

All Roads Lead to Convergence

Greg Scherer

VP, Server and Storage Strategy

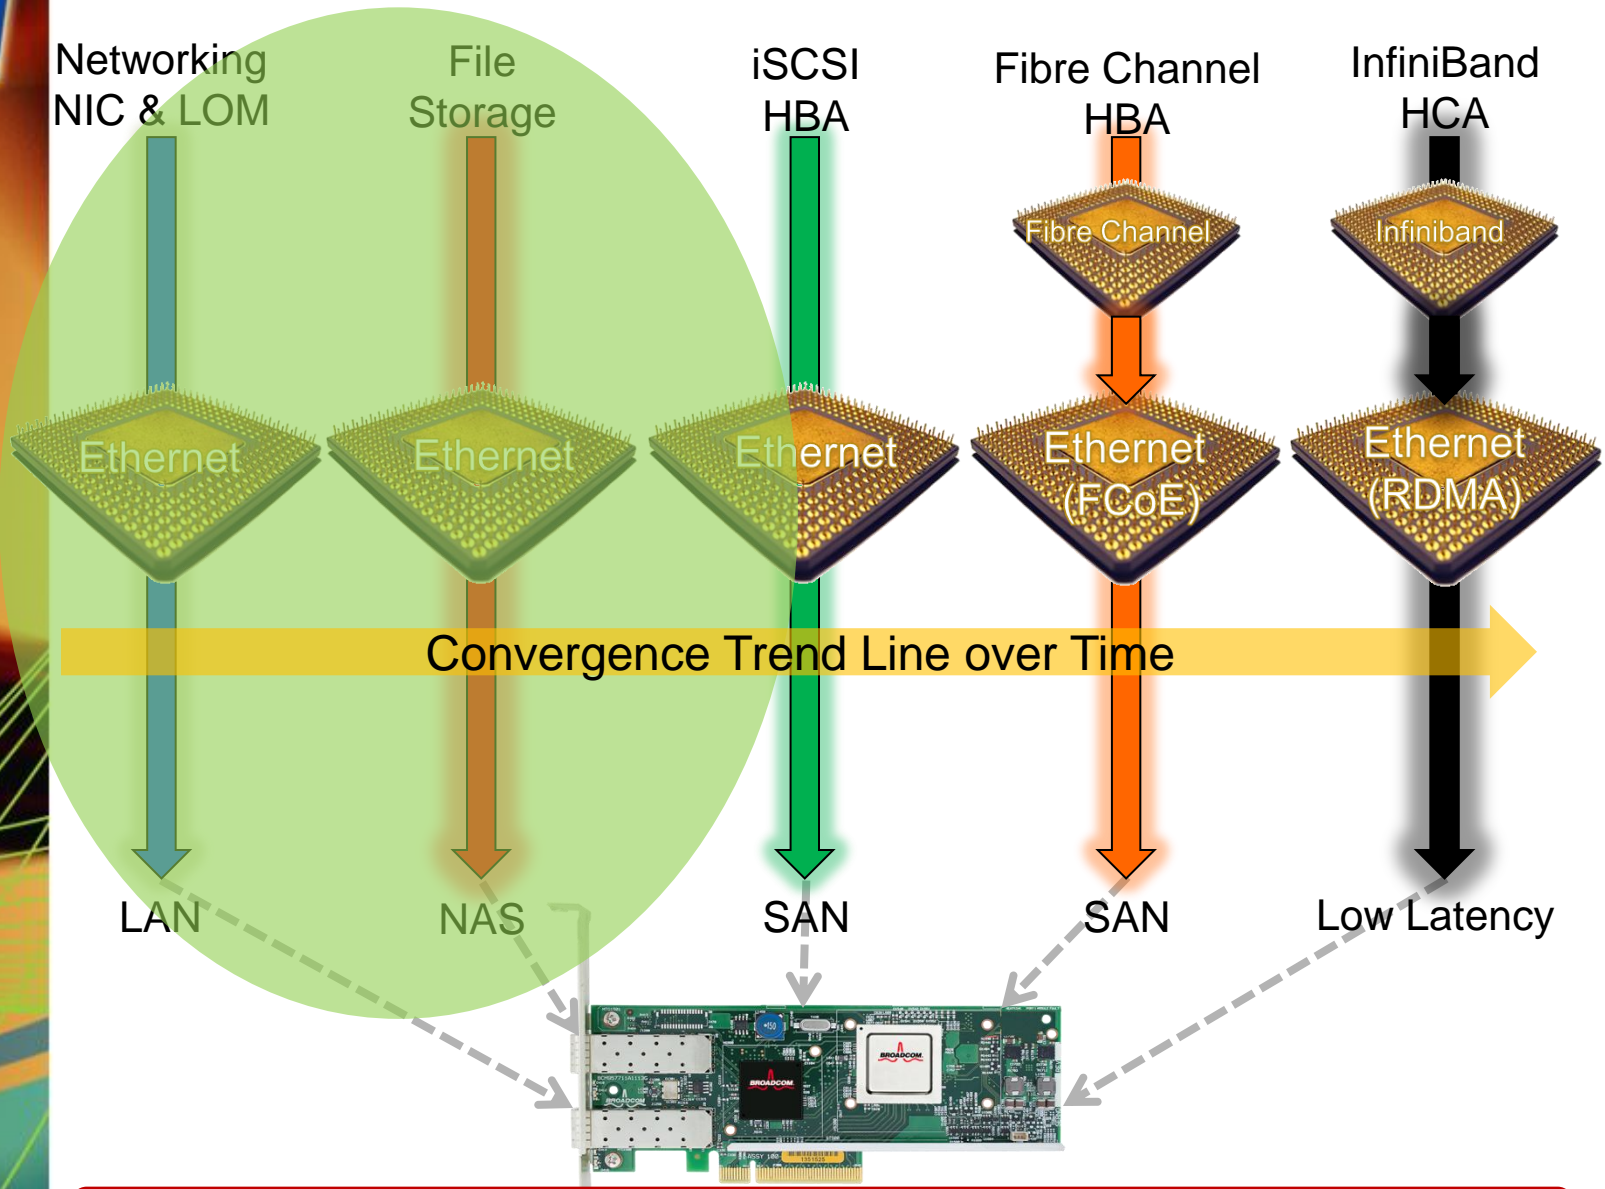
gscherer@broadcom.com

Broadcom Corporation

Agenda


- The Trend Toward Convergence over Ethernet
- Reasons for Storage and Networking Consolidation over Ethernet
- 10GbE Deployment Evolution
- Architectural Requirements for Convergence
- Summary

All Roads Lead to Ethernet for Enterprise I/O



Ethernet has Been a Converged Network Since its Inception

Networking & Storage Convergence – When?



The Trend is Clearly Toward Converging Storage over Ethernet
(by 2014 ~70% of all non-DAS Server Ports will be Ethernet)

Networking and Storage Convergence – Why?

1. *Ethernet*

Open, Interoperable, Simple, Low Cost, and High Volume!

2. *Unified Fabric & Management*

Significant Opportunity for Lowering OPEX/CAPEX Costs

3. *Virtualization*

Network Storage is Imperative for Seamless Migration, DR, etc.

4. *Scalable Bandwidth*

Clear Path Beyond 10Gb to 40Gb, 100Gb, etc.

5. *Continued Innovation*

Loss-Less (DCB), Energy Efficient Ethernet (EEE),
Congestion Mgmt (802.1Qau) , Advanced Routing (TRILL), etc.

Ethernet : Technology Leadership Driven by Economics

Server Virtualization is Driving 10G+ Technology vectors driving Convergence

1. VM Density

16+ VMs per server common
Usually requiring 1GbE per VM
CPU performance @ or >10GbE

2. Best Practices

4-8x GbE LAN / NAS / iSCSI
2x GbE Management
2x GbE VMotion

3. Network Storage

85% of VM installs use
Network storage
(allows for VM migration)


Adapter, cable consolidation
Reduced power and cost
Improved performance

NIC Partitioning w/ QoS allows
1Gb deployment model w/
better performance, cost and Mgmt

Enhances virtualized deployment
iSCSI and NAS 10GbE performance
Converged fabric – Cost/Power/Mgt

Networking:

