Simplifying Network Interface Virtualization

Joe Pelissier

Andrew Lambeth

new-dcb-pelissier-NIC-Virtualization-0908

Virtualization Results in Complex Bridge Hierarchy

The Hypervisor LAN Switches (soft bridges) are taking on increasingly complex functionality
-ACLs
-VLANs
-Security
-Congestion Notification
-Priority Flow Control
-Enhanced Transmission Selection
-Etc.

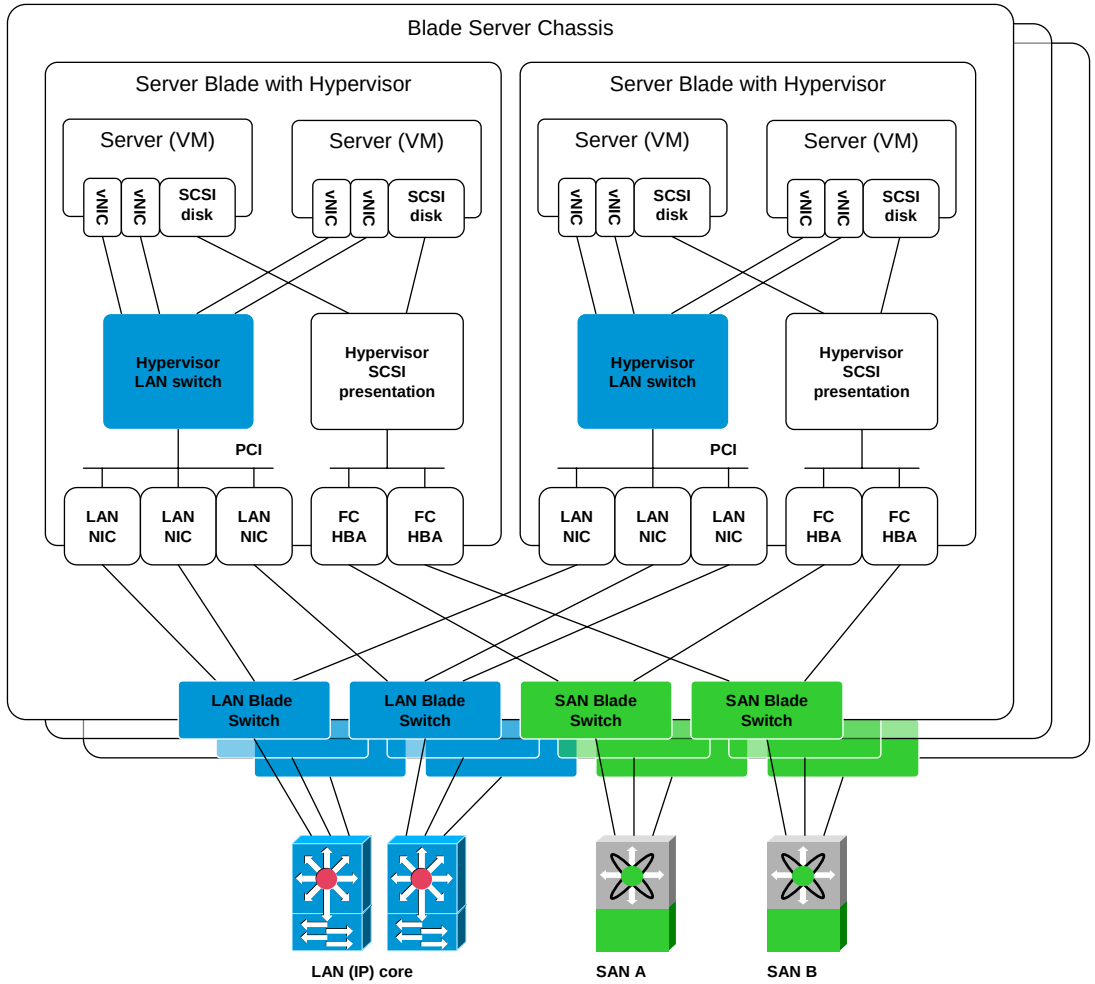

A Possible Solution…

Delegate complex and performance critical data path functions into upstream networking devices (bridges)

- Ensures feature consistency to all traffic
- Fewer bridges – simpler and more consistent management
- Better performance and scalability

NICs provide value add data movement and ULP features

- TCP offload, RDMA, FC/SCSI DDP, IPC queue pairs, etc

Hypervisors provide features based on visibility of host state

Embedded (soft and blade) Switches Replaced with simple “Interface Virtualizer”

