


Converged Fabrics



Lloyd Dickman, QLogic

www.openfabrics.org

“You can’t always get what you want”




HPC Adds a Third Fabric

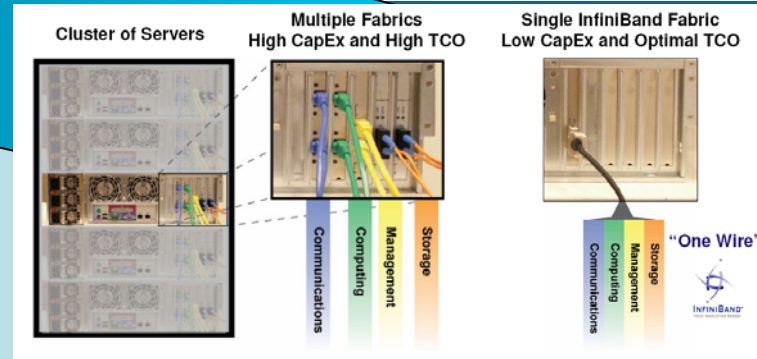
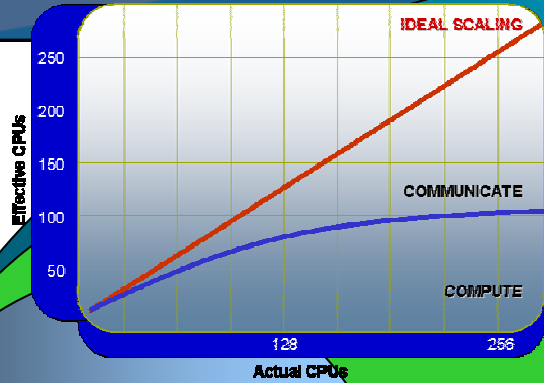
*Network Fabric
(Ethernet)*

**Server
Fabric**

*Storage Fabric
(Fibre Channel)*



Server Interconnect Market Growth



Interconnect market for High Performance Computing

Interconnect market for HPC plus DB Clustering

Interconnect market for HPC, DBC plus Storage Area Networking

Interconnect market for HPC, DBC, SAN plus Enterprise Grid



Compelling Economics for a Converged Fabric



➤ Reduce CapEX

- Deploy one fabric vs 3
 - One adapter type
 - One switching fabric
 - One cabling plant
- Increase network resource utilization
- Increase server resource utilization

➤ Reduce OpEX

- Manage one fabric vs 3
- Flexible resource allocation
- Easier virtualization

Why OpenFabrics for a Unified Fabric?



- Supports multiple Fabrics
 - InfiniBand
 - iWARP

- ULPs for HPC, Sockets and Storage

- Verbs and uDAPL provide APIs for new applications

- Market claims preference for
 - Application portability
 - Multiple vendors with no lock-in
 - Potential for pervasive adoption
 - Open source
 - Wire interoperability



Data Centers Are Challenging



➤ HPC cluster

- Homogenous environment
- Workloads
 - Many are physical simulations
 - Restartable
- Most network traffic driven directly from user
- Rocket scientists who enjoy the journey

➤ Production Data Center


- Heterogeneous environment and legacy
- Workloads
 - Numerous, changing
 - Many are real time, transactional
- Most network traffic driven from OS
- Mere mortals who want to reach the destination
- SLAs

Data Center Market Status ...



- Most current InfiniBand deployments are based on vendor stacks
- Market understands OpenFabrics value proposition
- Technology trials and some deployments are underway

Customer Trials



“But if you try sometimes”



OPENFABRICS
ALLIANCE



OFA needs to do more

- OFED – Moving too fast
 - Rapidly changing sockets implementations
 - Level of interface stability
 - Level of storage maturity
 - Level of robustness

- OFED – Moving too slow
 - QoS capabilities to support SLAs
 - Congestion management
 - Virtualization capable
 - Unified management tools
 - Gateways to storage and communications networks
 - Multiple subnets
 - Resilience
 - Security

Alternatives Exist

➤ iSCSI

➤ FCoE

These alternatives enable transparent deployment

- Do not need a new software stack
- Use existing fabric management

“You just might find you get what you need”



Who is OFA's Customer?

- End-user?
- Distro or OS Vendor?
- System Vendor?

What should OFA provide?

- **Product - OFED**
 - Current HPC direction, but is it working for the Data Center?

- **Technology – Verbs and ULPs**
 - Current direction, and somewhat working
 - Technology should be applicable to Linux, Windows, ..., VMware, Xen, ...

- **Reference Implementation**
 - In place, looking very promising
 - Changes development focus from rushed releases to stable releases with the objective of correctness

Thank You

