



# Thoughts on Network Interface Virtualization

Joe Pelissier


new-dcb-pelissier-NIC-Virtualization-0708

# Current state of evolution

## Standalone Servers and Individual Adapters

Servers connect to bridges with an OS interface per port

Network policies applied to bridge ports which are equivalent to OS interfaces




## Virtualized Servers and Consolidated IO

Virtual machines connect to bridges in hypervisors

SR-IOV adapters with integrated bridges

Multiple interfaces from multiple OSES per bridge port

# Use Cases for Virtualized Adapters



**Functionally consolidated IO devices**

**Ethernet & Fibre Channel**

**Multiple interfaces for single OS servers**

**Multiple LANs...**

e0 – system management VLAN

e1 – DMZ VLAN

e2 – oracle traffic



e3 – all other traffic

**... and SANs**

SAN A and SAN B

**Interfaces to virtualized servers**

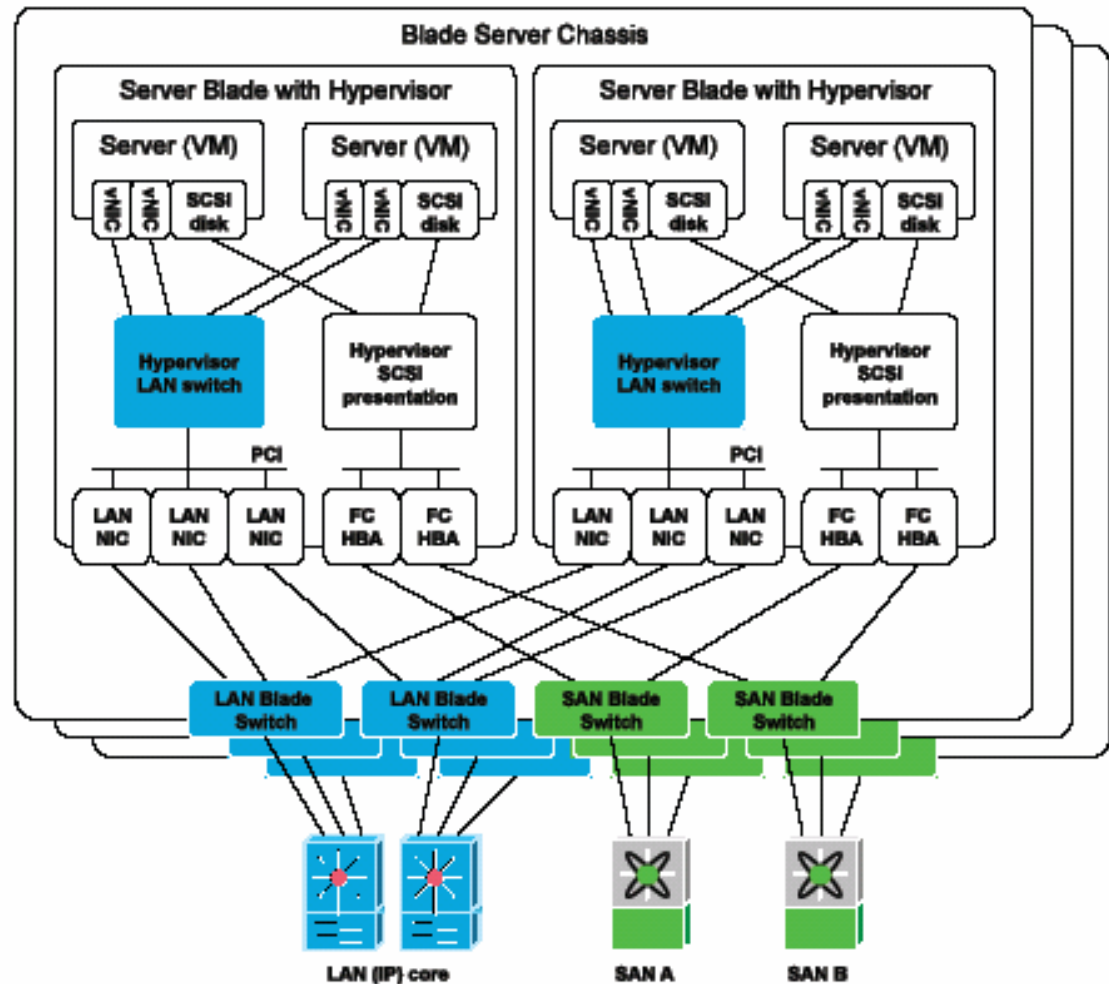
**Reduces overhead of data movement**



# Virtualization Results in Complex Bridge Hierarchy

The Hypervisor LAN Switches (soft bridges) are taking on increasingly complex functionality without the performance benefit of dedicated hardware:

- Scalability reduced
- Performance reduced
- Management complexity increased



# Evolving Issues

- **Softbridge performance is not scalable**
  - Especially when policy enforcement is involved
  - Use of more than 10% of CPU by softbridge becomes a significant barrier to acceptance
- **When VM migration is common, uniform policy enforcement is important:**
  - Across VMs to be same as across physical servers
  - Embedded softbridges complicate policy enforcement and reduce scalability

# A Possible Solution...

**Delegate complex and performance critical data path functions into upstream networking devices**

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**

# For Consideration...

- Introduce a new on-the-wire indication to identify a virtual adapter endpoint within a data frame

For example, a new tag similar to QTag

**Specify a the use of such an indication between NICs supporting virtualized interfaces and an upstream bridge**

Enable upstream bridge to perform data plane functions such as:

Filtering

Access Control

Security Functions

**Define appropriate supporting control plane and management protocols**

Thank You!