Interface Virtualizer (IV) aggregates downlinks (connections to physical or virtual NICs) into uplinks (connections to Virtual Interface Aware Bridges (VIB))

-Outbound traffic is simply passed through to VIB

-Inbound traffic is forwarded from the VIB to the proper downlink

-A link-local tag is used between the IV and the VIB to indicate the downlink port from the IV

- **-Eliminates need for CAM in the Interface Virtualizer**

-Complex bridge functions (ACLs, VLANs, etc.) remain in the VIB

-No fundamental change to Bridge forwarding or operation

- **-Downlink ports are learned along with MAC addresses**
- **-Frames transmitted to the IV receive a simple tag with the downlink port**

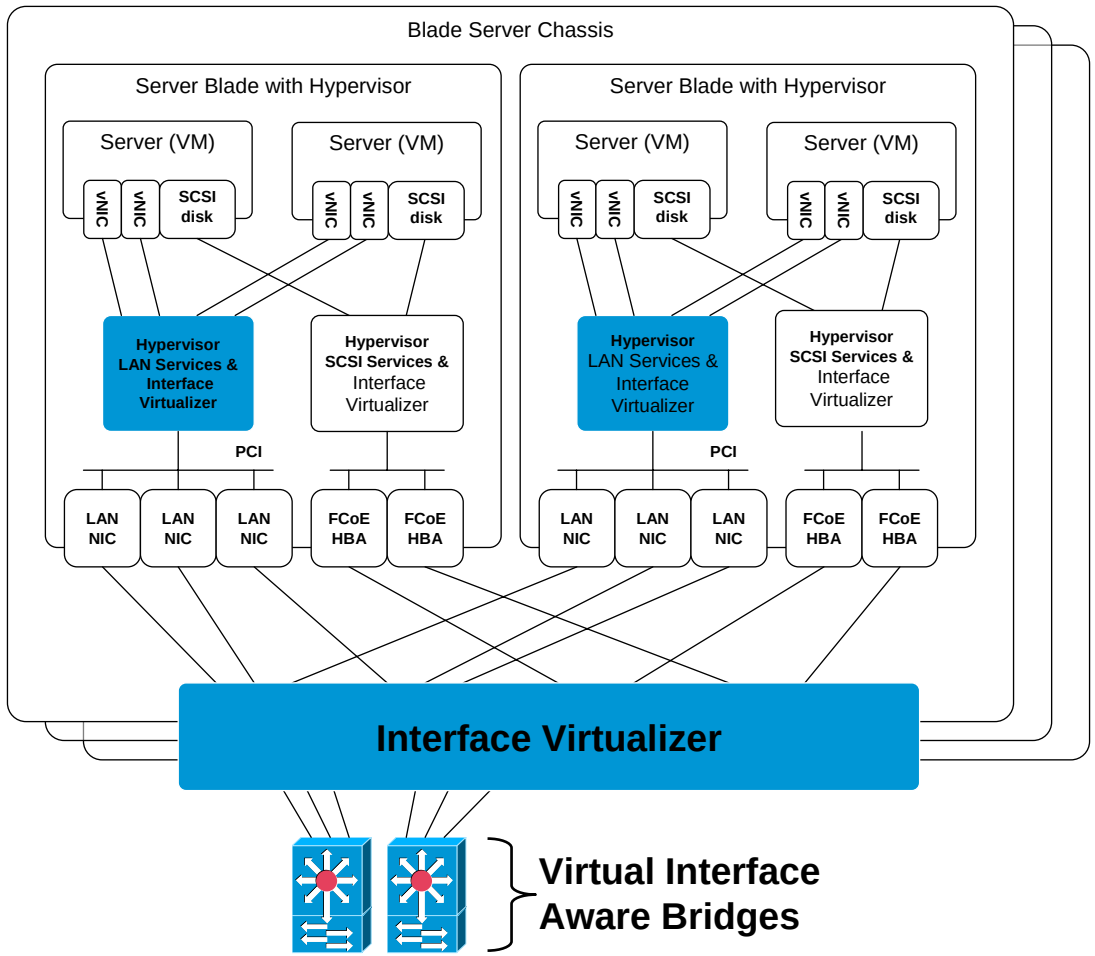

Interface Virtualizer

- **Additional Functions (to be discussed in greater detail during the Dallas DCB meeting)**
  - **Multicast Replication**
  - **PFC Queuing**
  - **COS capabilities**
  - **Protocol Communication with Virtual Interface Aware Switch**

Thank You!