



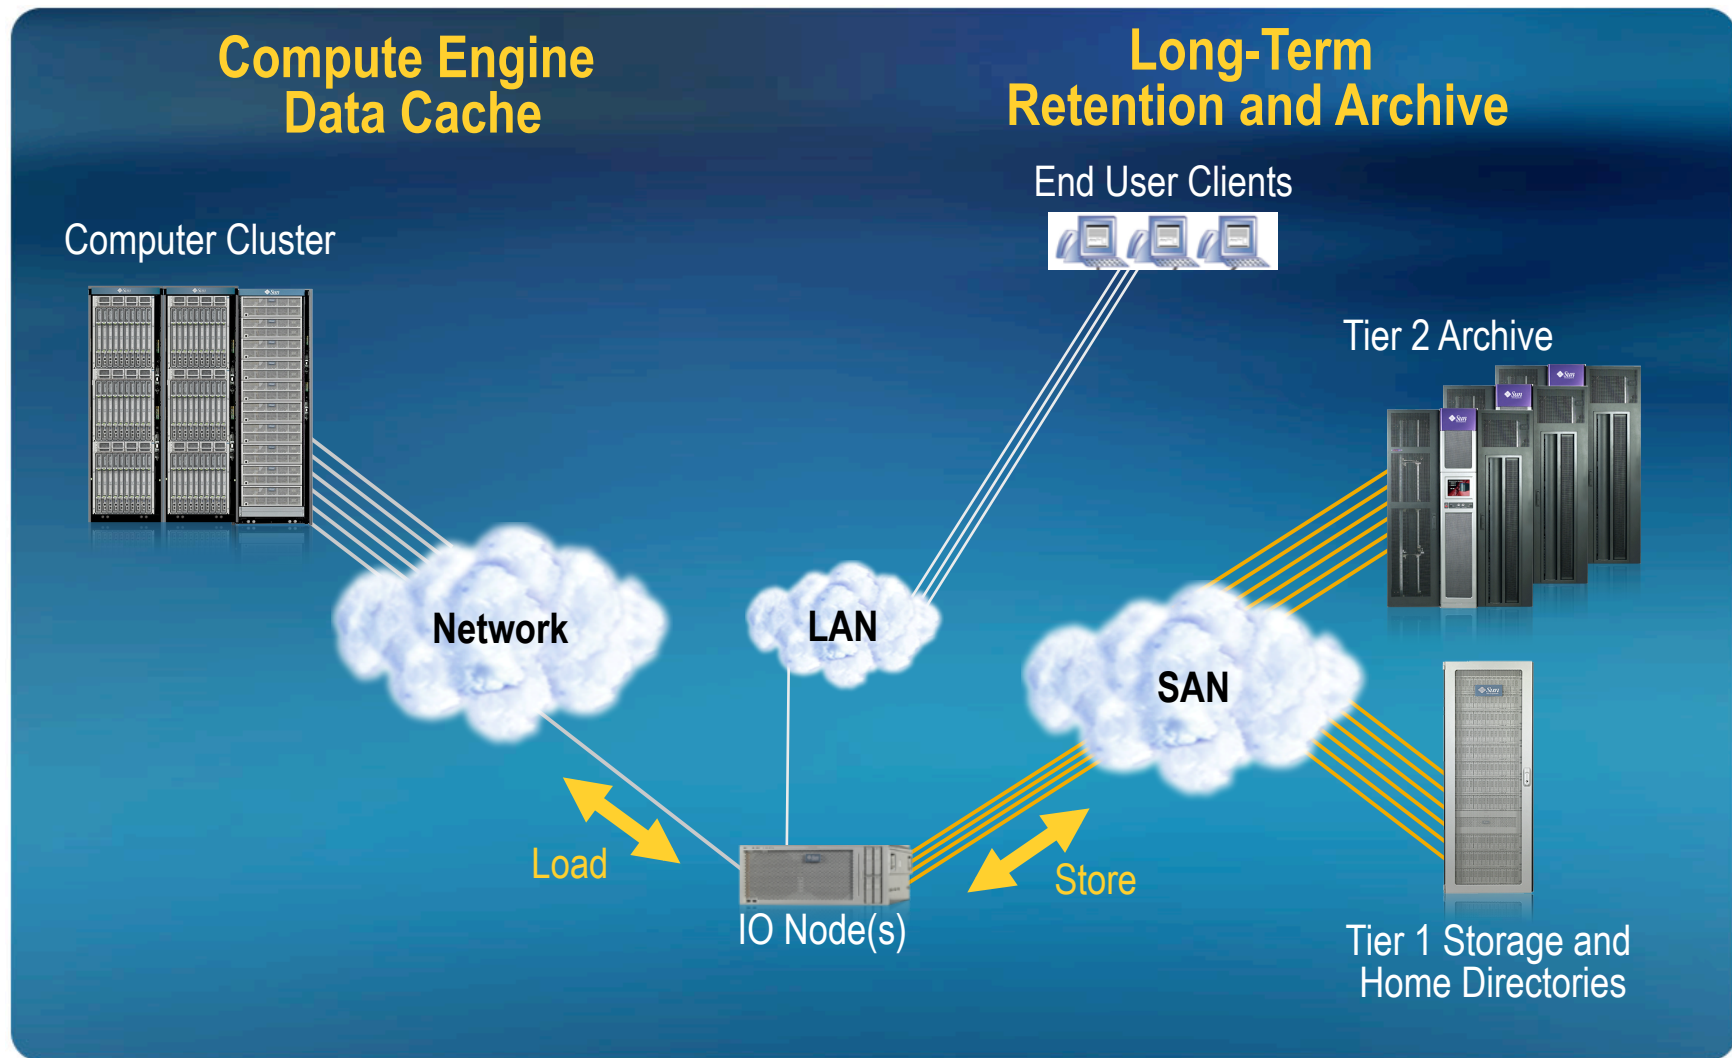
Managing Mountains of Data in Large Scale HPC Systems