Direct Queues
NIC Partitioning
SR-IOV

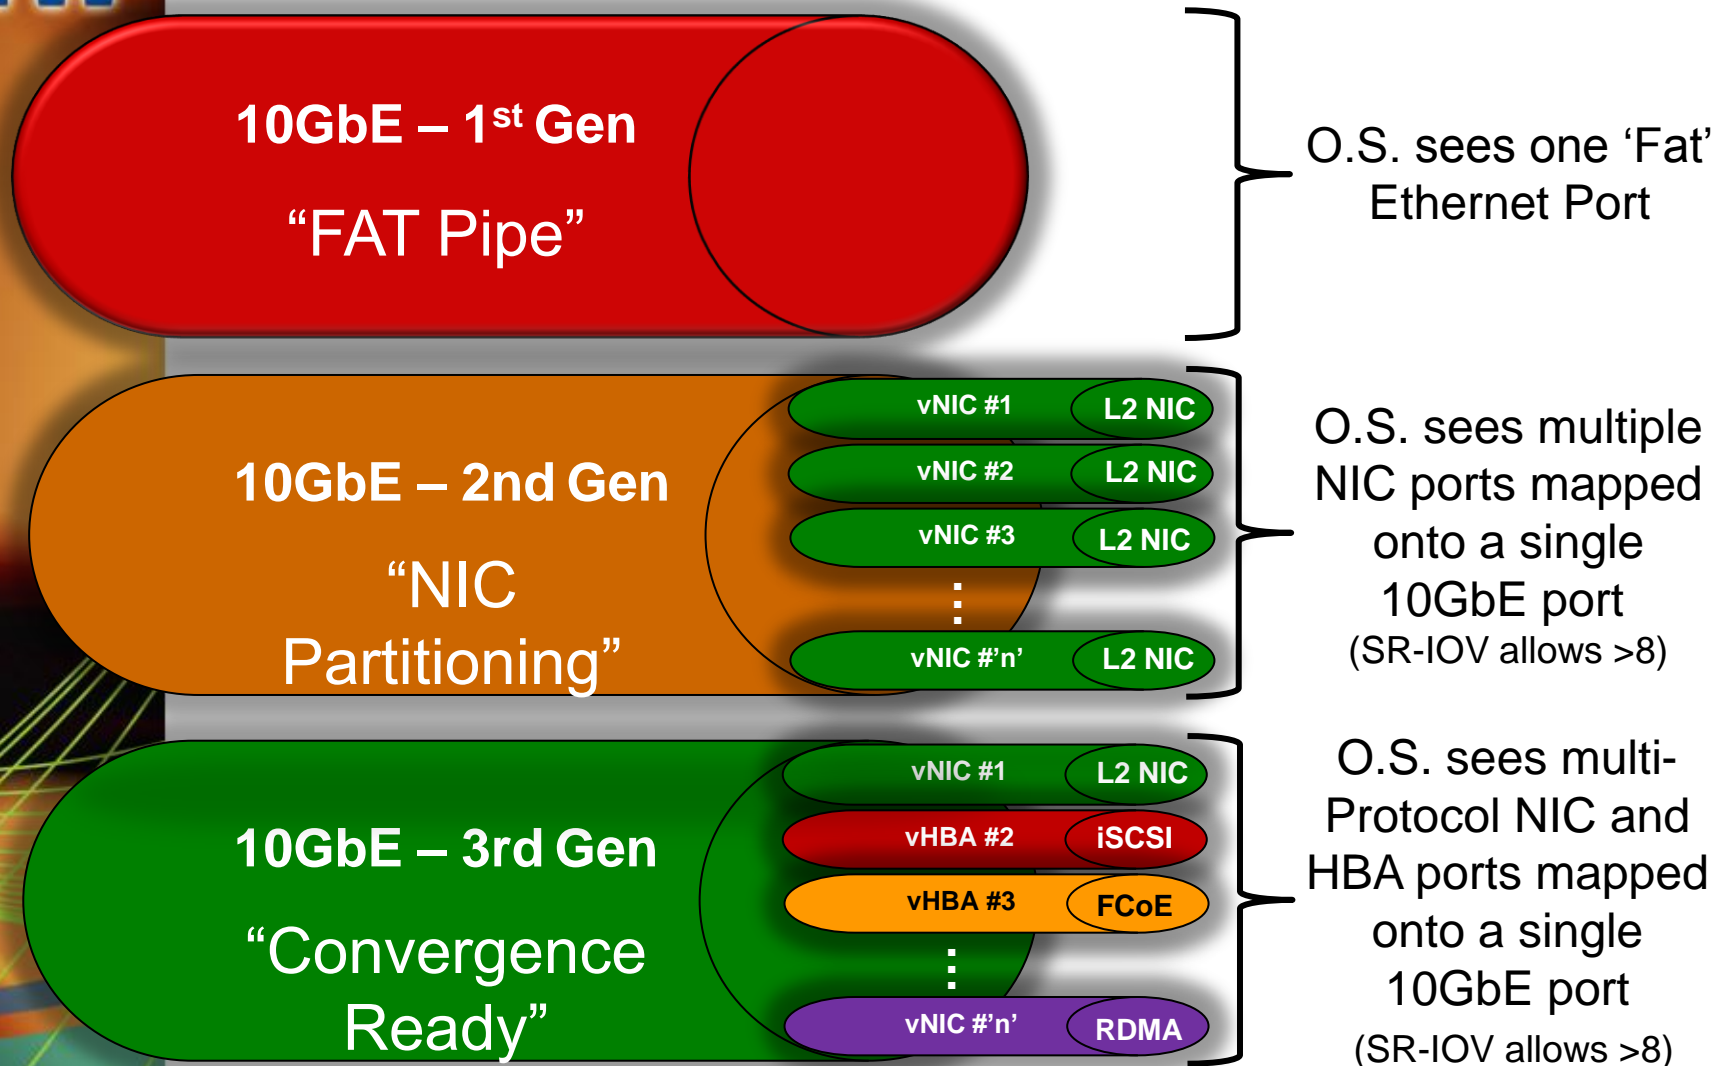


Storage:

