


SPBM in the Data Center
May 2010 802.1 Interim Meeting
Geneva, Switzerland

Bob Sultan (bsultan@huawei.com)

Some Supporters of this Proposal


- Paul Bottorff
- Nigel Bragg
- Paul Congdon
- Don Fedyk
- Ben Mack-Crane
- Bob Sultan
- Paul Unbehagen

Fat-tree Architecture



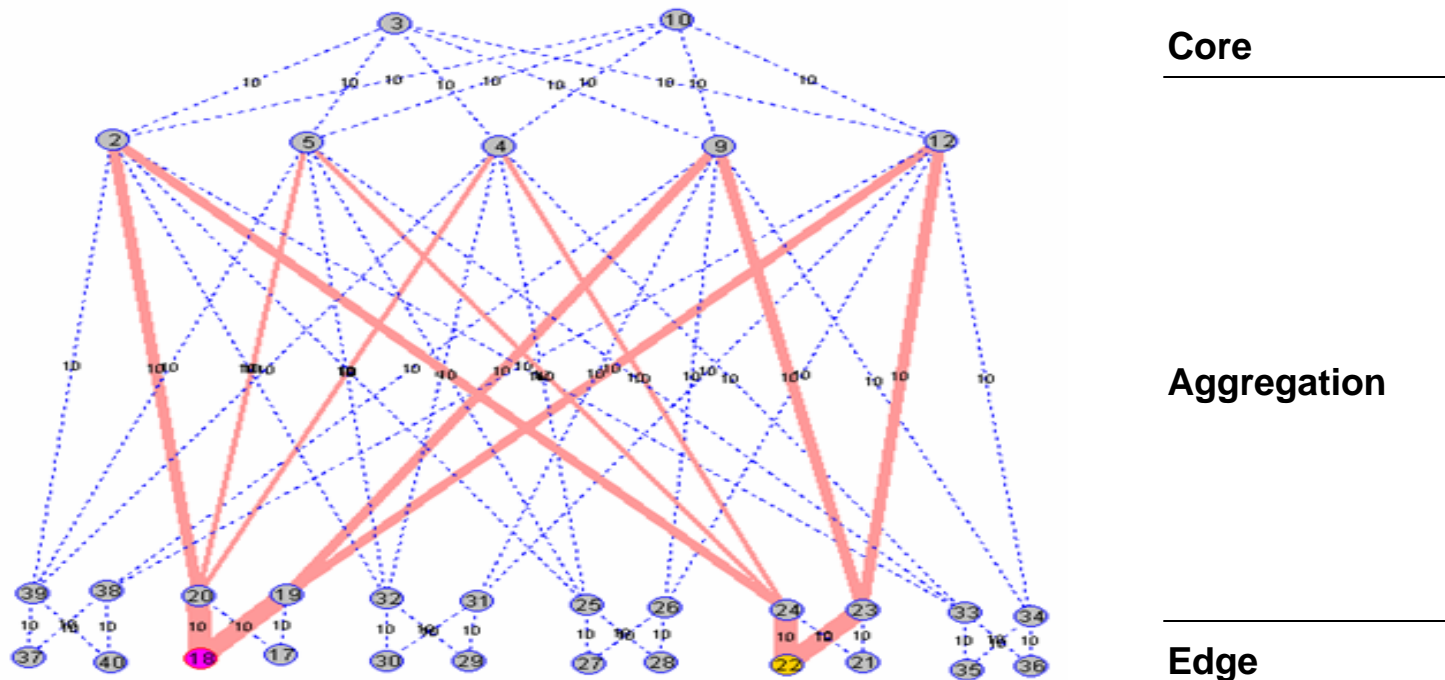
- Data center typically deploys a *fat-tree* architecture;
- The server may be further decomposed into a hierarchy of bridge-like devices (VEBs/VEPAs) providing connectivity to Virtual Machines at the network edge;

Fat-tree Architecture




- We ignore the decomposition of the server;
- We assume only that the servers exchange C-tagged traffic with the network-edge;
- Key objectives in this ‘wired’ portion of the data center network are:
 - Short paths;
 - Traffic load distribution among the short paths;

