

# The Low-Latency InfiniBand Leader

---




## InfiniBand Low-Latency Technical Forum

Jeff Broughton  
Senior Engineering Director  
InfiniBand Products  
Host Solutions Group  
QLogic Corporation

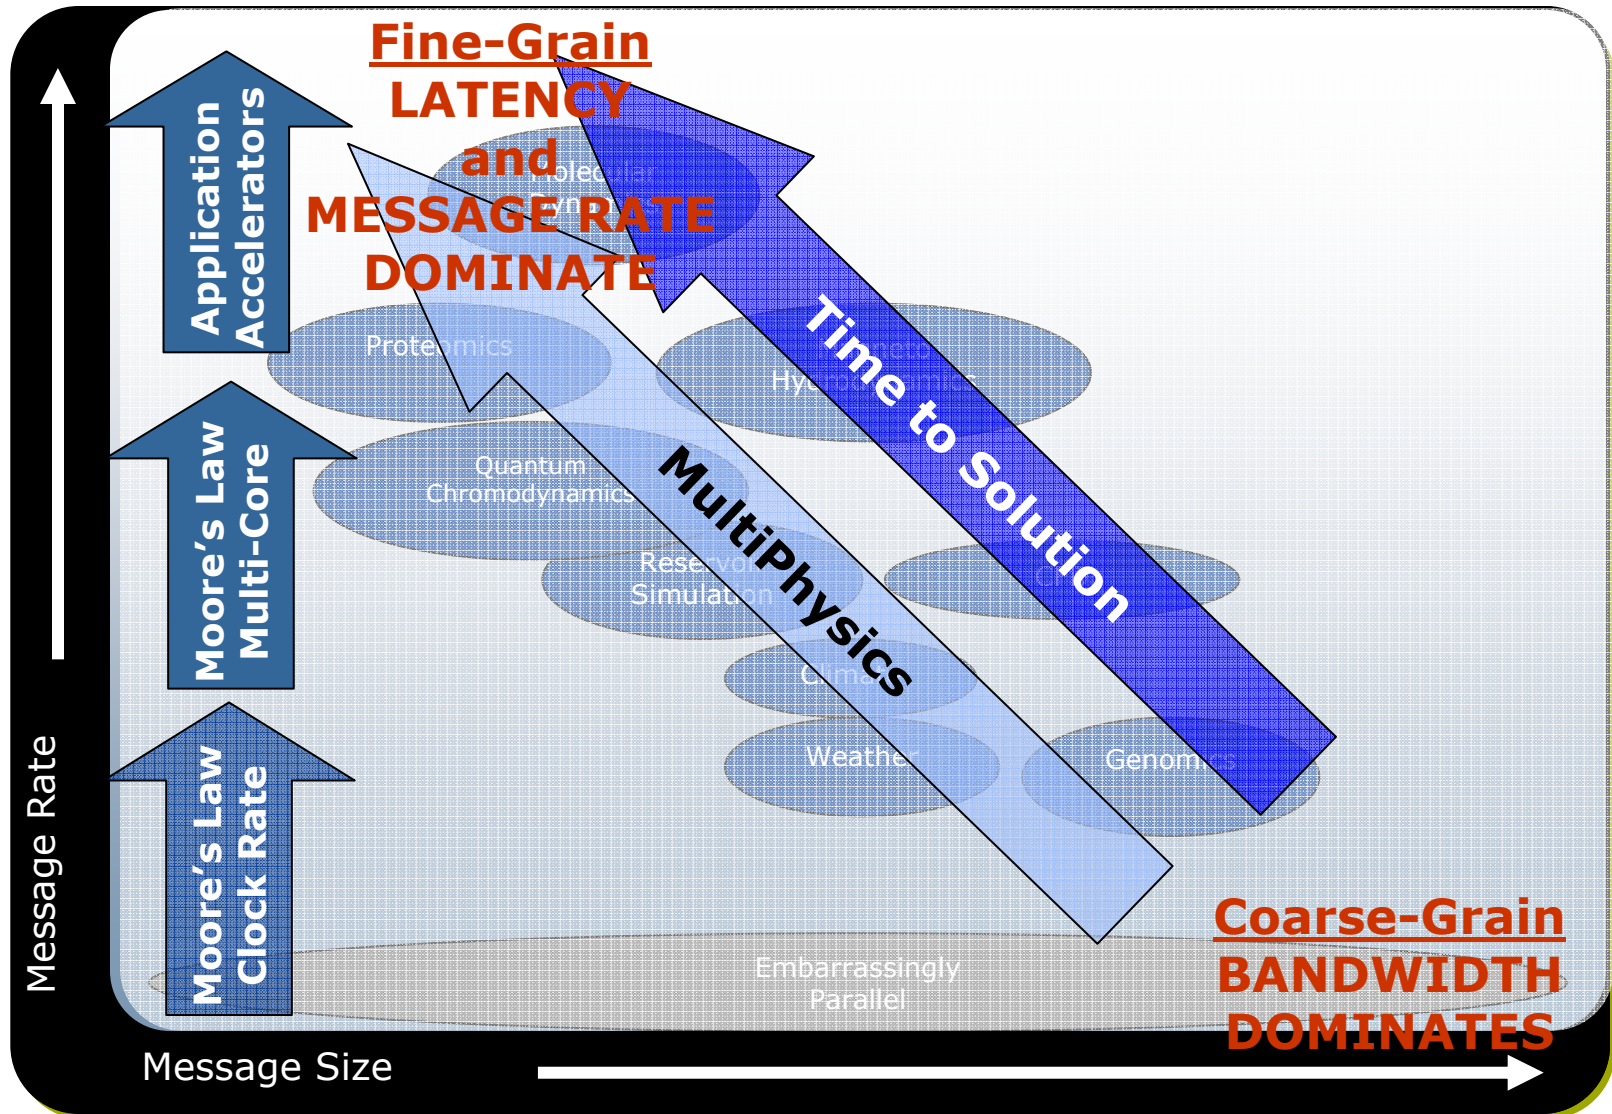