NAS, iSCSI &
FCoE

Server Virtualization

10Gb Ethernet Controller Deployment Evolution

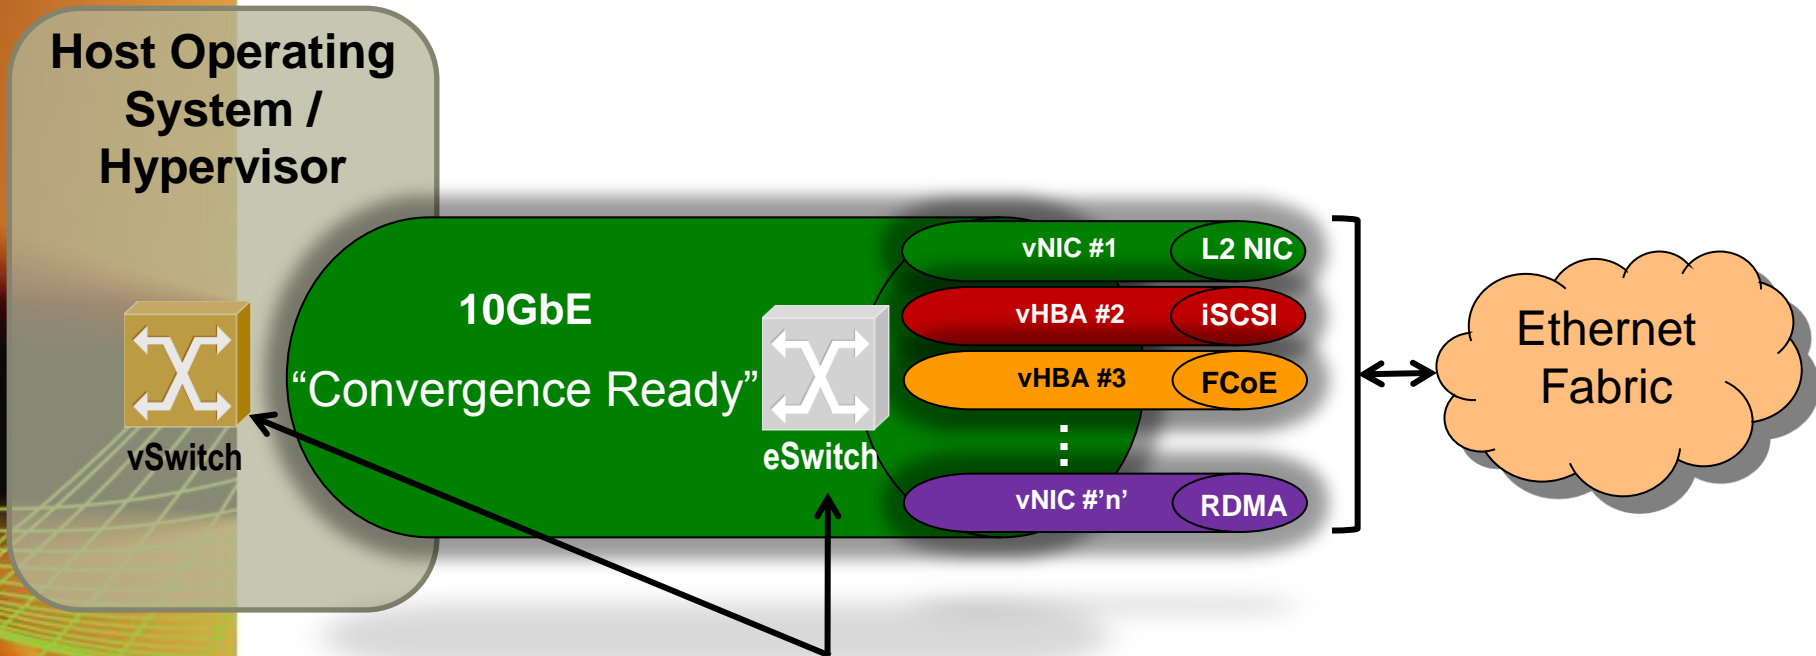


Virtualization of the Ethernet ‘Pipe’ is Required to *Digest* 10Gb -> 40Gb

Server Virtualization:

Migrating the new Switching “Edge” Inside the Server

- Support for internal switching between vNICs / vHBAs
 - In the Hypervisor “soft-switch”
 - In the Controller






New switching points inside the ecosystem