Roger Goff
HPC Architect
Sun Microsystems



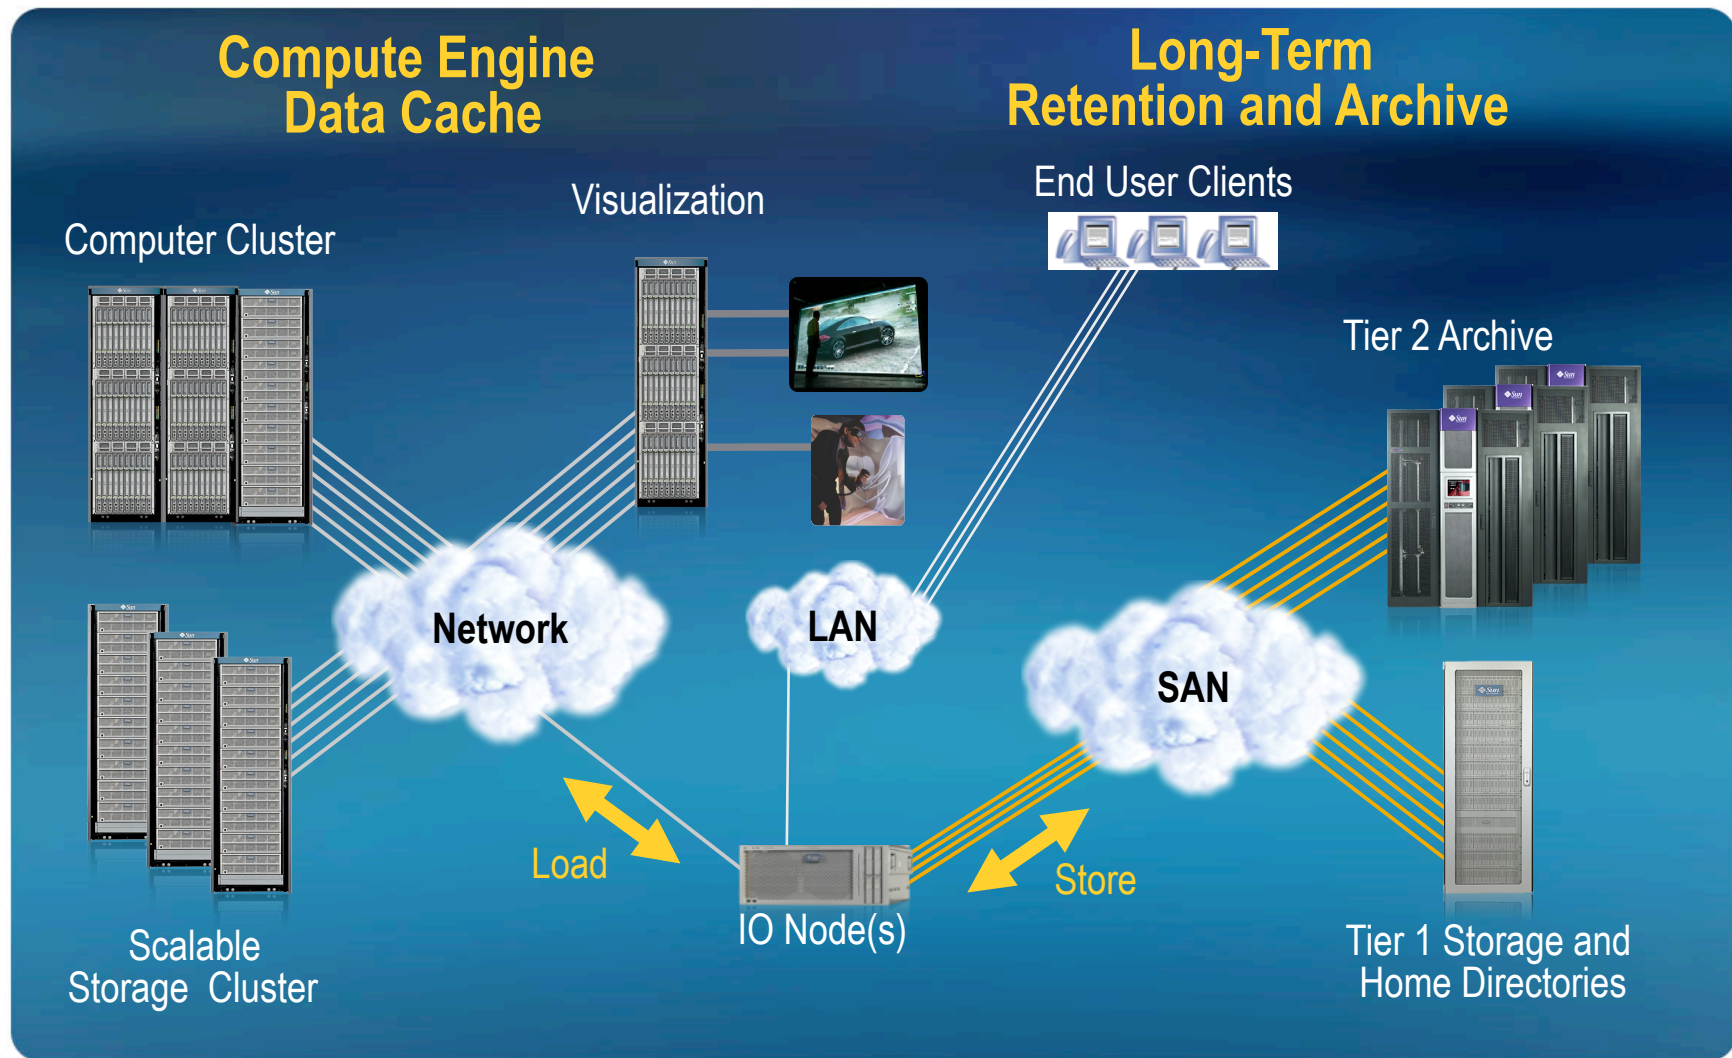
