

BB6 Requirements 11-218v0

Siamack Ayandeh
Chief Architect (Datacenter)
HP Networking
May 25th, 2011
Siamack@HP.Com



Datacenter Ethernet Trends

- Datacenters are moving towards large port count flat layer-2 Ethernet fabrics that:
 - Take advantage of regular flat topologies that are well known and understood within the telecom industry e.g. Clos and its variation Fat-tree
 - Deal with the table size limit of Ethernet switches using some form of encapsulation
 - May involve separation of host MAC from the location MAC to enable virtual machine mobility
 - May involve use of a hierarchical addressing format that implies the path & does away with transient loops
 - Discourage broadcast storms
 - E.g. ARP proxy/server to prevent ARP storms
 - Use path/multi-path calculations of IS-IS or OSPF or a path hunt algorithm for known topologies
 - Support virtualized hosts environments with many hosts sharing a physical port and handling encapsulations
- L2 fabrics are required to provide scale and services such as lossless transmission, load balancing, loop free, virtualized host mobility, security ACLs



Requirements for BB6

- ***Be compatible with and take advantage of emerging L2 Ethernet fabric services***
- ***Achieve separation of control and data plane***
 - Make FCF a fabric service rather than an element of the forwarding plane
 - Allow hosts to discover each other, perhaps using an FCF, and then have direct VN to VN port communications over Ethernet L2 fabric for data transmission
 - No more L3 forwarding during data transfer phase
 - Data transmission L2 path is under control of Ethernet fabric (NOT FSPF)
- Take advantage of the L2 Ethernet fabric using Ethernet switch capabilities as opposed to L3 FC functions embedded in Ethernet switches to support:
 - Virtualized hosts
 - Security – zoning
 - Path calculations



Trouble with multi vendor BB6 Dist-FCF

- Dist-FCF proposal breaks up a switch to control component and forwarders potentially from different vendors
- Based on past experience, this type of architecture ends up to be single vendor:
 - Very difficult to debug across multi-vendor boundary e.g. experience with FC E-ports showed how difficult it is to get multi vendor working for FC
 - *E-Ports are still not interoperable*
 - May have been workable a few years back with a different competitive landscape
 - *Multi vendor technical feasibility is currently in question*
 - Multi vendor results in loss of value-add features both FC and Ethernet removing incentive for customers to buy such a solution
 - *Questionable market potential*
 - Customers also double the number of software stacks that they need to maintain despite the de-featuring that occurs
 - *Questionable TCO*



Single vendor Dist-FCF, not an urgent topic for standardization

- In the absence of multi vendor implementations the internal control architecture of a switch is not a high priority subject for standardization
 - IEEE 802.1Qbh contributions on port expanders for an Ethernet switch; concentrate on server/virtualization interface; not internal control plane of a distributed switch



Incompatible with industry efforts to scale L2 fabrics

- BB6 ingrains FC control stack and forwarding within Ethernet switches
 - ***IT IS MERGING AND NOT CONVERGING***
- You have two of everything side by side; even with FCoE based storage
 - FC as well as Ethernet control stack
 - L3 FC DID based as well as L2 Ethernet based look up & forwarding
 - FSPF routing protocol as well as say IS-IS for Ethernet fabric
 - FC zoning based on WWN and FC addressing as well as Ethernet ACL's based on MAC
 - All above lead to two distinct management solutions with constant need of mapping from one to the other
 - Not only you need a SAN manager and a data network manager – THEY NEED TO TALK



Thank you

Questions on Design Intent (DI)

DI #1 – *BB6 deals with the domain ID count limitation of FC fabrics*

- BB6 simply expands the number of ports per FCF using a port expander concept
- The limited number of domain-IDs in a fabric, likely a control plane scaling issue, still remains

DI #2 – *FDF's can exchange traffic over virtual Ethernet A-links hence reducing the number of hops*

- Hop count reduction's vary and depend on the network design and patterns of communication
- The number of L3 hops does not change between BB5 and BB6
 - The number of L3 hops can actually increase in some scenarios
- The number of L2 hops depends on the design of the loss less Ethernet cloud and is un-known to BB6
- If I-T pair are out side the scope of a single switch inter FCF traffic still goes through FCF and virtual E-ports