New Switch Edge adds Key Capabilities *and* Management Demands

As Consolidation and Convergence Occur Performance is Key

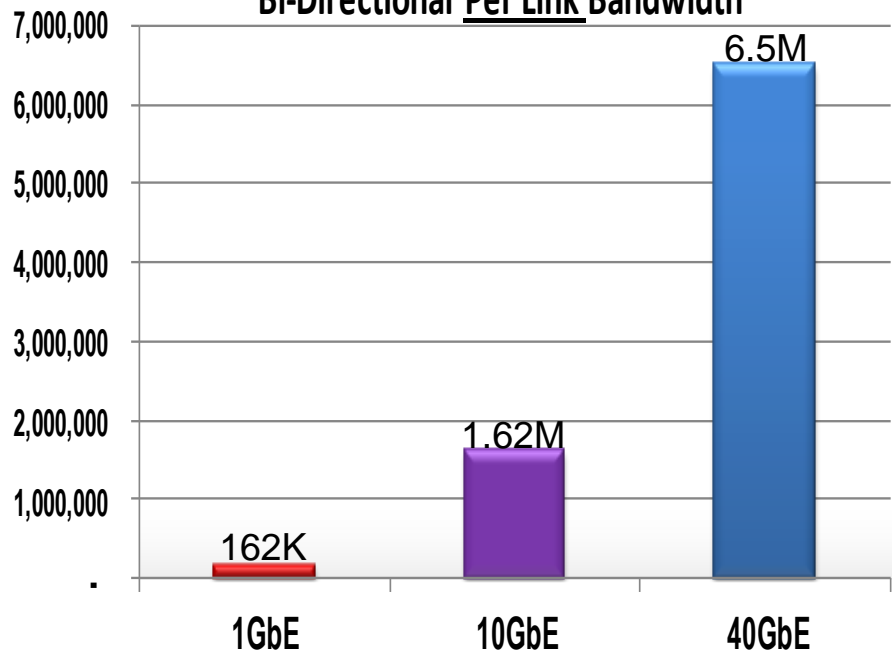
- Server Consolidation
 - Each of 16+ individual server workloads are now being consolidated on a single Virtual Server
- VM Densities are Increasing
 - The # of VMs per Virtual Server are growing at a large CAGR (>25%)
- Application I/O needs are Growing
 - Database
 - Clustering
 - “Data hungry” applications