**September 17, 2007**

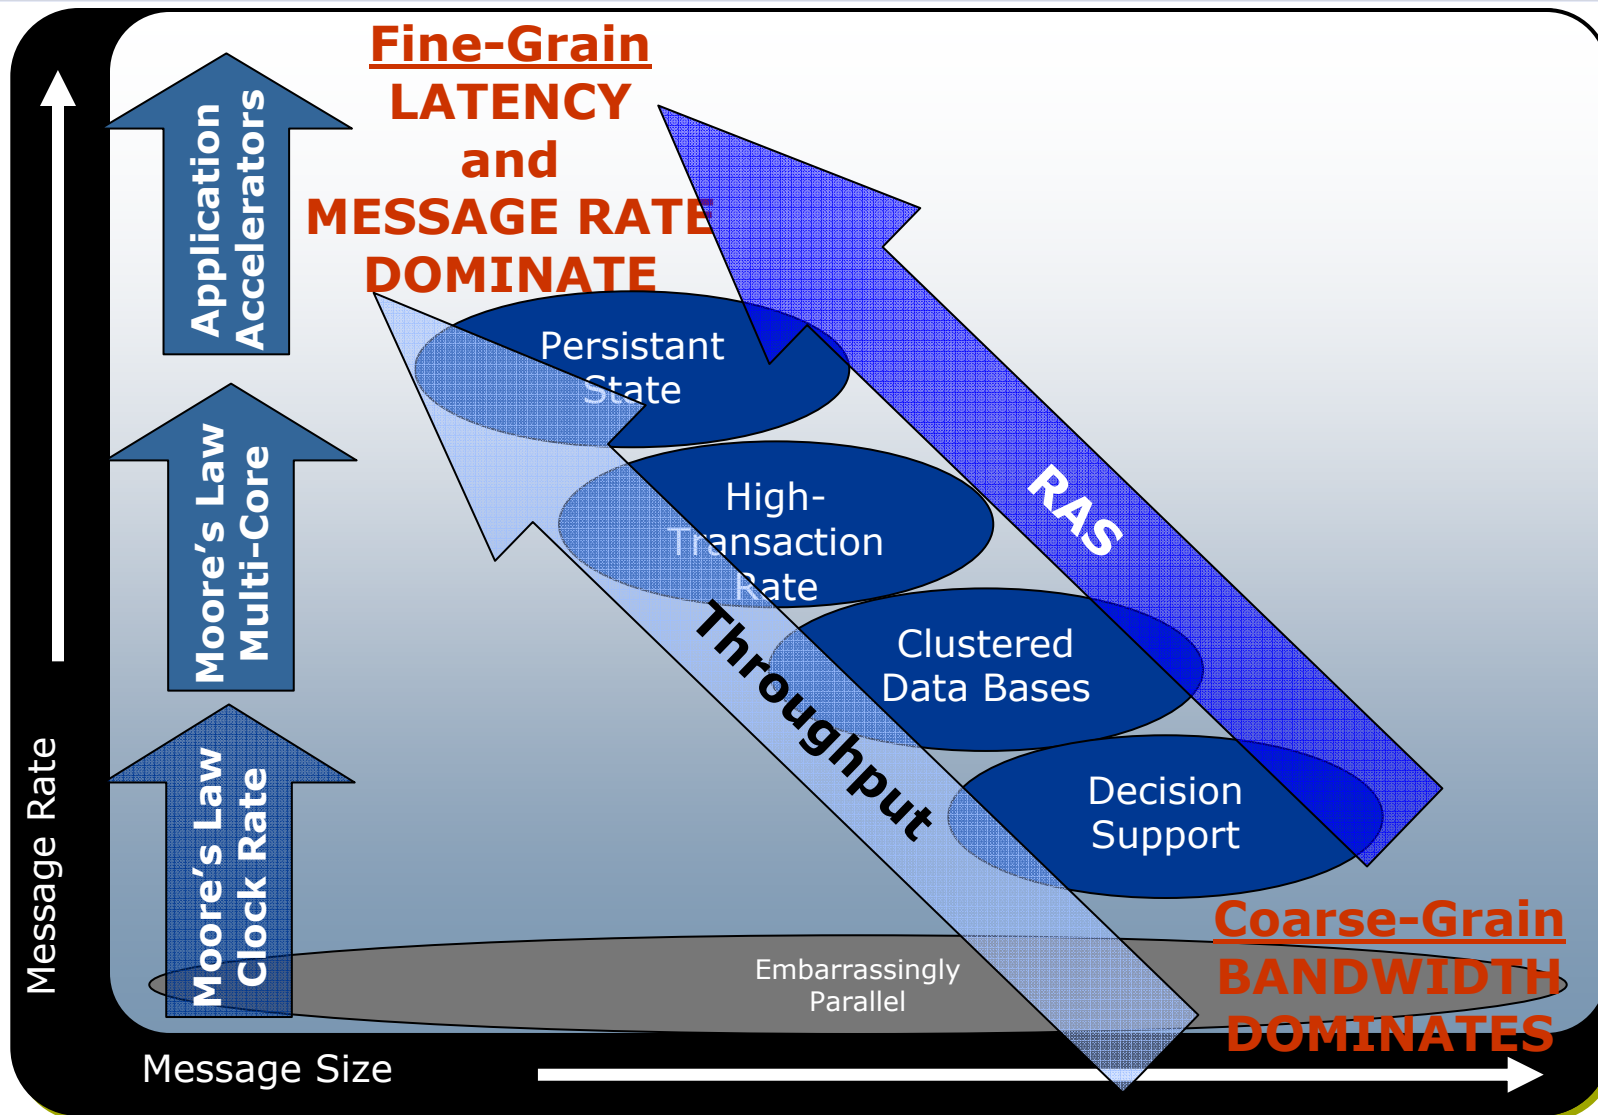
# QLogic Acquired Two Leading InfiniBand Companies in 2006



