



Thank You, Cloud Computing!

Lane Patterson
April 14, 2010

Key Biz Drivers for Next-Gen IX



- Give me more 10G ports, cheaper
- I'm desperate for 40GbE or 100GbE as soon as it is affordable on my router interfaces
- Can you extend your fabric to multiple locations that interest me, with good redundancy/scaling topology?
- Make sure there are no bottlenecks
- I don't want it going down, failures should heal fast
- Can you give me truly transparent virtual private interconnects on same interface/fabric as the public VLANs?

Or in short...

“Holy crap, transit is dropping fast—I need cheap, high-speed infrastructure!”

- **LEGACY:** A few decent legacy switches:
 - STP, LAG/LACP and other basic ethernet switch features
 - IX features like SFLOW (or equiv), Port Security
 - Port density up to 140 x 10G
- **VPLS:** Only one “switch-priced” VPLS-capable solution
 - AMS-IX has shown how to do multi-stage access/core to achieve strong scaling with VPLS
 - Can combine several VLANs as well as private peering pseudowires
 - Best proven platform for uber-scaling: 128 moving to 256 x 10G ports
- **EMERGING DATA CENTER FABRICS:** A few new solutions will provide an alternative to VPLS scaling:
 - Data center ethernet features like PBB-TE or Trill variants (MAC-in-MAC), Ethernet OAM, low latency with per-priority queuing and flow control
 - Excellent densities: 300-600 ports/chassis
 - About to commoditize price/port of 10G to sub-\$1000

The MPLS/VPLS Approach



- **Scalability:**
 - MPLS can provide multipath trunking topology, buying you L2 ECMP
 - This is critical for multi-stage scaling
- **Resilience:**
 - MPLS provides FRR
 - This is critical for 50ms protection
- **Topology control/virtualization:**
 - VPLS allows VLAN mapped over MPLS
 - VLL (Pseudowire) allows point-to-point interconnect mapped over MPLS
- **Future Evolution:**
 - MPLS-TP, Seamless MPLS, etc. will continue to enhance MPLS approach for some time

- **IEEE 802.1ah MAC-in-MAC: PBB**
 - Defines basic approach of Ethernet-based Backbone “tunnels”, and bindings between customer-domain MACs and Backbone MACs for forwarding
- **IEEE 802.1Qay: PBB-TE**
 - Provides “Traffic Engineering” for these Tunnels, e.g. can define their path, setup backup paths, L2 ECMP, etc.
- **IEEE 802.1aq: SPB (SPBV for Q-in-Q; SPBM for MAC-in-MAC)**
 - Provides the control plane to compute PBB-TE paths
 - IS-IS for MAC address routing
- **IETF TRILL**
 - Similar alternative to the IEEE standards
 - Purists argue that PBB handles multicast better

Metro + Data Center Ethernet (cont'd)



- **802.3ah, 802.1ag, ITU-T Y.1731**
 - Ethernet OAM, notably low-granularity (10-100 ms) Ethernet Continuity Check Messages (CCM)
 - Keepalives that trigger sub-50ms failover using TRILL or PBB-TE
- **IEEE Data Center Bridging (DCB) standards:**
 - Needed for stringent forwarding requirements of low-latency, converged (Fibre Channel-capable) Ethernet fabrics
 - IEEE 802.1Qbb: Per-priority PAUSE frames
 - IEEE 802.1Qau: Per-priority Congestion Notification
 - IEEE 802.1Qaz: Enhanced Transmission Selection—framework for assigning the P-bits
 - IEEE 802.1Qaz: DCB Exchange—discovery of above capabilities between switches in the same fabric
- **Multi-chassis LAG**
 - Allows a LAG to uplink to 2 switches for access redundancy

Why do we care?



- In short, why are these standards of interest?
- Because, within 6-12 months, best hope of commoditizing the IX requirement set
- Because, these platforms will be the first with 40/100GbE density at a good price

We Love You, Cloud Guys!