Ethernet Performance Requirements: 1Gb, 10Gb and 40Gb Links

Full Size L2 Packet Performance


#of 1518 Byte Packets/Second Required to Fill
Bi-Directional Per Link Bandwidth



6.5M 1.5KB L2 Pkts/Second @40Gb

Medium-size L2 Packet Performance

#of 536 Byte Packets/Second Required to Fill Bi-
Directional Per Link Bandwidth



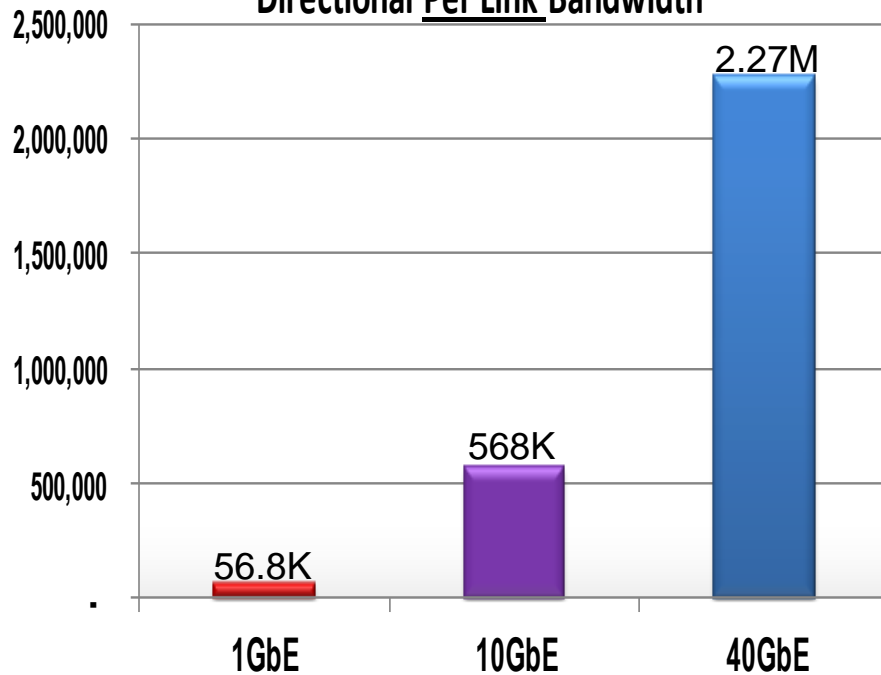
18M 536B L2 Pkts/Second @40Gb

Layer 2 Packet Performance is Driven up Dramatically with 10Gb and 40Gb

SAN I/O Performance Requirements

4K Byte I/O Performance

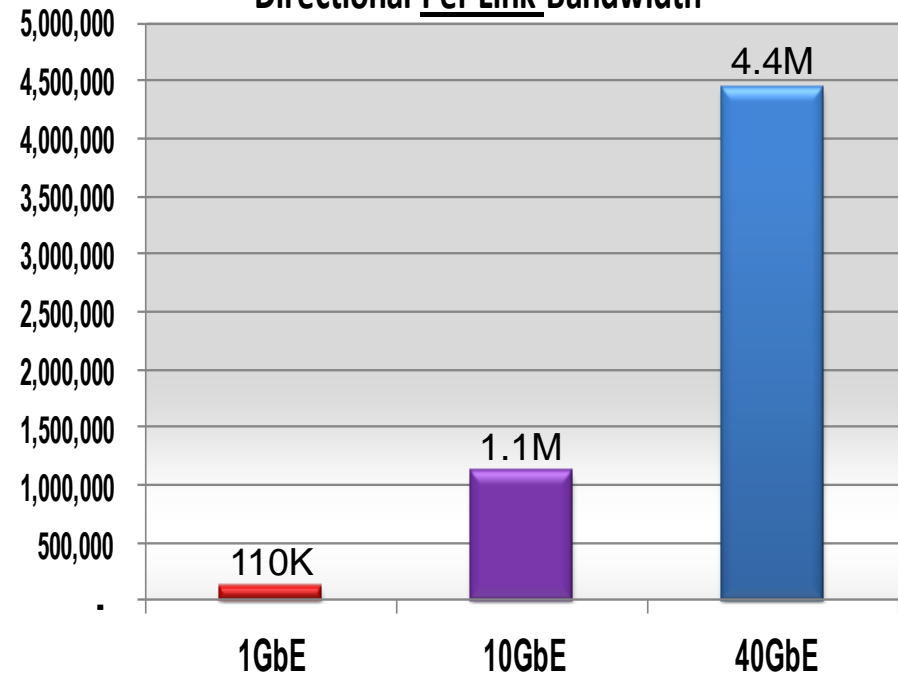
#of 4KB I/Os per Second Required to Fill Bi-Directional Per Link Bandwidth



