



Voltaire and the CERN openlab

**collaborate on Grid technology project using
InfiniBand**

May 27, 2004


Patrick Chevaux
EMEA Business Development
patrickc@voltaire.com

VOLTAIRE



Scaling Out Using Clusters

Super Computers
and Mainframes



Paradigm
Shift



Bunch of interconnected
Linux machines



✓ **Much Lower cost**

✗ **But, lower reliability/MTBF, underutilization,
Higher Complexity, Storage bottleneck**

Our challenge: minimize the Scale-Out overhead

InfiniBand value proposition

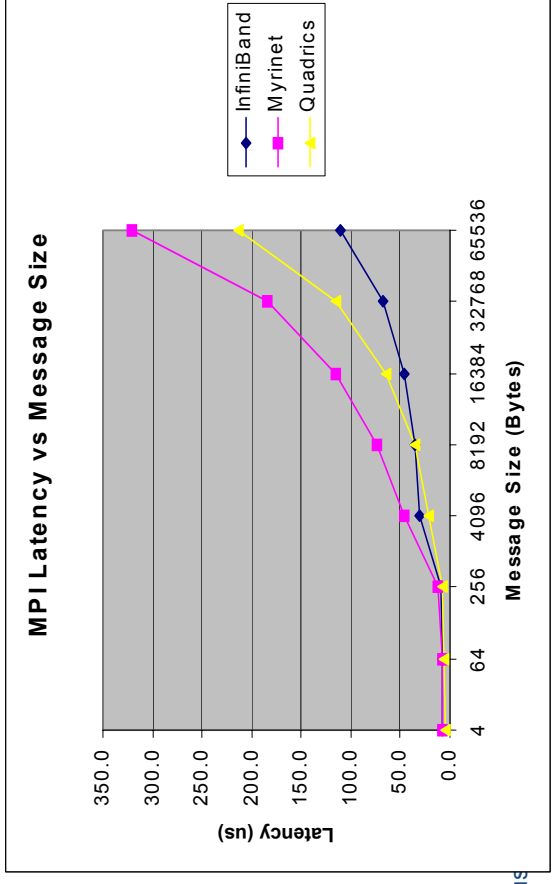

- **Open Interconnect Standard**, designed from ground-up for High-Performance Compute & I/O Clustering
- Significantly lower cost/performance than any other technology due to its architecture
- **Serial + Switched Interconnect**, scales to demand
- **High throughput, Low latency, Low CPU utilization**
 - 850 Mbytes/sec MPI throughput
 - 140ns per switch, 5.8µs MPI end-to-end
- **Enable high-speed file & block I/O, network, and IPC traffic using a single technology with RDMA**
- **Built-in Traffic Classes, QoS, Flow Ctrl, Partitioning, etc..**
- **Already supports 30Gbit/sec links TODAY**
 - New standard for 10, 20, 30, 60, 120Gbit/sec defined

More on InfiniBand

- Intel has been and still is THE main driving force behind InfiniBand
- InfiniBand design was the result of the combined efforts of
 - Compaq, Dell, HP, IBM, Intel, Microsoft, Sun !!!
 - V1 spec issued in 2000
- IBTA, InfiniBand Trade Association, was created to promote the technology and provide the IB « ecosystem », see www.infinibandta.org
- OpenIB www.openib.org has been recently (early 2004) created to promote open source software for InfiniBand

InfiniBand vs. proprietary interconnects


	InfiniBand	Myrinet
Data Rate	10Gbps	2 Gbps
Multiple Protocols	Yes	No
Managed	<p style="text-align: center; background-color: yellow;">Much better performance and multipurpose at lower prices</p>	
Dynamic		
Proprietary		
MPI Latency (uS)	5.9	6.3
Bandwidth (Mbytes/sec)	879 (2500 pci-ex)	248



For discuss

Voltaire: Fast Facts

- **Locations**
 - Business HQ: Boston, USA
 - Sales Reps: Japan, Europe
 - R&D: Herzeliya, Israel
- **Headcount : 70 (May 2004)**
- **Financing**
 - Strategic Investors: Hitachi and Quantum
 - Top US and Israeli VCs
 - Recently raised 15M\$ + cash in bank
- **Partnerships and on-going developments**
 - Hitachi
 - HP, IBM, Sun, Apple, SGI
 - CERN





