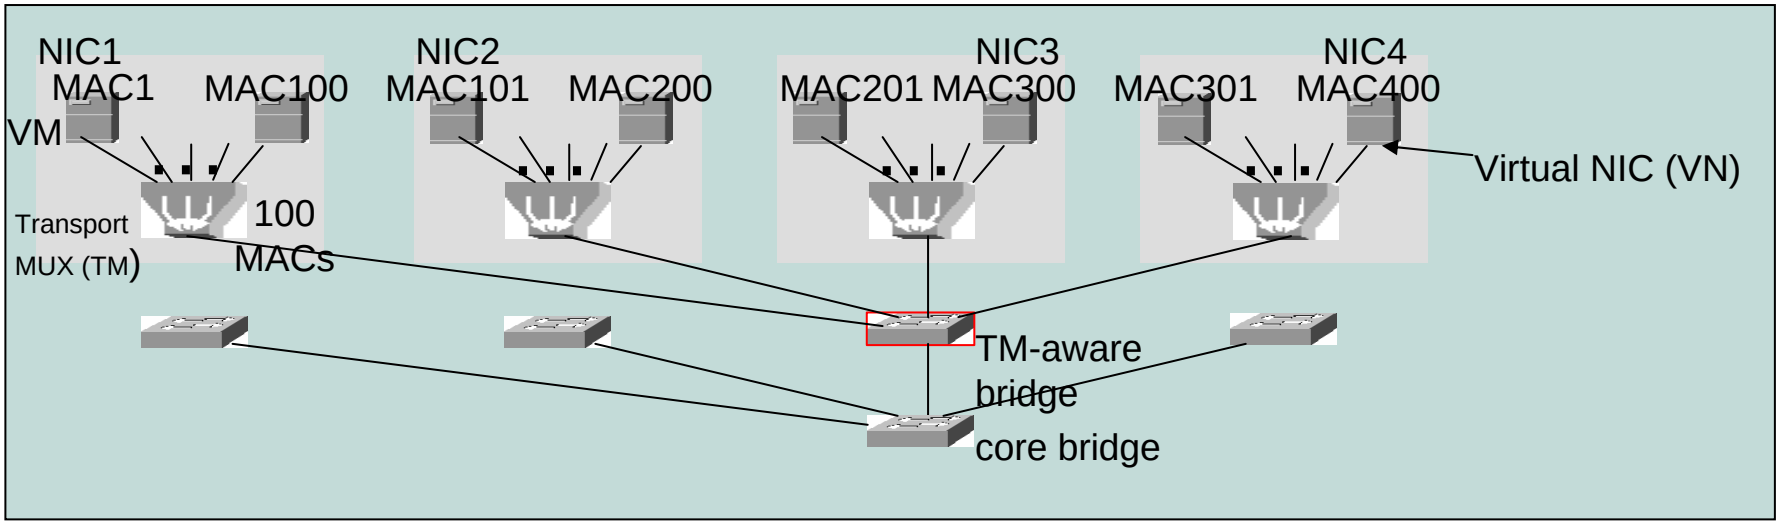


- **Application Scenario:**
  - VM-to-VM data exchange can't be performed on the server because the special advanced policy for data exchange among some VMs are not supported by the NIC or the Virtual (simple) Ethernet Switch on the server.

# Difference in requirements between the two problems

- **Problem #1:**
  - Provider A doesn't want Provider B in any learning or provisioning beyond the basic access (E-Line or E-Tree from provider B is good enough).
  - Replicating frames in provider A over ENNI towards B is acceptable.
- **Problem #2:**
  - Replicating frames by external bridge towards NIC is not acceptable.
  - Provisioning on NIC is acceptable