SPBM in the Data Center



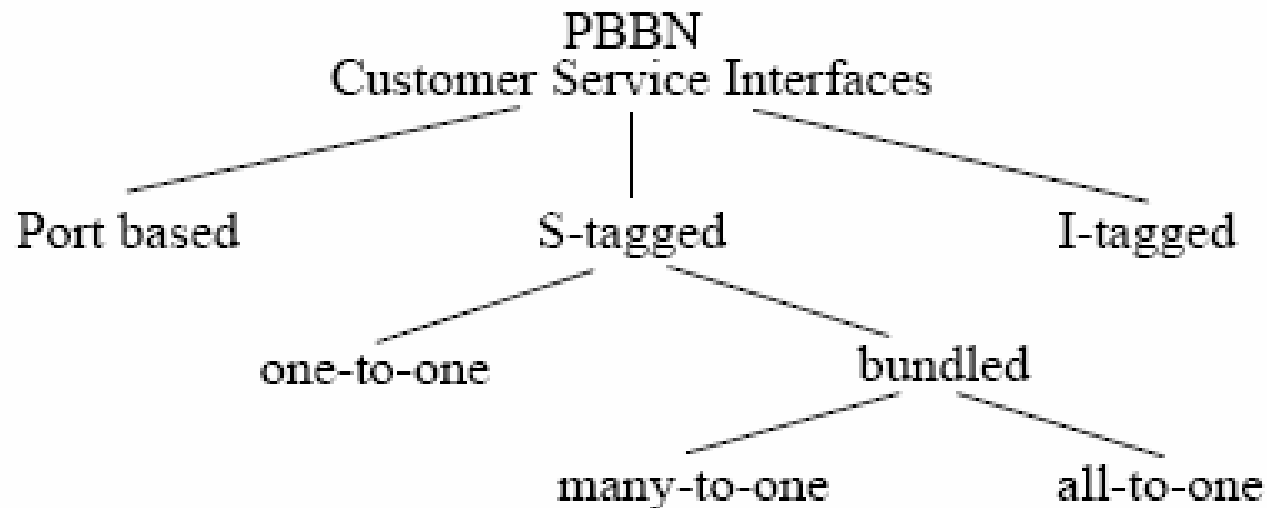
- It has been demonstrated that the MAC (group address) form of SPB (SPBM) provides a practical and efficient means spreading traffic load across a set of 'shortest paths';
- See, for example <http://www.ieee802.org/1/files/public/docs2009/aq-ashwood-ECT-framework-1109-v2.ppt>

The Problem



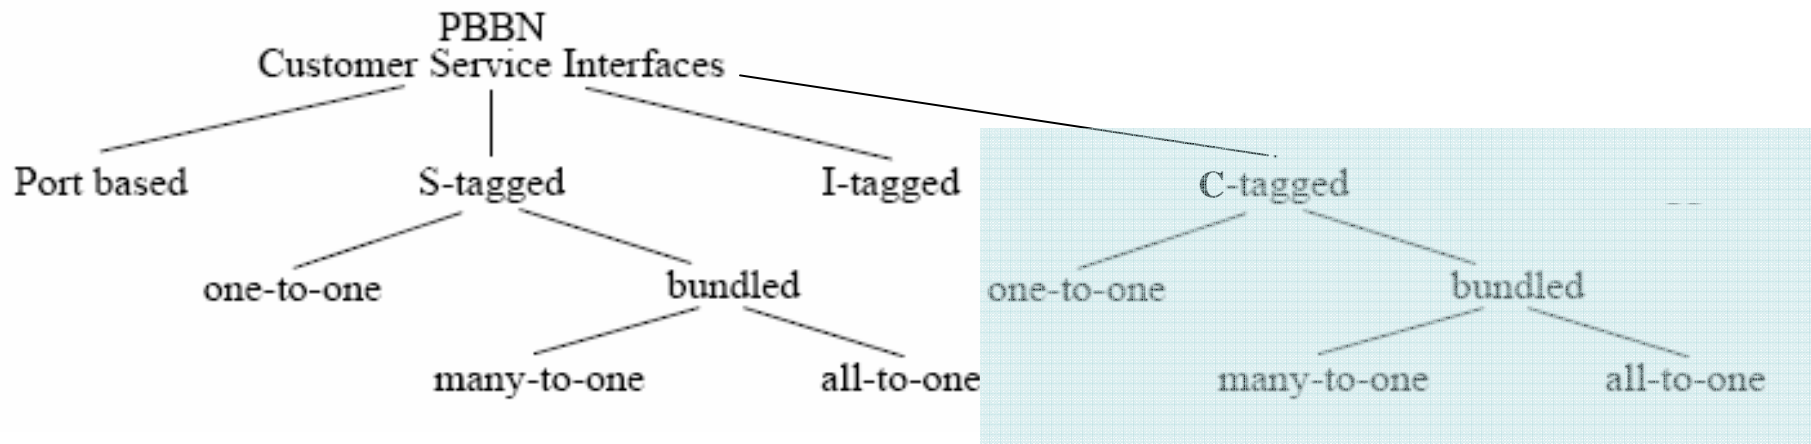
- The data center is a *Customer Bridged Network*;
 - Traffic exchanged with the network edge is C-tagged;
- SPBM requires that the red circled edge bridges be PBBN Backbone Edge Bridges;
- Currently, a BEB supports C-tagged traffic only by means of a Port-based Service Interface;

The Port-based Service Interface



- Port-based service interface causes *all* traffic received on a port to be mapped to one Backbone Service Instance (ISID) associated with that port;
- SPBM traffic load distribution (ECMP) is defeated;

Proposed Straightforward Fix



- Extend 802.1ah with amendment to .1Q adding C-tagged Service Interface with function analogous to S-tagged interface;
- Architectural approach is to define a type of I-component supporting C-VLAN relay;

How much effort?

- A draft of what this amendment might look like is posted as `new-sultan-cbb-0510-v01`;
- Please have a look – the changes appear to be reasonable with this approach;
- A preliminary proposed PAR is posted as `new-sultan-cbb-par-0510-v01`;

The Reality

- The problem described here is ***not*** a problem in the field;
 - Current deployments of PBBN/SPBM *do* map C-VLANs to Backbone Service Instances;
- This PAR is a matter of bringing the architecture in line with the reality;

Conclusion

- This PAR should be of great interest to anyone anticipating the use of SPB in the Data Center;
- It provides a straightforward fix to a problem that would greatly reduce the benefit of deploying SPB in the Data Center;