Multicast pruning in .1ah Provider Backbone Bridged Network

May 2005

Muneyoshi Suzuki
Paul Bottorff
Michael Chen

Background

- Wide Area Ethernet was replacement of point-to-point leased line, ATM, and FR in early stage of service deployment
  - Thus, customers used Wide Area Ethernet as point-to-point links
- Since Wide Area Ethernet services can support point-to-multipoint communications, customers adopt multicast applications now
  - OSPF broadcast mode
  - Video distribution (Replacement of satellite communications)
  - Some VoIP control signals
- Multicast capability is a distinct feature of Wide Area Ethernet, however it is burden for networks
  - Pruning is indispensable to maintain communications quality such as frame delay and jitter
  - Especially, subscriber lines may not be broadband, so QoS is strongly affected by flooded traffic

An example of unnecessary flooding

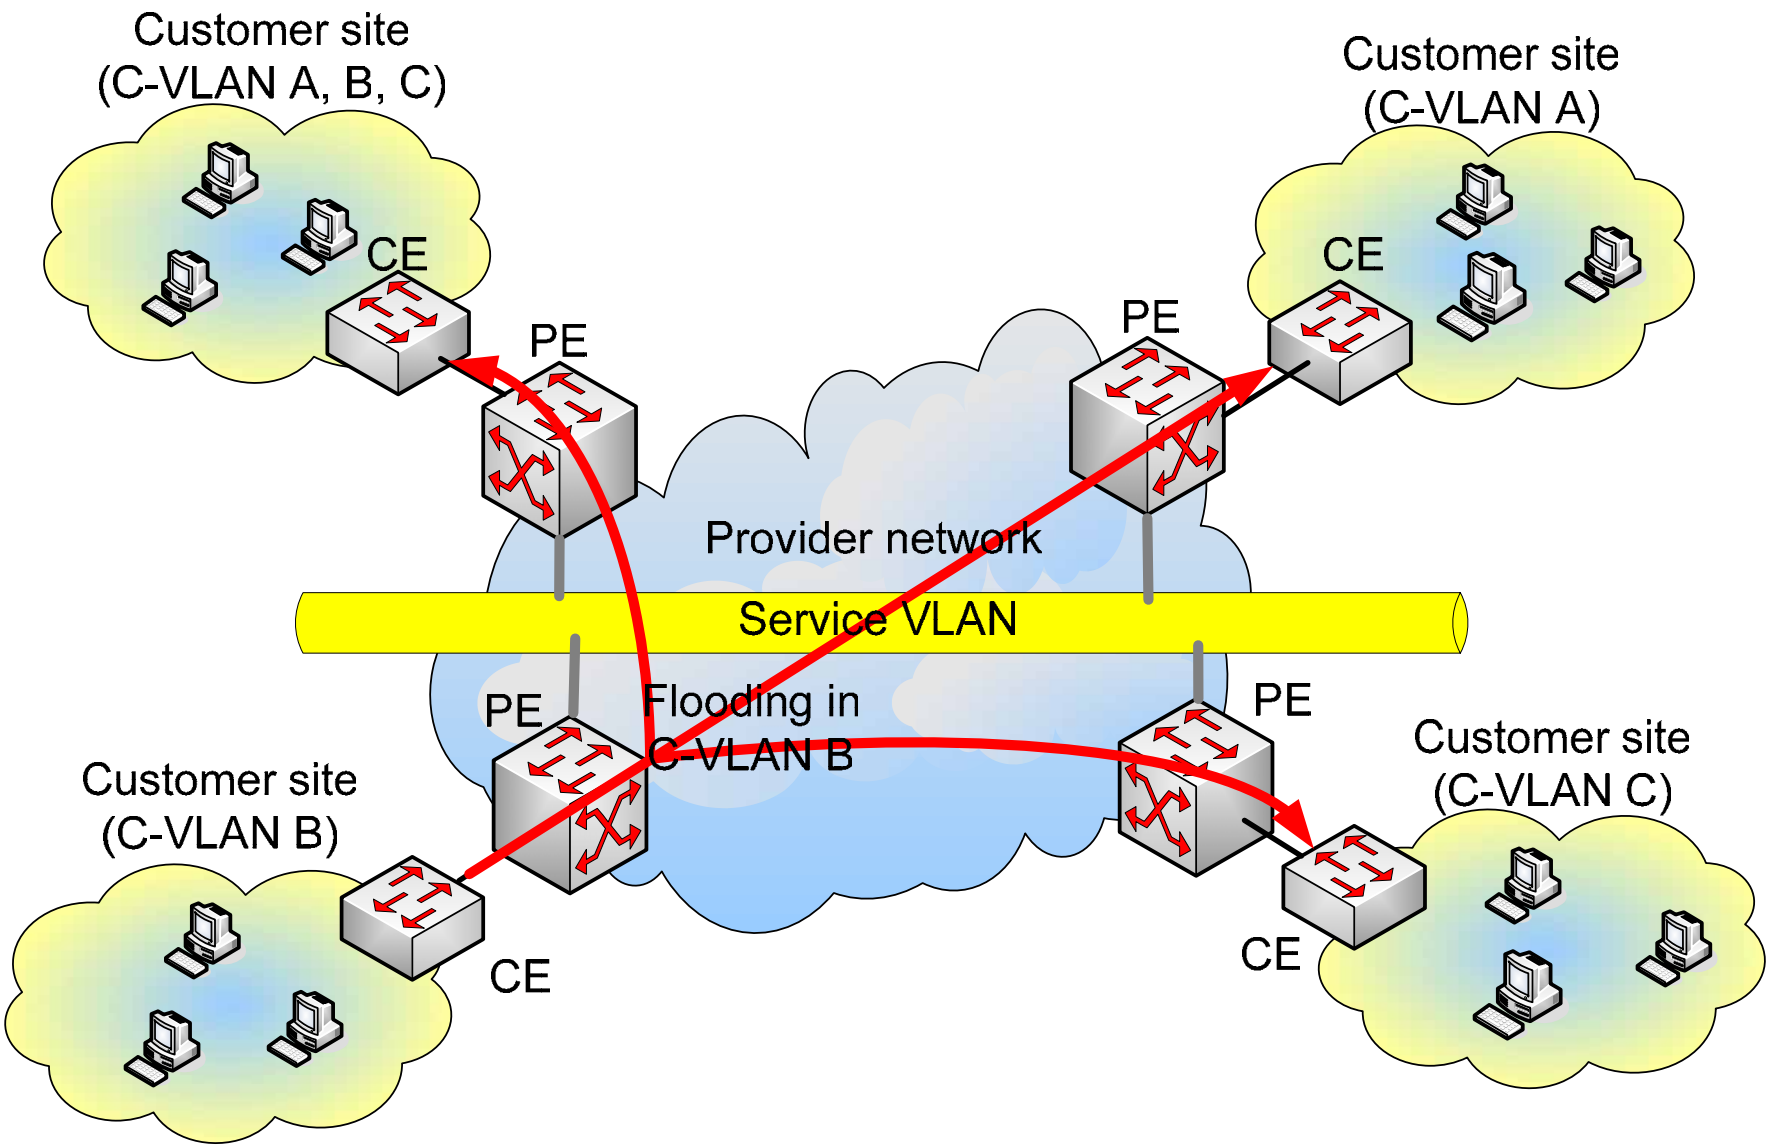

Requirements for Pruning

- In terms of customers, a "service instance" in provider network is logically equivalent with a "long yellow cable"
- So, unnecessary flooding is propagated to all customer sites as well as subscriber lines
- Thus, pruning in provider network/provider backbone network should be aware of C-VLAN
- Therefore:
  - Pruning in .1ad provider network should be based on DA and S-VID as well as C-VID
  - Pruning in .1ah provider backbone network should be based on B-DA and B-VID as well as C-VID
- However, pruning in current stacked VLAN implementations is unaware of C-VID