Voltaire InfiniBand Product Family



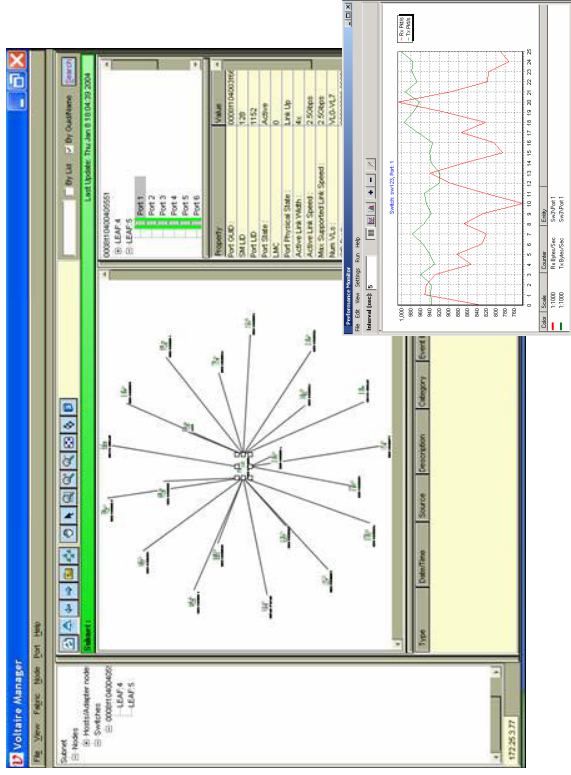
Voltaire InfiniBand Switch Router Family

- 6, 24, 96, 288 port non-blocking, multi-protocol connectivity
- Integrated wire-speed network and storage Virtualization
- No single point of failure – Hot-swappable FRUs
- Non-disruptive software update, fail-over
- Modular elements for investment protection
- Advanced integrated management

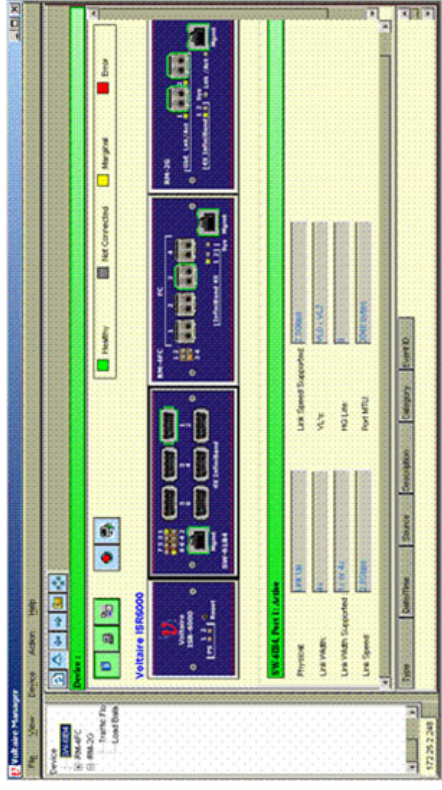


Complete grid infrastructure management


Voltaire Fabric Manager (VFM)



Voltaire Device Manager (VDM)




TCP/IP Provisioning



Load Balancing / NAT, Filtering, VLAN's, QoS

Storage Provisioning



Recent Voltaire Successes in HPC

- **Largest InfiniBand installed base, Most field tested**
- **Voltaire's competitive edge:**
 - Most scalable switch family, Largest IBTA Certified Switches
 - Scalable HPC focused software: Stacks and Fabric Management
 - Highly scalable, Integrated GbE and Storage connectivity
 - Genuine Open Source Strategy



Mississippi State University Installation



Mississippi State
UNIVERSITY

TOP 500[®]
SUPERCOMPUTER SITES

192 Dual Xeon Servers
1.4 TFlops

Voltaire provided:
Switches, Adapters,
Software and Advanced
Fabric Management

“ ... with the power, performance and scalability of Voltaire's InfiniBand switches and the robust functionality of the Voltaire Vision management software, we are well enabled to achieve significant throughput and performance gains...” Trey Breckenridge, HPC Resources and Operations Administrator, Mississippi State University ERC.



Voltaire and the CERN openlab

- **CERN openlab is working on architectures for high performance**
 - Compute clusters
 - I/O subsystems (file I/O, block I/O)
- **Voltaire develops HW+SW solutions**
 - Based on InfiniBand
 - Standard interconnect ideal for HPC
 - Scales well in large compute clusters
 - Low latency, High Throughput, low CPU utilization
 - Atomic and Collective operations
 - ... and ideal for high performance I/O
- **We look forward to a fruitful cooperation**



Thank You !

Have you seen our Web today ?

www.voltaire.com

patrickc@voltaire.com

VOLTAIRE