- Great Latency & Throughput
- Scalable
- Reliable & Easy to use

# For HPC, Application Messaging Profiles Matter!

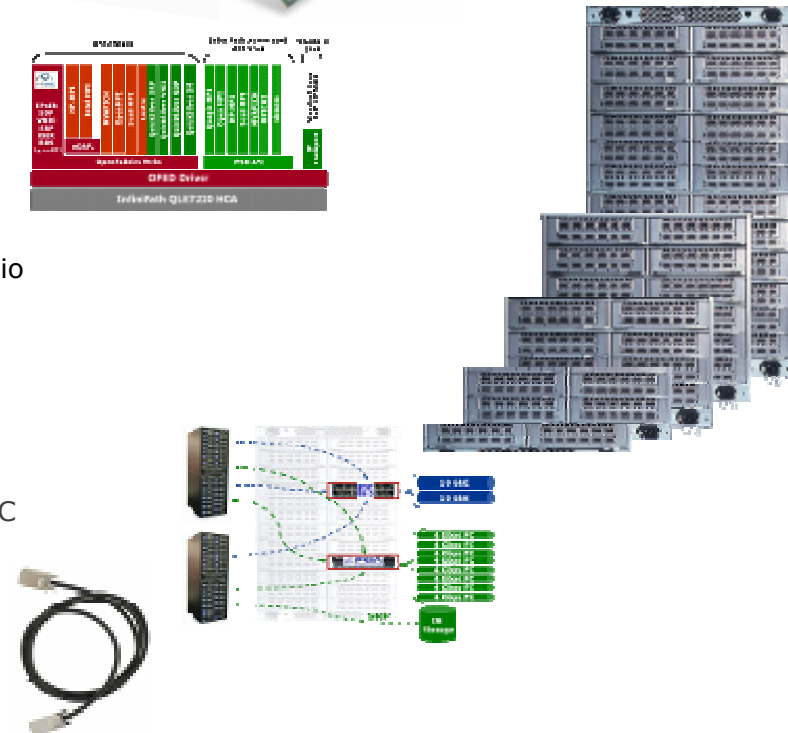





# Full InfiniBand Product Portfolio for Low-Latency Solutions



- **Widest range of host InfiniBand adapters**
  - Highest performance, most scalable
  - SDR and DDR capable
- **OFED+ software stacks**
  - OpenFabrics / OFED
  - Accelerated standard MPI stacks
- **High Capacity Multi-protocol Director**
  - Scalable InfiniBand Core Switches
  - 11.52Tb/s throughput in a single chassis
  - InfiniBand DDR supported across full fabric switching portfolio
  - Supports 20Gb/s node-to-node, 60Gb/s switch to switch
- **InfiniBand Edge Switches**
  - SDR and DDR capable
- **Multi-protocol gateways**
  - Virtual I/O Controller (VIC) Functionality – 10 GbE and 4G FC
  - Scale I/O and servers independently
  - Pool and share I/O between servers
- **Cables**
  - SDR and DDR
- **Comprehensive InfiniBand Fabric Management**
  - Centralized configuration and fabric initialization
  - Accelerated fabric deployment and verification
  - Fabric health and performance monitoring
  - Fabric-wide diagnostics and maintenance
  - Centralized management of virtual fabric services





- **Lowest latency in the industry**
  - Lowest point to point – MPI 1.3 us (with switch)
  - Lowest scalable latency of any cluster interconnect: 2 us at 512 cores
- **Highest bandwidth in the industry**
  - Full line rate InfiniBand SDR and DDR
- **Highest packet rate in the industry: 20M packets per sec**
- **Strongest RAS in the industry**
  - “Mem-Free”
  - No firmware
  - Simple to deploy and run
  - Strong data integrity and recovery
  - Strong diagnostic tool set
- **Scalable MPI**
- **Persistent state applications**
- **Full DDR capability**
- **Headroom to maintain performance under demanding transaction loads**
- **Support large cluster uptime**
- **Mission critical, real time processing**




# InfiniBand Switch Products: Low Latency and Outstanding RAS



- 150~450 ns latency
- > 1 TB/s aggregate bandwidth
  - Single and Dual data rate
- Fabrics scalable to thousands of ports
- RAS
  - Modular Director Switches
  - Common hot pluggable modules for all chassis

**24 Ports**  
1U  
**Non-modular**





# Low-Latency, High IOPS Connectivity to ALL Storage




- >40K IOPS @ 512 bytes (target limited)
- >500 MB/s @ 256 bytes
- >700 MB/s @ 1 MB/s

