



2001 HPSS User Forum
June 5-7, 2001
San Diego, California

Sun Microsystems, Inc.



Preface

HPC SAN AUGMENTING HPSS

AGENDA

HPC NEEDS

HPC STORAGE CHALLENGES

HPC SAN SOLUTION

HPC SAN HARDWARE AND SOFTWARE STACKS

HPC PRODUCT COMPONENTS

VALUE OF HPC SAN

SUMMARY AND DISCUSSION

HPC NEEDS

Massive scalability

Continuous uptime -- 24/7

Shrinking Backup Window

Very Fast Restores Expected

Centralized data/resource management

Terabytes/Petabytes and beyond data requirements

Heterogeneous data sharing beyond NFS

Thousands to millions, unequal capability, single file access rates needed into the GB/sec range





HPC NEEDS Continued

HPC Customers Also Need:

COTS costs and vendor viability

Minimum integration challenges except for their mission-unique areas

High single file to single process I/O rates

Maximum file system efficiency under load

Easy evolution to new computing technology



HPC

Storage Challenges

High data rates and large file sizes cause extraordinary data-handling requirements:

Mission Critical timelines require low-latency access to most important data first

High bandwidth, low latency access to shared data

Some files are VERY large!

There may be millions and millions of smaller files

All must be tuned and serviced concurrently



HPC SAN Solution

Connect multiple, potentially heterogeneous, servers to a

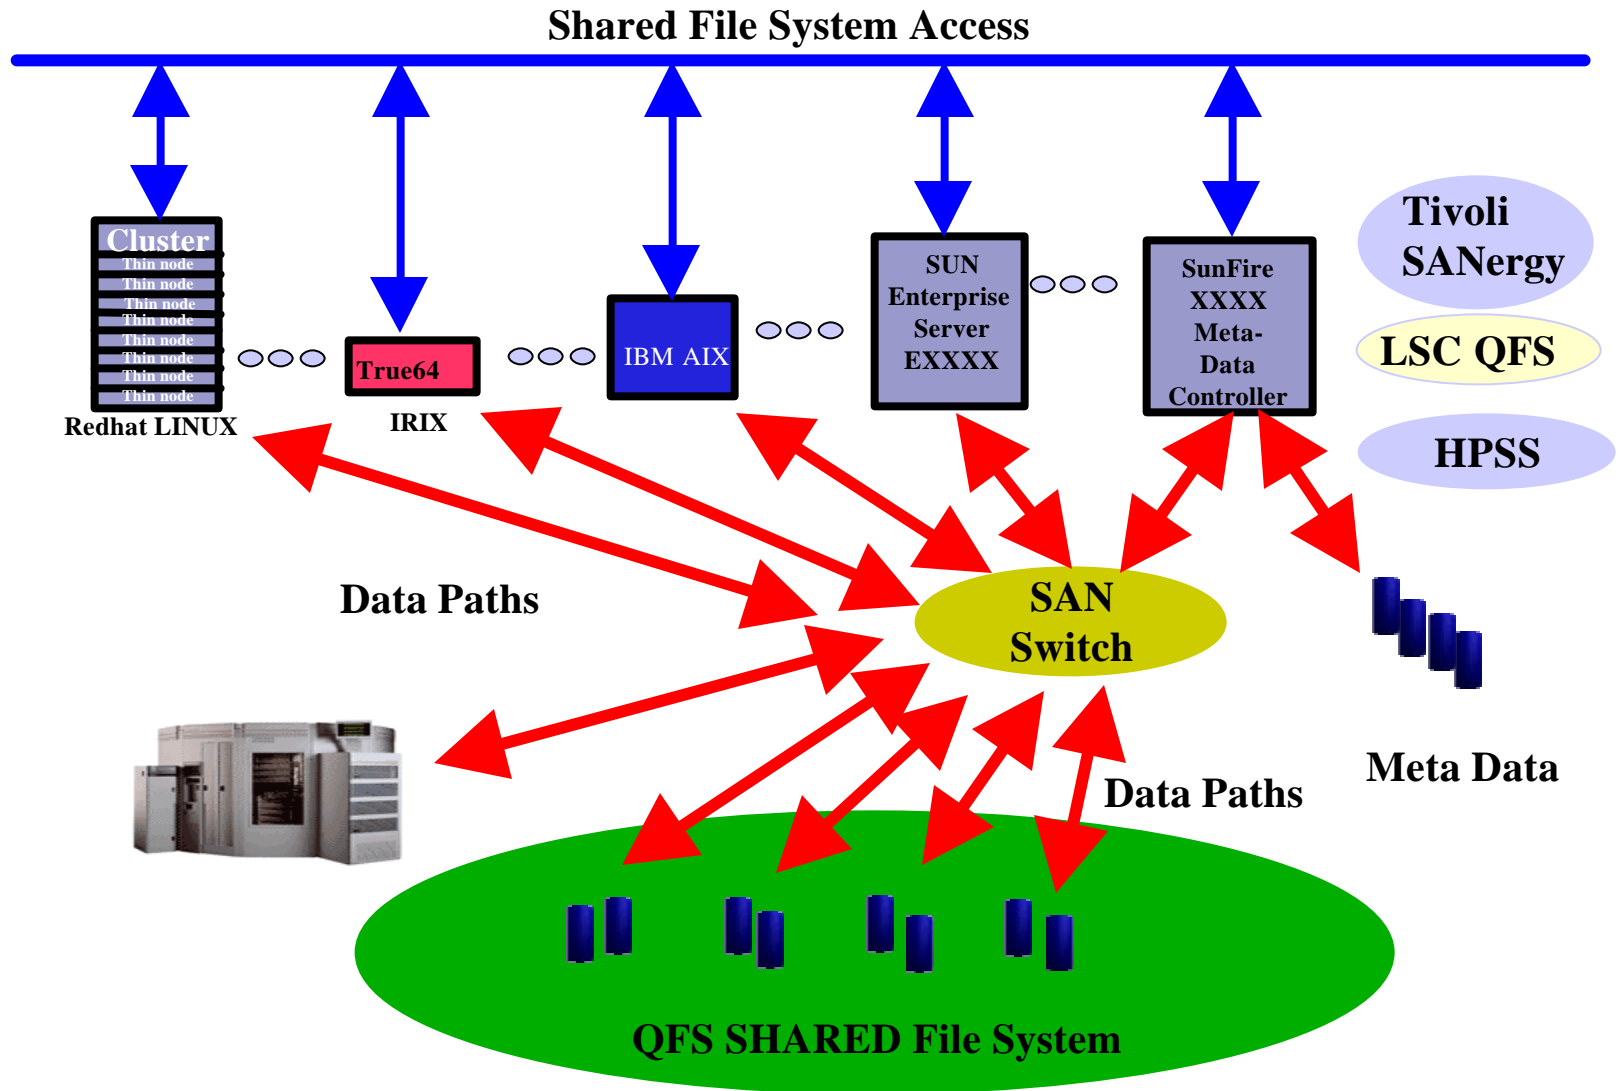
High Performance Computing Storage Area Network