HPC Storage Ecosystem

NFS + SAN = Small to medium systems or small data



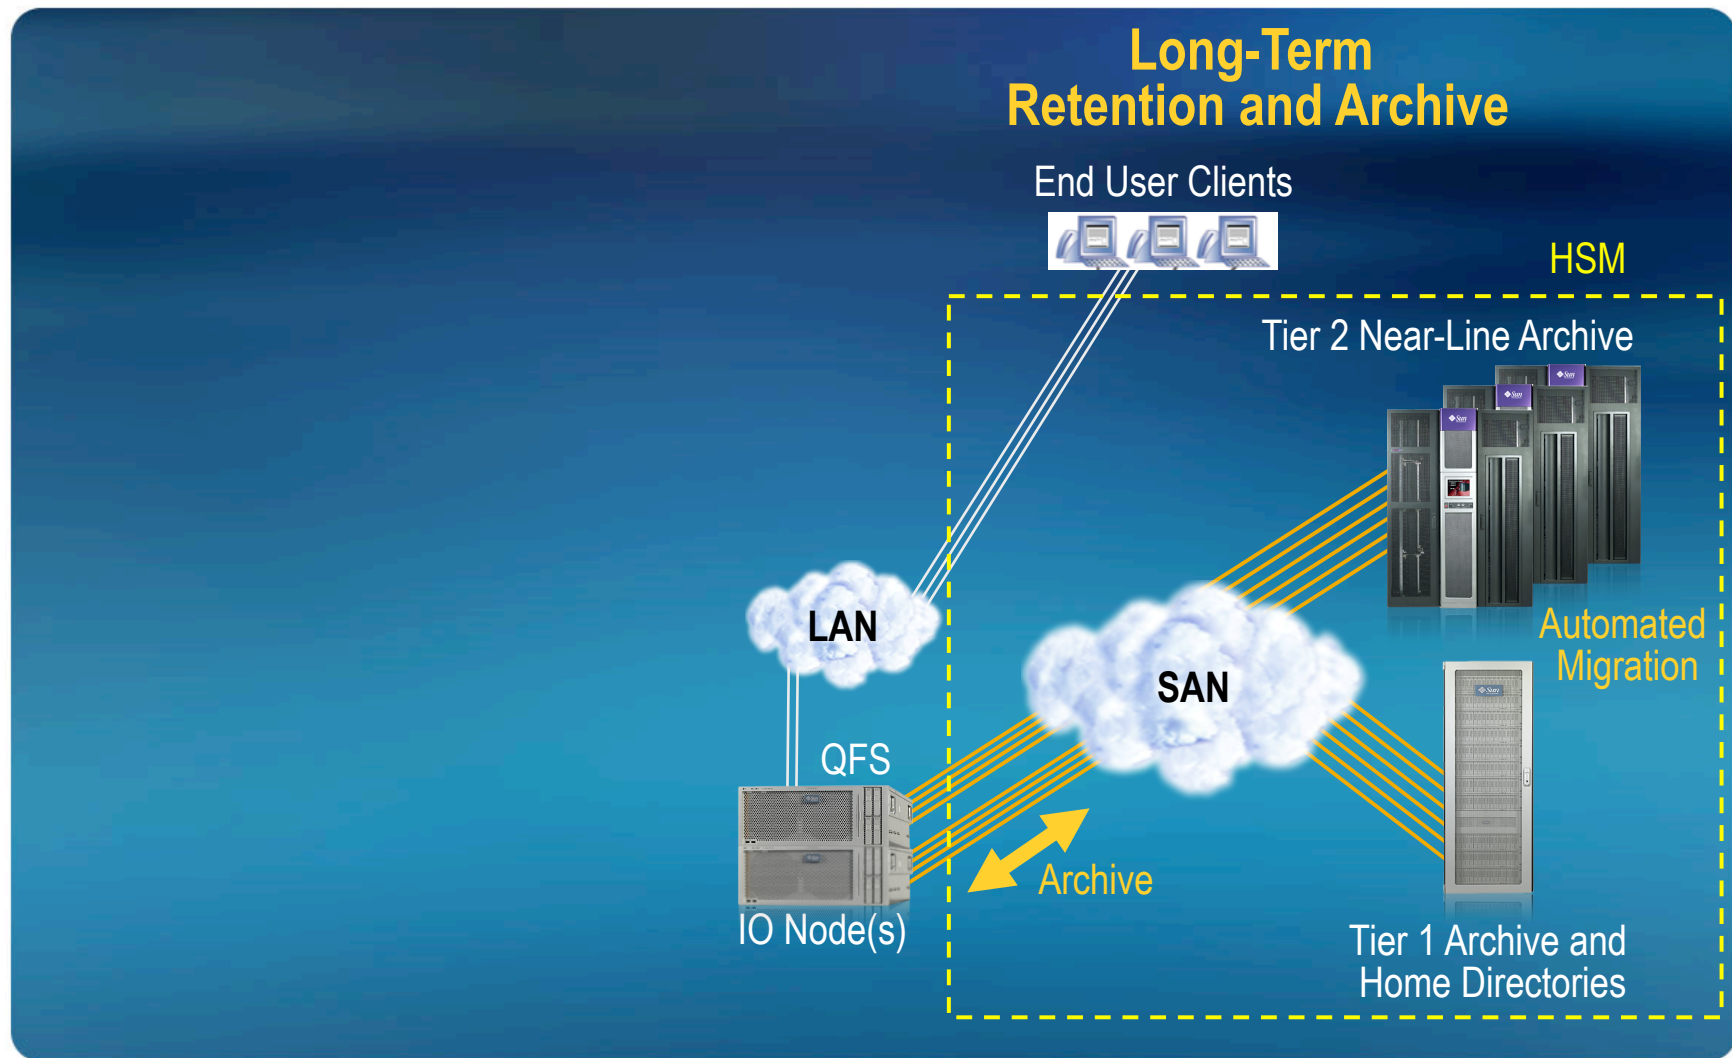
HPC Storage Ecosystem

Lustre + NFS + SAN = Large systems or large data



