

DataDirect[™]
N E T W O R K S

Reducing Service Latency in Large Storage Systems

Dave Fellingner, CTO
September 28, 2010



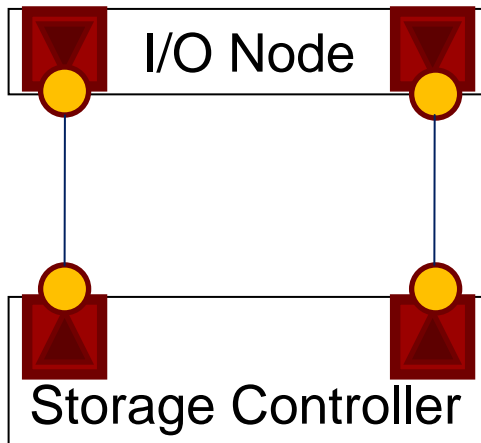
Bottleneck analysis

- Hardware chain
 - Drive servo operation
 - Multiple SCSI layers
 - Multiple bus transitions
 - Memory bandwidth limitations
- Software chain
 - Memory copies
 - Kernel operations
 - Layers of consecutive operations including Vnodes, Inodes, and FAT

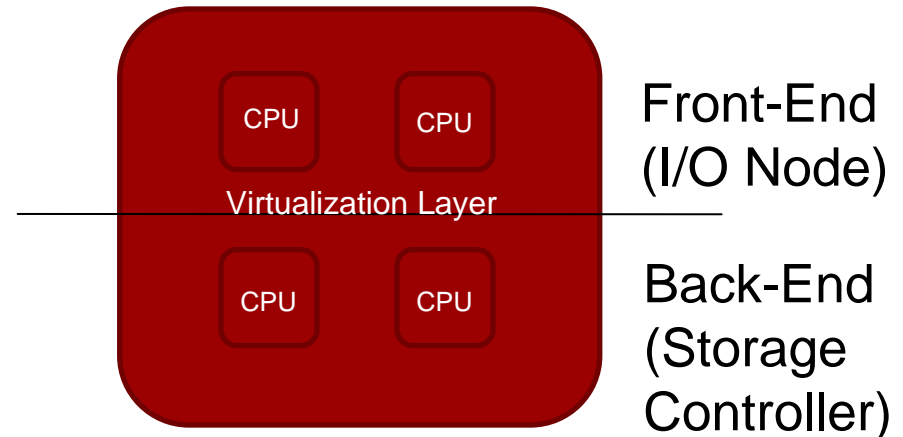
Hardware Simplification




System Simplification/Cost Reduction

Today



Simplified I/O Path



-  HBA/HCA
-  SFP/XENPAK/XFP
-  Cable

- All I/O Node to Storage Controller interconnect components ***eliminated***
- Protocol overhead and conversions ***reduced***
- Shared memory space
- High-speed internal connections

Software Simplification

QDR IB/10GbE

QDR IB/10GbE

Applications,
File Systems
Database, etc.

Applications,
File Systems
Database, etc.

Failover

Native PCI-e Drivers

Native PCI-e Drivers

MMAP'd
Hi-Speed
Direct
Disk I/O

MMAP'd
Hi-Speed
Direct
Disk I/O

High Speed I/O
Virtualization Hypervisor

High Speed I/O
Virtualization Hypervisor

DDN RAID Stack

DDN RAID Stack

Real-Time Storage OS

Real-Time Storage OS

Cache
Coherency &
Mirroring High-
Speed
Interconnect

SFA Controller

SFA Controller

System Simplification

- Use virtualization technology to eliminate layers
 - File system services
 - Metadata services
 - Data analysis
- Use Object based file system technology to reduce software layers
 - Eliminate the need for FAT
 - Eliminate the possibility of fragmentation
 - Force one disk seek for small to medium objects
 - Enable external metadata search engines

Conclusion

- Simplifications must be employed at every level of the system
 - Local networks can be simplified and can employ RDMA technologies
 - Concentration should be place on the end goal of file reads per second rather than system IOPs
 - All legacy protocols should be reviewed and optimized for the current task requirements
 - Analysis tools should be as close to the storage as possible in terms of processor and network layers

DataDirect[™]
N E T W O R K S

Thank You