Provides usual UNIX file system functionality across multiple fabric links

Provide access to a truly shared, potentially heterogeneous, storage subsystems in order to backup, store, and share data

While most provide functionality,
THEY GENERALLY DO *NOT*
PROVIDE
HPC-CLASS PERFORMANCE!

HPC SAN Hardware And Software Stacks





HPC Product Components

IBM -- HPSS

LSC Inc. – Sun Acquisition -- QFS

HPC Best-of-Breed Performance File System

Tivoli SANergy

Proven Heterogeneous SAN Support

Architecture Reduces Impact of Vendor OS Changes on Porting

Teaming with Sun to push SANergy into terascale solution space and HPC class performance

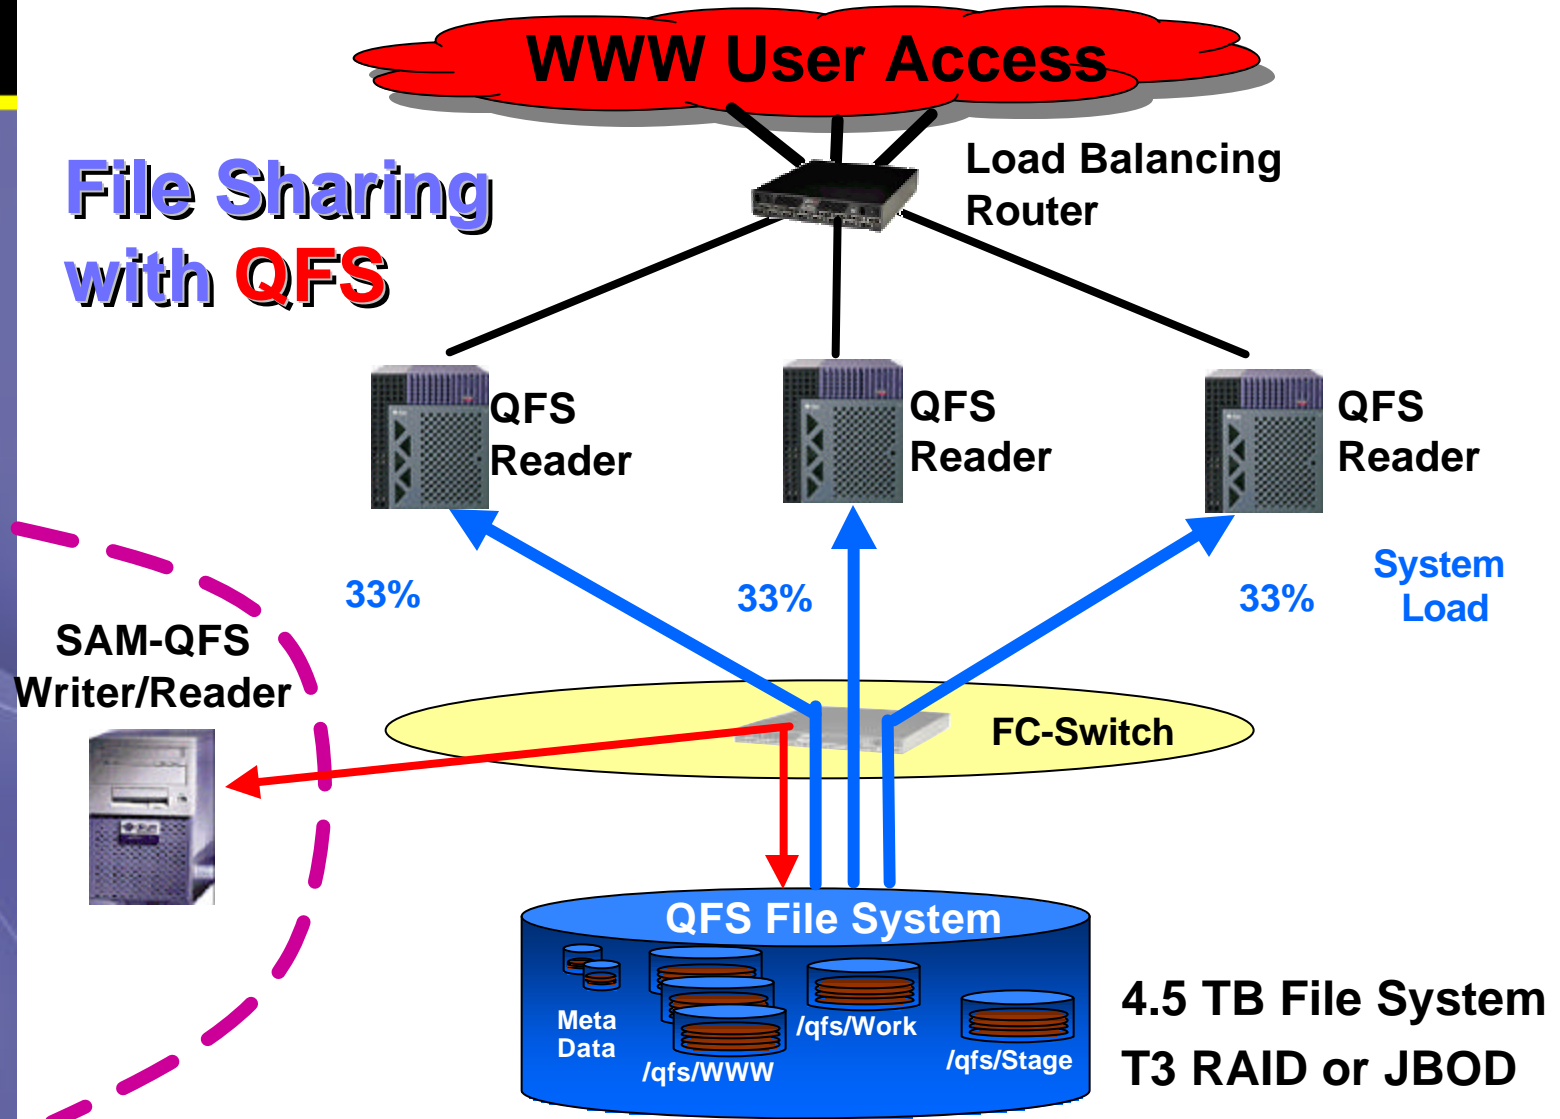
Qlogics and Brocade

Fibre Channel HBAs & Switches

Veritas

HA for Metadata Controller

File Sharing with QFS



Value of HPC SAN

Critical **Performance** Link in Throughput Computing

Reads and Writes @ Near Local Disk Speed.

Increased Efficiency = Increased Application Cycles

Decreases Time to Solution

Cross-platform **Connectivity**

Cross-platform **Data Access**

Shared Single File System

Multiple Heterogeneous Compute Engines

Data Management **Control**

Stable Shared Data Environment

Consolidation of Data Services and Management

Value of HPC SAN

“Lab Results”

Multiple High-Speed Data Streams

Scalable HPC Performance **Demonstrated**—Per 12-Wide File System/Stripe Group - single file to/from single process

> **1100 MB/s read !**

> **1000 MB/s write !**

Previous Generation Hardware and Software

E6500, E10000

T3 RAIDs

FC-AL—1 Gb/sec

LSC QFS/Tivoli SANergy (pre-releases here only)

Multiple Simultaneous Readers and Writer

Read-while-Writing At Maximum Efficiency



HPC SAN

Summary/Discussion

Provides scalable, reliable, and manageable storage services to meet the exploding data requirements of the distributed marketplace

Massive horizontal and vertical scalability

Continuous uptime

Highly reliable storage of irreplaceable data

Fully integrated set of products and services for manageability

Provides HPC-CLASS performance for homogeneous/heterogeneous configurations