## Direct Connect InfiniBand Storage



**SRP**




# Intra-Cluster Sockets Communications

<8 us Sockets latency  
>500K UDP packets/s @ 256 bytes



**VNIC**



**10 GbE Gateway**

**IEEE 802**

**IPoIB**

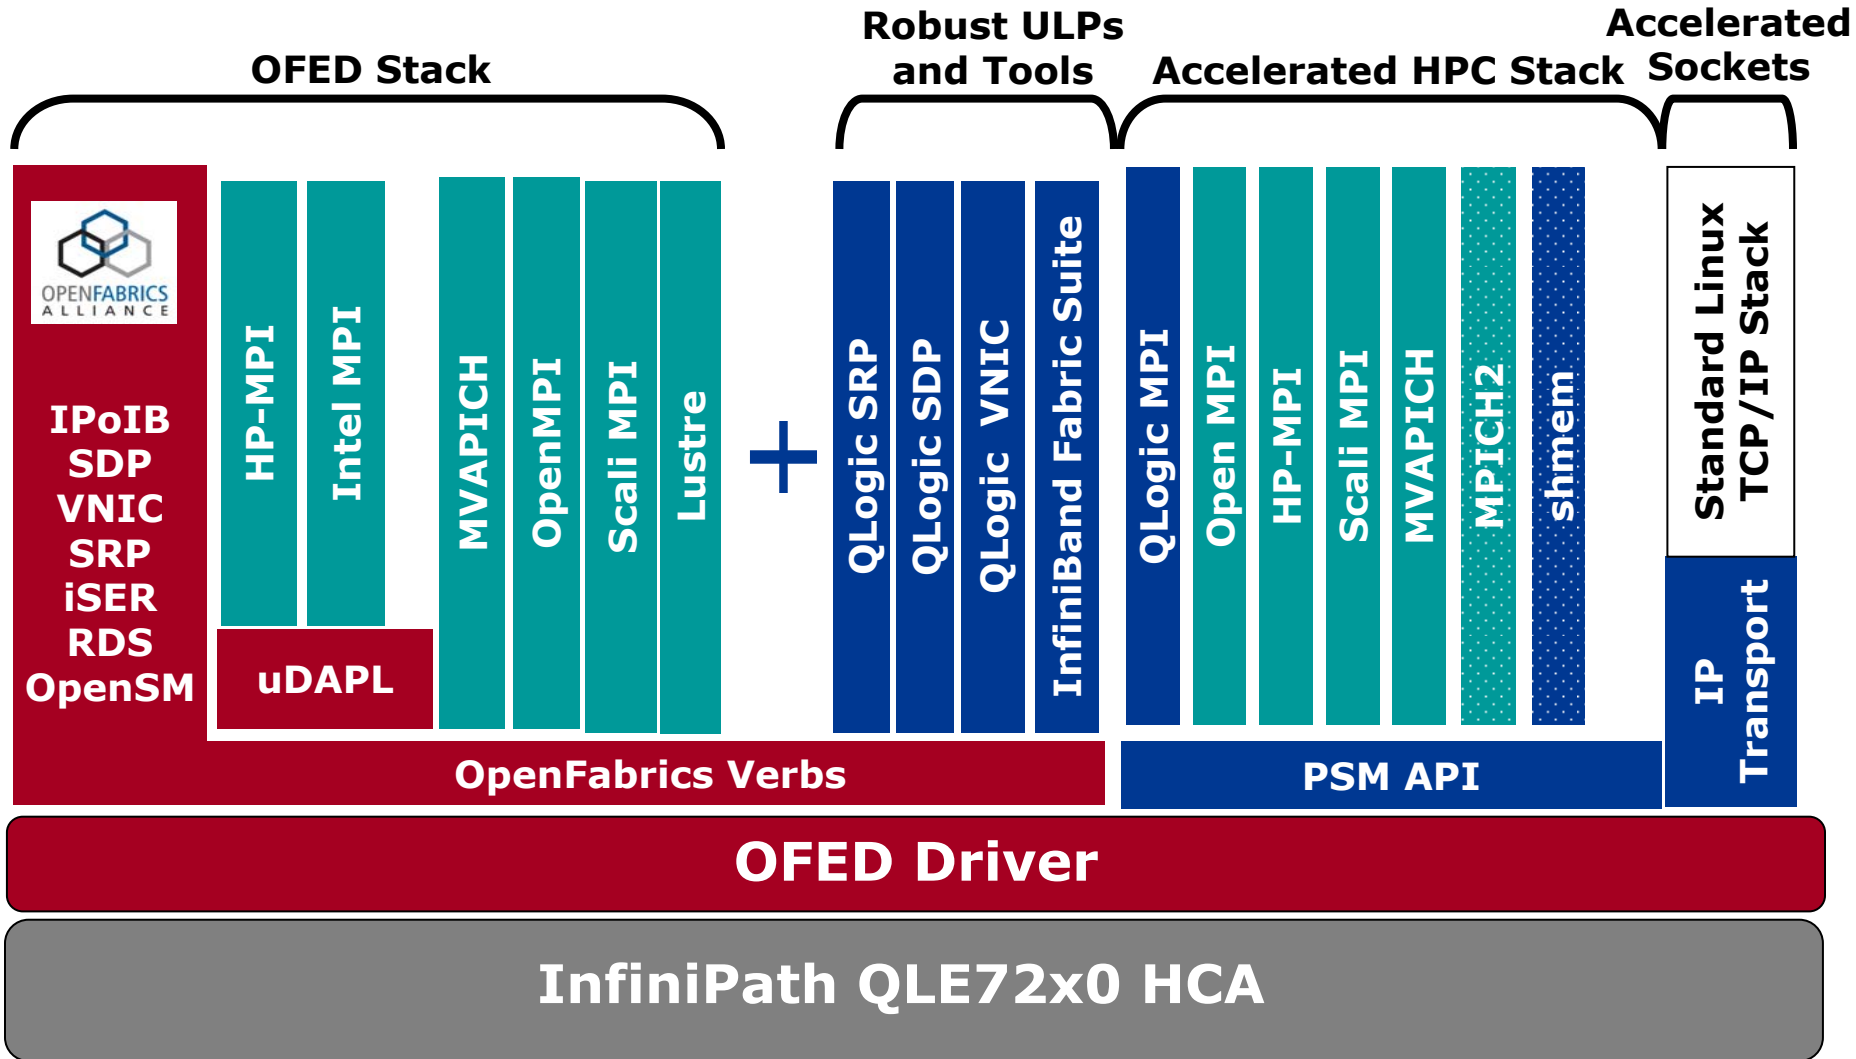


**SDP**

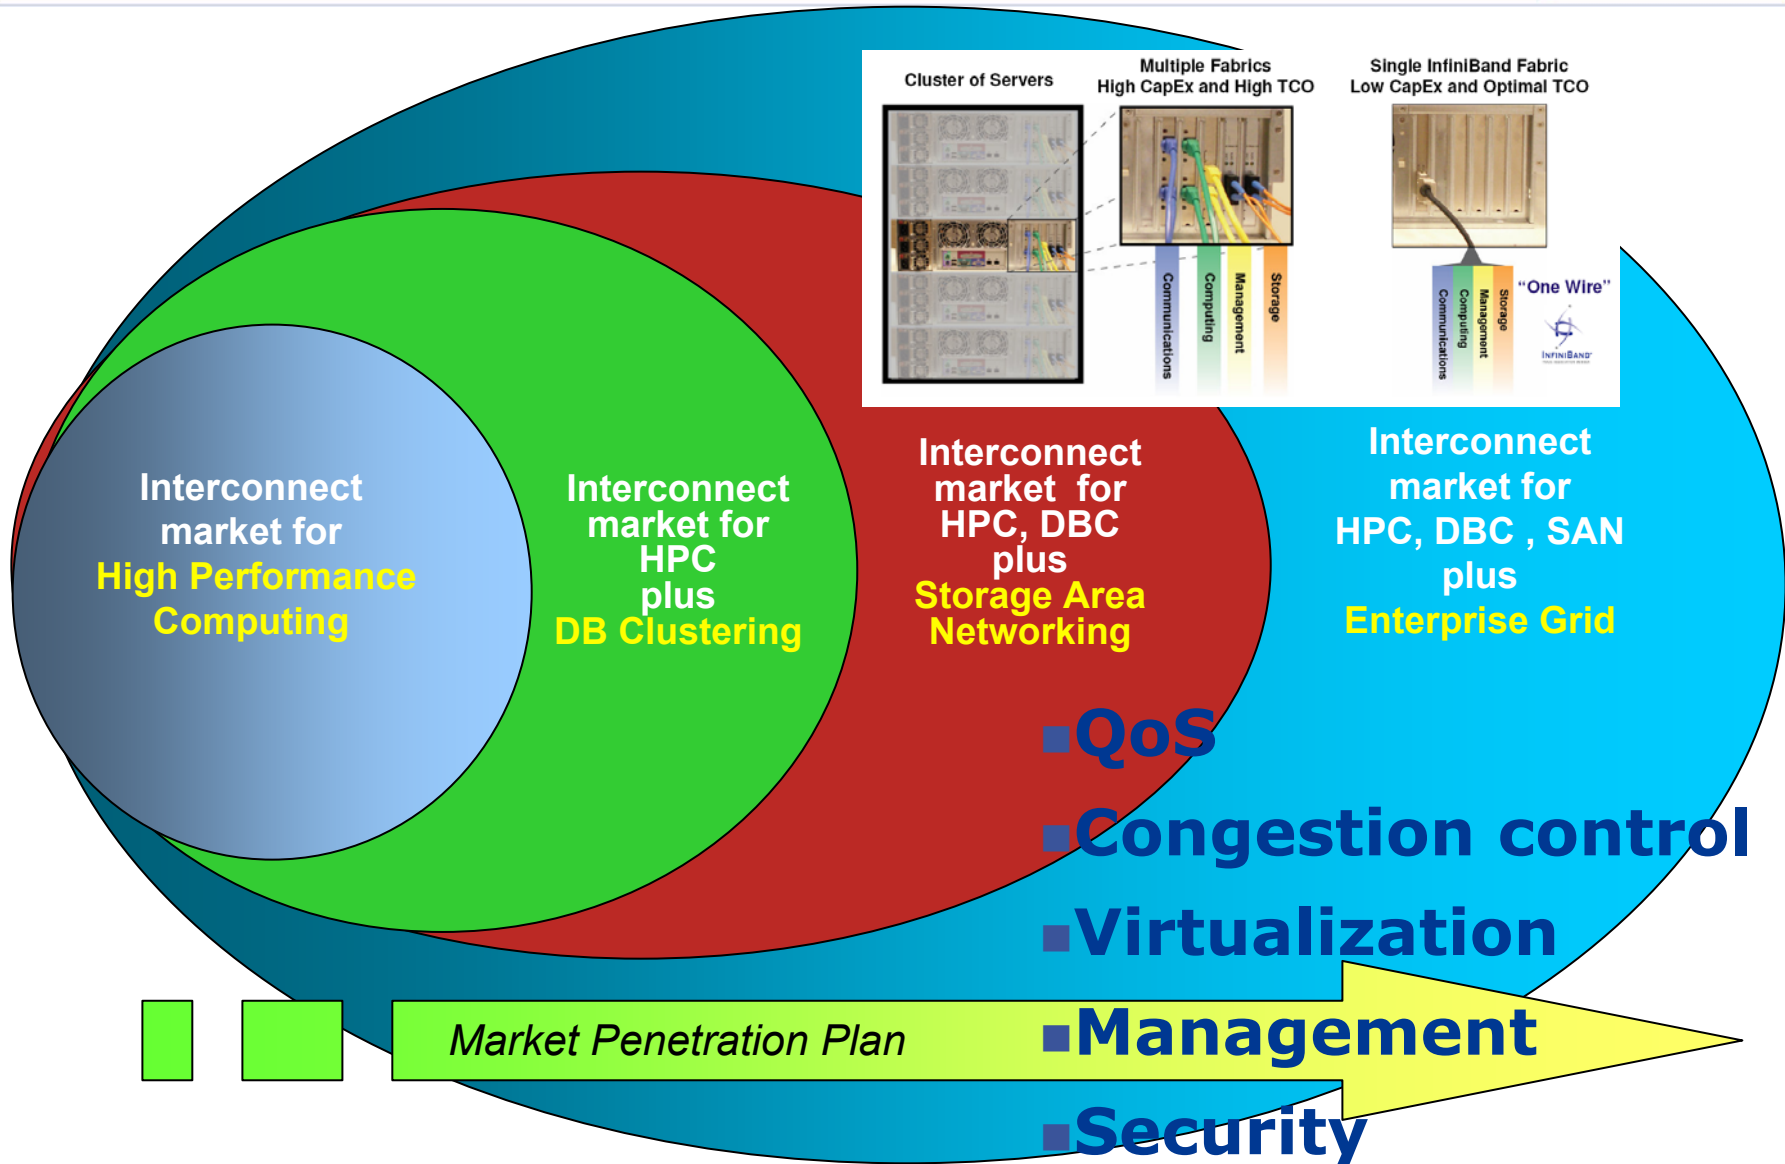