2.27M 4K I/Os per Second @40Gb

2K Byte I/O Performance

#of 2K I/Os per Second Required to Fill Bi-Directional Per Link Bandwidth



4.4M 2K I/Os per Second @40Gb

SAN I/O Performance Requirements are Enabled by SSDs and Fast Arrays

Requirements: Architectural Implications

Converged and Virtualized Networking Platform

Performance and Efficiency Across all Protocols and Environments


L2 Networking

Layer 2 – Data Link

Layer 1 – Physical


- Latest PCIe
- Stateless offloads
- Multiple teaming
- DCB
- RX/TX coalescing
- Deep RSS
- Management

Virtualization



- MSI-X (Multi-core)
- NIC Partitioning
- Multi-queue
- SR-IOV
- Internal switching

Convergence



- NAS & CIFS
- iSCSI HBA
- FCoE HBA
- TOE
- RDMA
- 1588 and 802.1AS

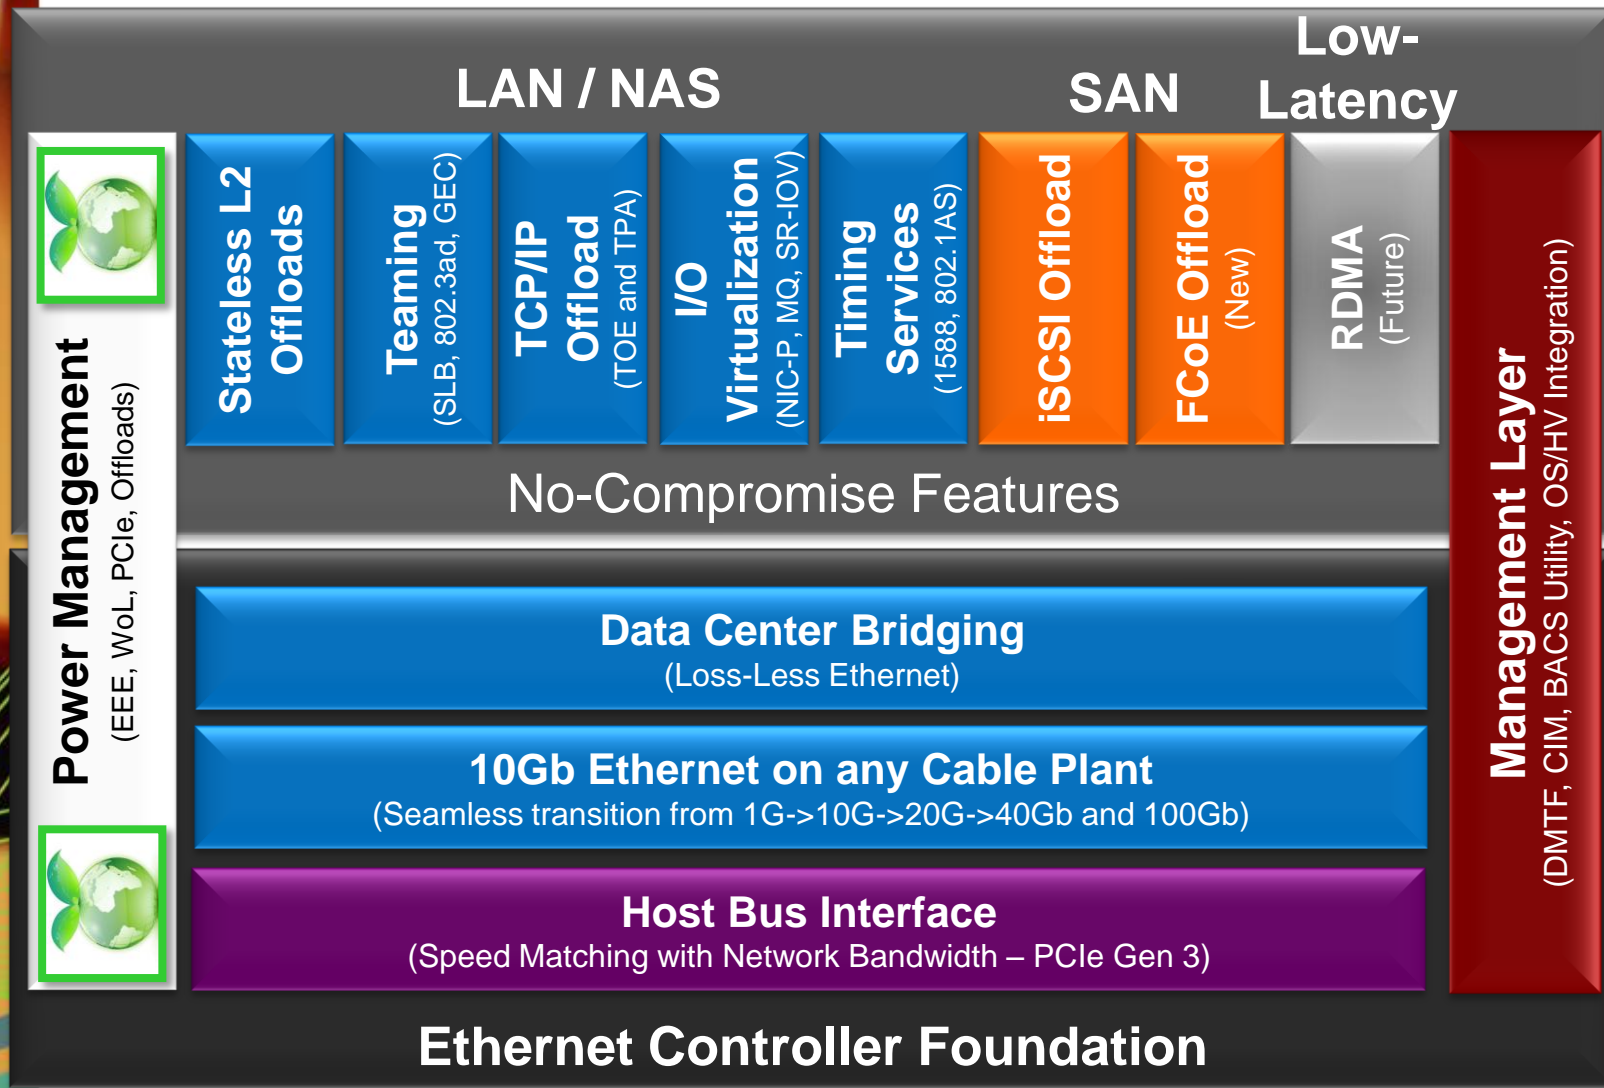
Power Management



- Hardware offloads
- PCIe ASPM and CLKREQ#
- WoL (Wake on LAN)
- Energy Efficient Ethernet
- Passive cooling

Industry Standard Management (DMTF, SNIA, EVB, DCBX, NC-SI, etc.)

Controller Architectural Vision to Support Convergence



Common Platform Architecture and Compatibility 10G -> 40G

Summary

- Storage and Networking ARE and HAVE been Converging
 - Infrastructure comes first; deployment behavior comes later
 - FCoE is just the next workload to migrate to Ethernet
- Ethernet Continues to Provide Leading Innovation at Great Economics:
 - Ever increasing BW, New Features, and Virtualization
- Convergence Requires new Architectures to be Deployed
 - Loss-less Ethernet, L2, Virtualization, and Storage Workloads
 - Must be Wrapped in Industry Standard Management
- *Eventually* Networking and Storage WILL Converge on the Same Network
 - Mid-term: *Common* but Independent Networks will be Deployed (*Same H/W*)
- No One Knows when Full Convergence will Take Place, so...
 - **Future-proof with Adapters and Switches Capable of Full Convergence Today!**



Thank You!

Q&A