Pruning in .1ad provider network

- Option 1
  - Develop S-VLAN Bridge that supports DA, S-VID, and C-VID based pruning
  - Perfect, but too unrealistic scenario
- Option 2
  - Use C-tagged service interface to map a C-VLAN transferring multicast traffic to an S-VLAN
  - Prune traffic based on DA and S-VID in frame
  - Practical solution for current Bridge implementations
  - However, it wastes S-VID space

Pruning in .1ah provider backbone network

- Option 1
  - Develop core Bridge that supports (B-DA, B-VID, and C-VID) or (B-DA, I-SID, and C-VID) based pruning
  - C-VID snooping from .1ad frame have to be expensive
- Option 2: Ingress Bridge maps a C-VLAN to a B-VLAN
  - Prune traffic based on B-DA and B-VID in frame
  - B-VID space may be too small to enable this scenario
- Option 3: Ingress Bridge maps a C-VLAN to an extended S-VLAN
  - Prune traffic based on B-DA and I-SID in frame
  - Approximately 2^{28} I-SID space is required for 1M instances
  - Current .1ad Bridges don't support I-SID based pruning
- Option 4: Ingress Bridge maps a C-VLAN transferring multicast traffic to a multicast B-DA
  - Prune traffic based on B-DA and B-VID in frame
  - {S-VID, C-VID} -> multicast B-DA table is required
  - Multicast address space is 2^{47}, so large number of pruned C-VLANs in backbone could be supported