# Potential solution to solve problem statement #1 with minimum change to IEEE802.1Q

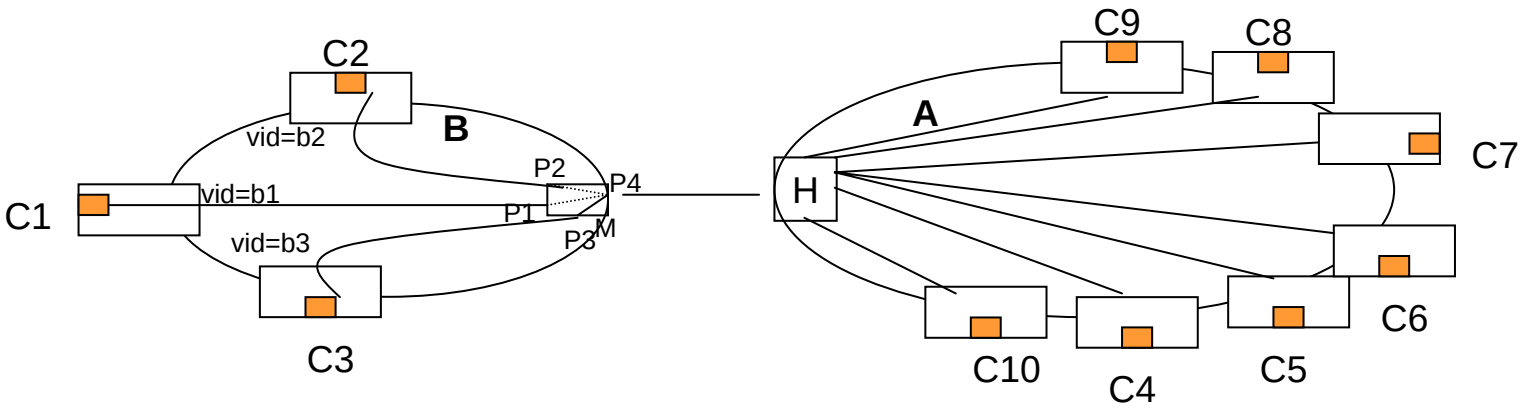


- **B needs those basic configuration (without any change to 802.1Q):**
  - Traffic from C1, C2 and C3 all have distinct vid, so that they are not switched anywhere within Domain B. E.g. Traffic from C1 has vid = b1, traffic from C2 has vid=b2, and traffic from C3 has vid=b3
  - At the location where the C1/C2/C3 traffic are aggregated together, say M, P1 only allows vid b1, P2 only allows vid b2, P3 only allows b3. P4 allows b1/b2/b3
- **One port on Domain A (the first port facing ENNI) has to support**
  - Multiple virtual ports with each vid mapping to one virtual port. (New to IEEE802.1Q)
  - Each virtual port has a S-component to change the vid of the incoming frame to a common vid (say a1) in Provider Domain A (No change to IEEE802.1Q)
  - When the frame exits the virtual port, the S-component will change the common vid (a1) back to the vid for the virtual port. E.g:
    - When a frame from C1 needs to reach nodes in C2, the frame will have <DA, vid=b1>. When the frame traverses the S-component on the virtual port for b1 on Node H, the vid will be changed to a1.
    - When the FDB determines the egress port is the virtual port for b2, S-component on the b2 virtual port will change the vid from a1 to b2.

# Similar Solution to Problem Statement #2

- **Each VNIC port adds an unique vlan_identifier (without any change to 802.1Q)**
  - so that the traffic from NIC carries a vlan_identifier which the external bridge can use to distinguish which VNIC the traffic is from.
- **The external bridge (the first port facing the server) has to support (same changes to 802.1Q as solution to problem #1)**
  - Multiple virtual ports are created on the first external port facing the server (“Virtual Port capable port” which is not necessarily supported by every external bridge port)
  - Same behavior as Problem #1

# Ways to avoid multiple (broadcast unknown) frames sent back over the ENNI (or the external link to Server)

- Have all the VMs register themselves at the external port via 802.1ak to avoid broadcast unknown with multiple frames (one frame for each virtual port) looped back over the external link.
- The number of broadcast unknown frames looped back over ENNI is limited by the number of UNI ports which Provider A leases from Provider B. Usually the number is not large. So it is not a big deal.

# Ways to avoid multiple multicast frames sent back over the ENNI (or the external link to Server)

- For the proposed solution to Problem #1, one multicast frame originated from C1 will cause two multicast frames sent back by the Virtual Port on H (with vid=b2 and vid=b3 respectively).
- A simple way to avoid multiple multicast frames sent by H over ENNI:
  - Problem #1: Provider A basically needs a Rooted Multipoint service from B. Multiple vids from B to A, and single vid from A back to B.
  - Problem #2: multiple vids towards external bridge port and single vid from the external bridge.

# Ways to avoid sending Multicast frames (sent back from Virtual Port) to their senders

- **The sender doesn’t register the group_address which he doesn’t want to receive.**
  - This method works when one single vid coming from external bridge (or from A back to B).
- **The leaf ports (or the UNI Port) perform SA checking to filter out frames back to sender.**

# Proposal:

## Create a PAR in Internetworking group to provide a solution to force traffic to be switched at a different domain.

- **Potentially there is common approach which can be used for both problems.**
- **There are two components to the solution:**
  - Component 1: Mux function which doesn't require any change to IEEE802.1Q
  - Component 2: The concept of multiple Virtual Ports on the external bridge port.