Thoughts on Edge Virtual Bridging

Joe Pelissier

new-pelissier-EVBThoughts-0509

EVB Weekly Discussion

- EVB has been meeting weekly
- Three major problem / solutions have been discussed:
  - Embedded Bridges (aka VEB)
  - VEPA
  - Interface Virtualizers
- In addition, various management paradigms have been discussed
  - Applicable to all of the above

Embedded Bridges

- **In virtualized environments, bridges may be embedded in servers**
  - **Hardware or Software**
- **These bridges can be fully .1Q compliant**
- **However, many operate a little bit different from the standard:**
  - **Many have taken advantage of close integration with the hypervisor and VNICs**
    - For example, eliminates need to learn / age
  - **In addition, being edge devices, these bridges may not forward between uplinks**
  - **There seems to be commonality in these functions and such bridges have been deployed and have been proven useful**
    - Might make sense to standardize
    - However – these bridges do operate a little bit different, so they do add somewhat to management complexity

VEPA

- **VEPA modifies the behavior of an embedded bridge**
  - **In general, embedded bridge in ingress performs its normal functions, then forwards frame to adjacent (external) bridge**
  - **External bridge *augments* functionality of embedded bridge**
    - Packet processing (TCAMs, ACLs, etc.)
    - Security features such as: DHCP guard, ARP monitoring, source
    - port filtering, dynamic ARP protection/inspection, etc.
    - Enhanced monitoring capabilities e.g. statistics, NetFlow, sFlow, rmon, port mirroring, etc.
  - **External bridge forwards frame back to VEPA ("Hairpin turn")**
    - VEPA forwards frame to destination similar to any other bridge
      - MAC/VLAN lookup, etc.
- **Defines two new (relatively simple) behaviors**
  - **Embedded bridge: forward frames externally**
  - **External bridge: hairpin turn**
- **New behavior complicates network management**

Interface Virtualizer

- High density server deployment (including but not limited to virtualized servers) creates a proliferation of bridges in the network
- Many of these bridges are operating largely as a simple mux
  - Essentially operate as fan in / fan out to higher level bridges
- Yet these bridges are responsible for a significant proportion of the network's capital expenditure, operational, and management costs
- Interface Virtualizers replace these specific bridges collapsing the number of bridges in the network
  - IVs essentially become ports of the bridge to which they are attached
  - Not independently managed; managed much like a line card in a bridge
  - Much simpler (i.e. more cost effective device)
- Intended for use in the "branches and leaves"
  - Not exclusively used at the end station
- Reduction in network complexity and associated management

Observations

- All three devices provide independent and valuable benefits to networks
- All three solve separate problems
  - VEPA more-or-less a superset of embedded bridging function
- None of these devices effectively address the issues addressed by the other two
- All three compliment and interoperate in the same network cleanly
  - Various combinations may be mixed and match for optimal usage in any given environment
- All three appear to have strong commitment by individuals to complete standards work

Proposed Next Step

- **Develop appropriate PAR and 5C (or set of PARs and 5Cs) for next meeting**
- **Work together to make all three efforts successful**

Thank You!