A new table for Option 4

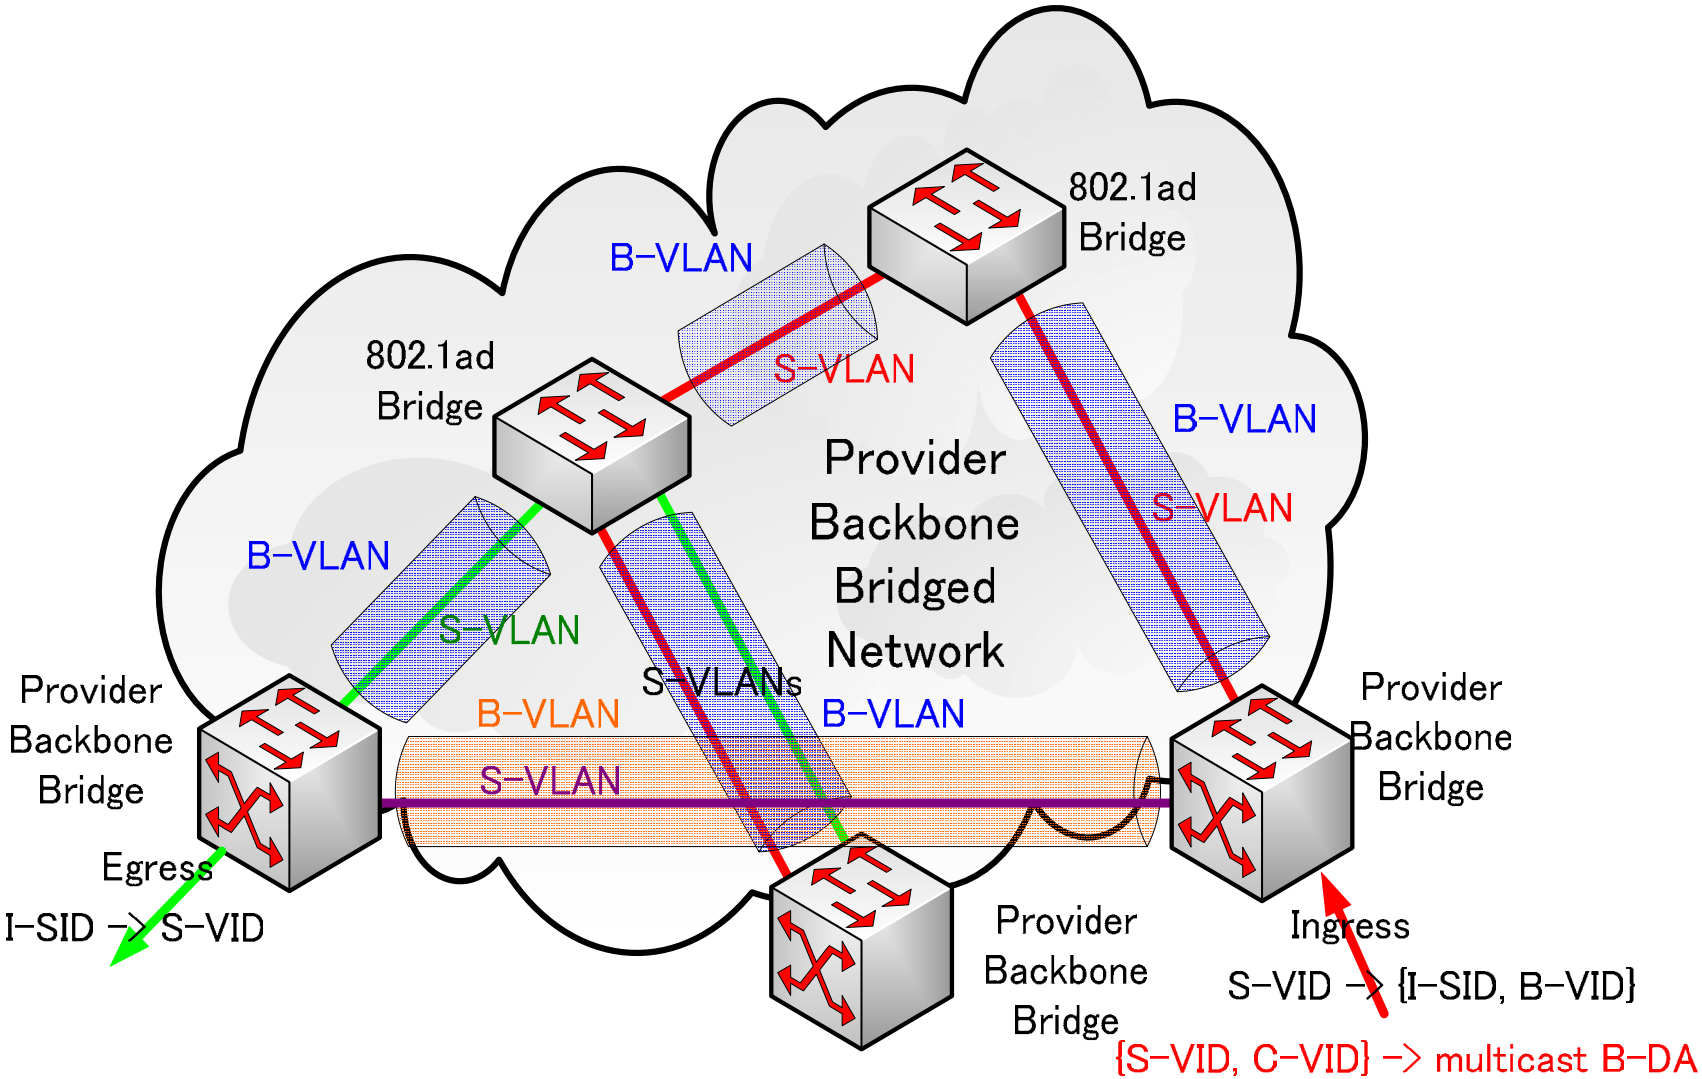

Discussions

- Option 4 may be feasible for C-VID based pruning in .1ah backbone
  - Options 2 and 3 may be alternatives
  - These options could support large number of C-VLANs
  - I-SID space should be large enough to work with all options
- To enable Option 4, a pre-configured {S-VID, C-VID} -> multicast B-DA translation table for multicast traffic to be pruned is required in Provider Backbone Bridges
- In this case, B-DA assignment scheme for ingress frame is:
  - If DA in .1ad frame is a learned unicast address, B-DA is assigned from FDB
  - Otherwise, search the translation table with {S-VID, C-VID} in .1ad frame,
    - If it succeed, B-DA is assigned from the entry
    - Otherwise
      - B-DA is a pre-configured default multicast address, or
      - If DA is a multicast address, B-DA is the same as DA
      - If DA is an unlearned unicast or broadcast address, B-DA is the broadcast address