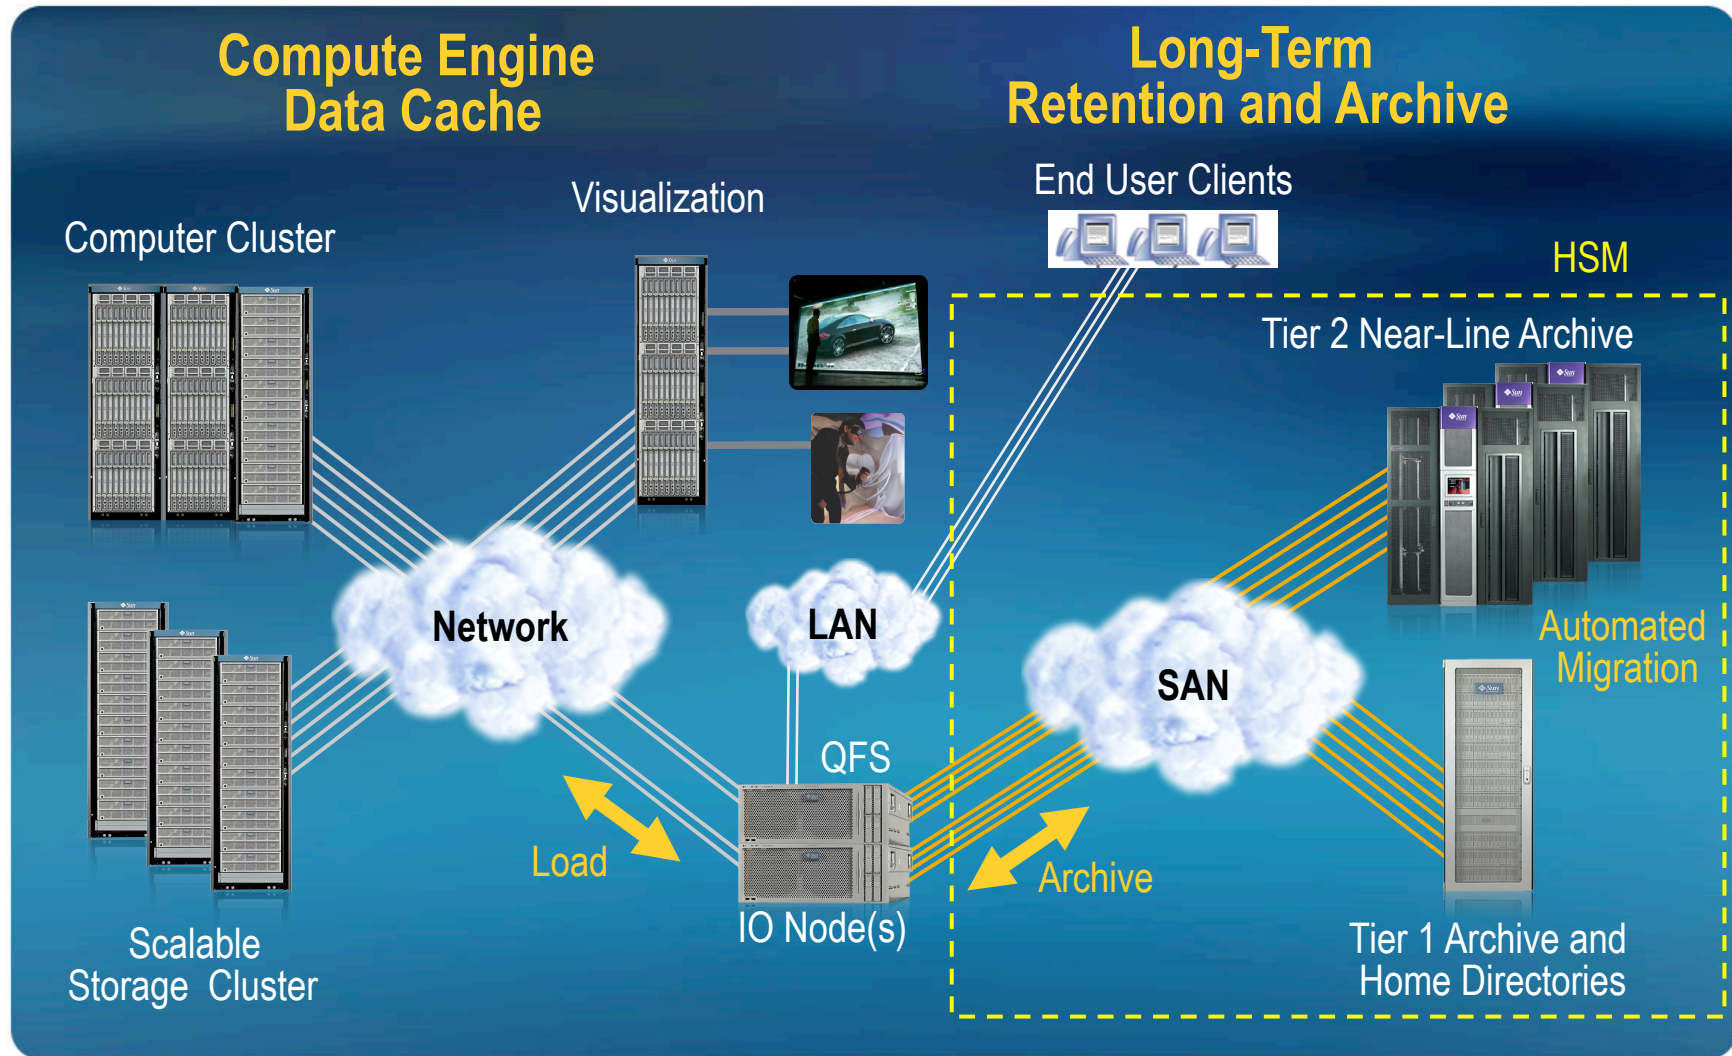
Hierarchical Storage Management (HSM)