**EtherNet Emulation**






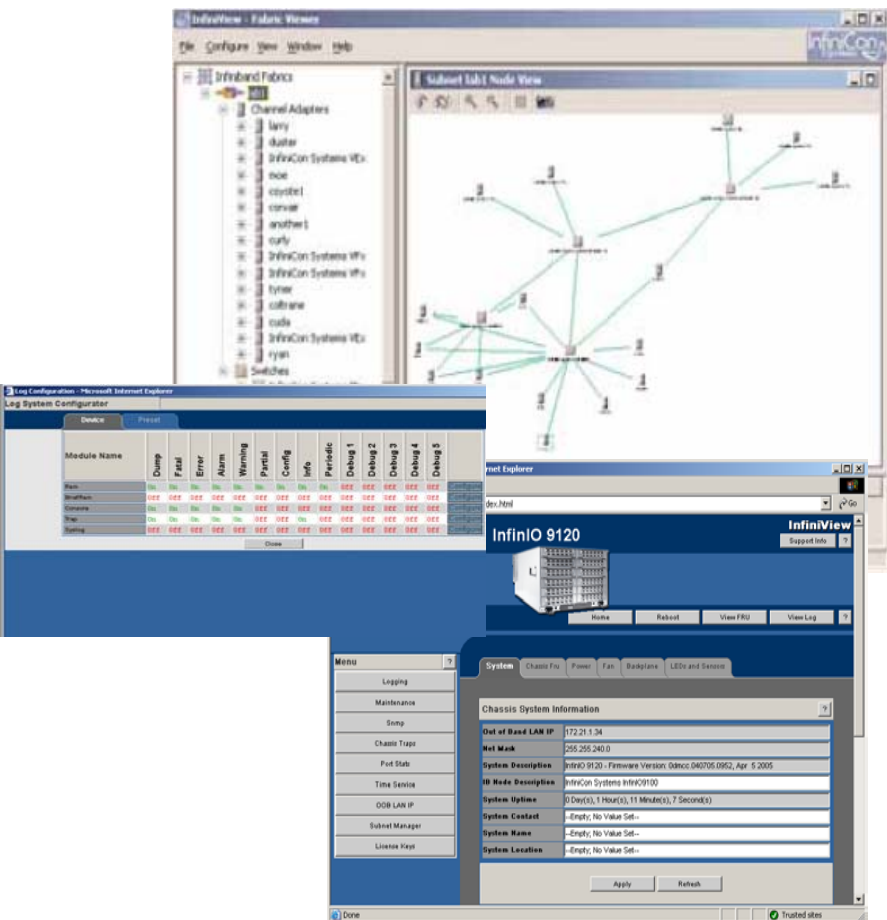
# Unified / Converged Fabric Places Increasing Demands on the Fabric



# Actively Working with InfiniBand Cabling Suppliers to Solve Customer Cable Management Needs for DDR



Excellent BER → Cluster RAS at scale



## Lowering InfiniBand OpEx

### Fabric Management

- No user intervention required
- Scalable to thousands of nodes
- Fast subnet manager
- Configurable QoS
- Optional redundancy


### Chassis/Element Management

- No user intervention required
- Hot swap FRU(s)
- Optional redundancy
- Common feature set, look and feel across all chassis/switch products
- Extensive trap generation and logging

### IO Management

- Flexible IO configurations
- Hot swap IO
- Virtual / Scalable IO
- Fabric Bootable Servers

Thank You



**QLOGIC<sup>®</sup>**