Ex. Sun Storage Archive Manager (SAM) + QFS



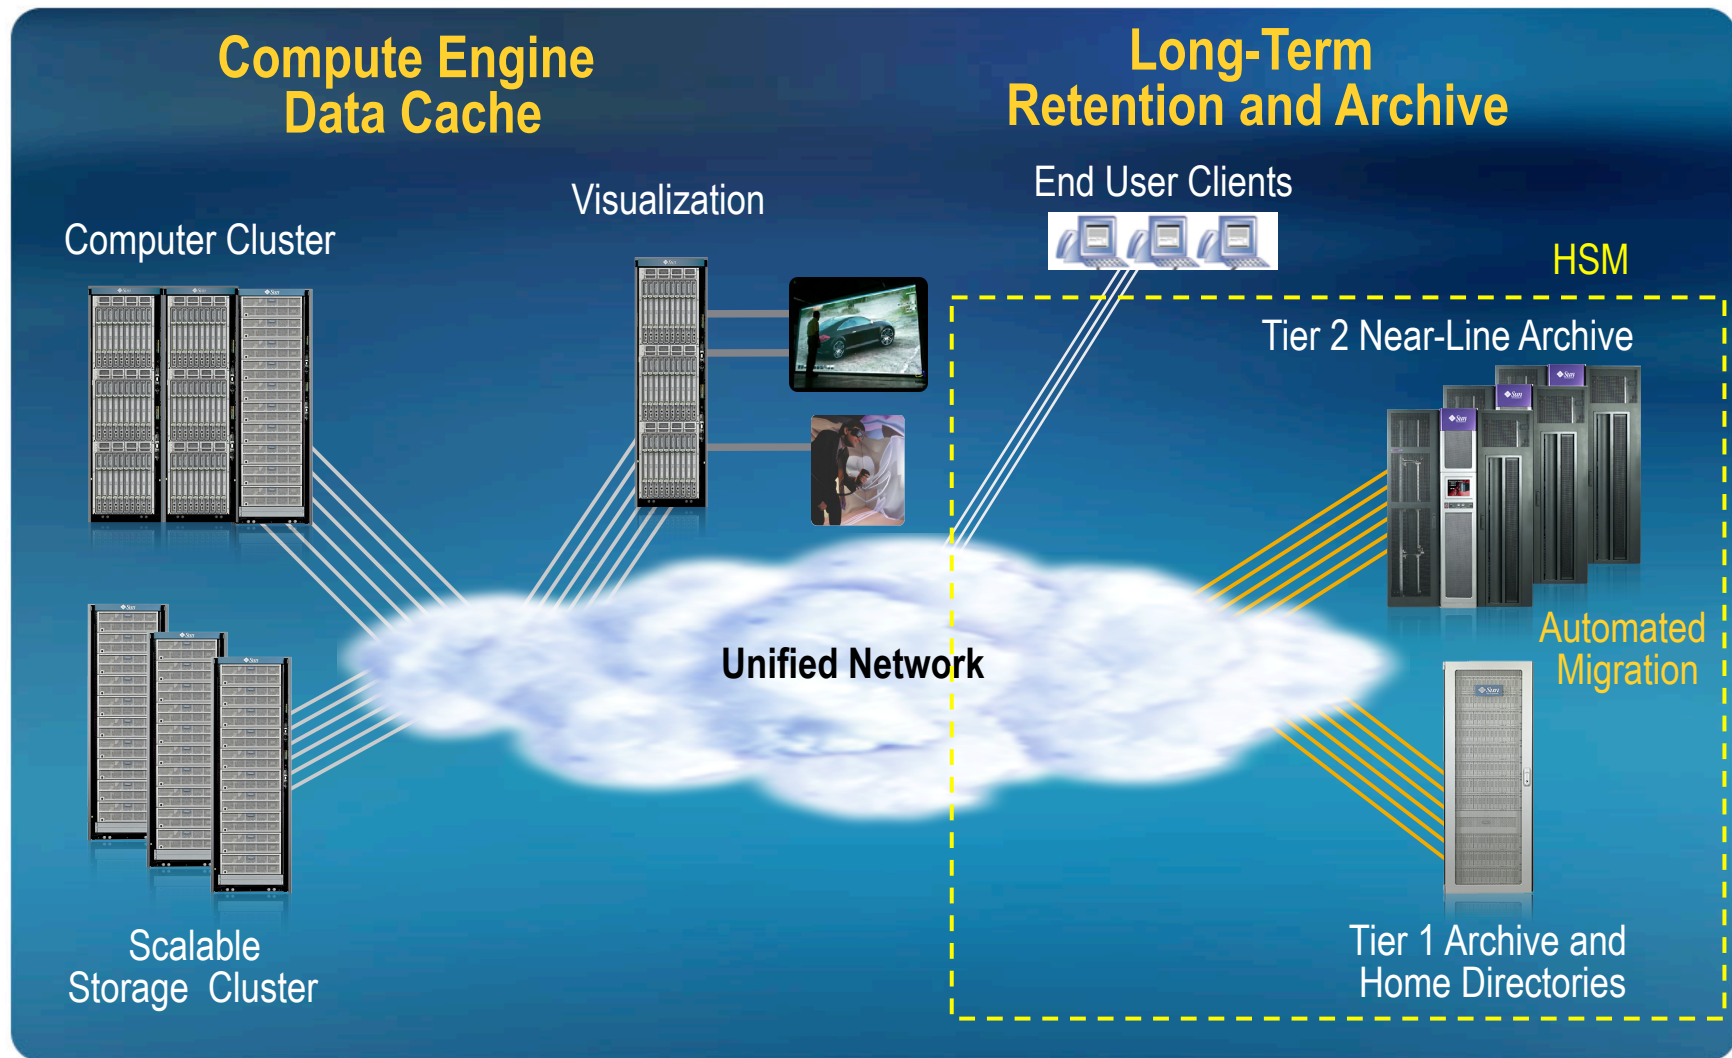
HPC Storage Ecosystem

Lustre + NFS + HSM = Large systems or large data with archiving



HPC Storage Ecosystem – In the Future

Lustre + HSM = Large systems or large data with archiving





The Changing Storage Landscape

Where to Store Data?

Optimization Trade-Off



Cost Effective Performance

SSDs are 70X more cost effective



- Enterprise HDD
 - > 180 Write IOPS
 - > 320 Read IOPS
 - > 300 GB
 - > ~18W
- \$ per IOPS: 2.43



- Enterprise SSD
 - > 7,000 Write IOPS
 - > 35,000 Read IOPS
 - > 32GB
 - > ~3W
- \$ per IOPS: 0.04

ZFS Hybrid Pool Example



- 4 Xeon 7350 Processors (16 cores)
- 32GB FB DDR2 ECC DRAM
- OpenSolaris with ZFS

Configuration A:



(7) 146GB 10,000 RPM SAS Drives

Configuration B:



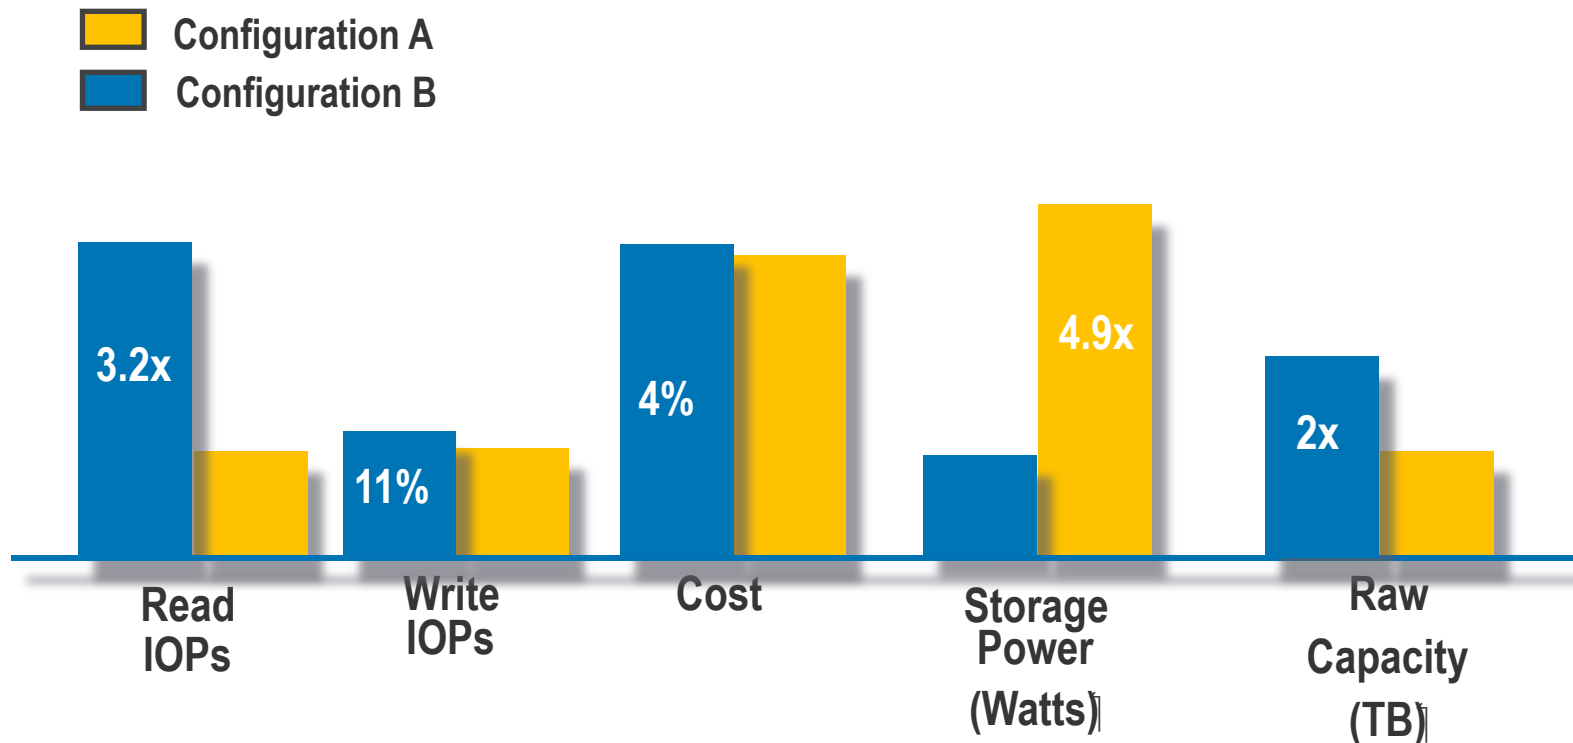
(1) 32G SSD ZIL Device

(1) 80G SSD Cache Device



(5) 400GB 4200 RPM SATA Drives

ZFS Hybrid Pool Example



- If NVRAM were used, Hybrid wins on cost, too
- For large configs (e.g. 48T–750T+) cost is entirely amortized

References

- Implementing the Lustre File System with Sun Storage
 - > <http://wikis.sun.com/display/BluePrints/Implementing+the+Lustre+File+System+with+Sun+Storage>
- Sun Storage and Archive Solution for HPC
 - > <http://wikis.sun.com/display/BluePrints/Sun+Storage+and+Archive+Solution+for+HPC>

Lustre Architecture

