

# Spider: A Global File System with Lustre and LNET Routers

Shane Canon  
**Group Leader for  
Technology Integration**

Collaborators:  
David Dillow  
Kenneth Matney  
Sarp Oral  
Galen Shipman

**April 28, 2008**




# Outline

- **Overview of Spider and SION**
- **Current Status of Spider**
  - Evaluation of Storage Systems
  - Router Status
- **Spider Plans**
- **SION Status and Plans**
- **Risks and Mitigation Efforts**
- **Other Efforts**

# Project Overview

- **Shared File System (Spider)**
  - Accessible by all major systems
  - Will serve as the file system for Baker (1PF system)
  - Eliminates islands of data
- **Scalable I/O Network (SION)**
  - A high performance backbone for I/O (Beyond just the file system)
  - Allows for new capabilities
  - Scales as our compute resources grow

# Spider



# Spider - Benefits

- **Accessible from all major LCF resources so users don't need to move data around**
- **File System is de-coupled from Computational Systems**
  - Accessible during maintenance windows
  - Can evolve on an independent trajectory
  - More flexibility in hardware (i.e. fatter SMP nodes from MDS and OSS)
  - Easier to incorporate site customizations for specific needs
  - Accessible during maintenance windows
- **Maximizes the impact of storage investment**
- **Enhances manageability**

# Spider Status

- **A smaller scale test-bed is currently available**
  - **5 GB/s of aggregate bandwidth**
  - **650 TB usable space**
  - **Accessible from a number of compute resources**
    - **Jaguar (Cray XT)**
    - **Ewok (“End-to-End” cluster)**
    - **Lens (Analysis Cluster)**
    - **Smoky (Development Cluster)**

# Router Status

- **Currently testing with six routers on 3500 node XT4 System (PCI-X)**
- **Worked through several issues**
  - **Deadlock in router configuration**
  - **Still some timeout issues to resolve**
- **PCI-E based router also tested**
  - **~ 1.5 GB/s for RDMA**
  - **~1.25 GB/s for LNET routing**

# Storage Bakeoff

- **Completing evaluating two storage technologies**
  - **DDN 9900 (3 Couplets) with 30 8+2 tiers (RAID 6) of 1TB SATA**
  - **LSI XBB2 (4 Controllers) with 24 8+2 tiers (RAID 6) of 1TB SATA**
  - **Both systems around 1 PB of raw storage**
- **Evaluating on**
  - **Performance**
    - **Random and Sequential Read/Write under multiple configurations**
  - **Reliability**
    - **Verification of Parity Check on Read**
  - **Robustness**
    - **Time to rebuild LUNs including performance under rebuilds**
    - **Performance impact of LUN scrubbing, Parity Check on Read, etc.**
  - **Packaging**
    - **Density, Serviceability**




# Performance Results

- **XDD was used to evaluate performance of these systems**
  - Random and Sequential Read/Write concentrating on spindle performance (long lived I/O)
  - 1MB request size (mimics Lustre stripe width)
- **Hero numbers are best case**
  - LSI: disable parity check on read, disable mirroring, cache enabled
  - DDN: disable controller coherency, cache enabled on sequential / disabled on random
- **Operating configurations numbers are a POSSIBLE configuration for production**
  - LSI: enable parity check on read, cache mirroring, cache enabled (doesn't impact random workload)
  - DDN: enable controller coherency, cache disabled (can't mirror cache for failover)

# Cray “Baker” - 1 Petaflops System

## FY 2009: Cray “Baker”

- 1 Petaflops system
- 37 Gigaflops processor
- 27,888 quad-core processors  
Barcelona 2.3 GHz  
(111,552 cores)
- 2 GB/core; 223 TB total
- 200+ GB/s disk bandwidth
- 13,944 dual-socket, 8-core  
SMP “nodes” with 16 GB
- 6.5 MW system power
- 150 cabinets, 3,400 ft<sup>2</sup>
- Compute Node Linux operating system
- Torus interconnect
- Planned install in Nov. 2009



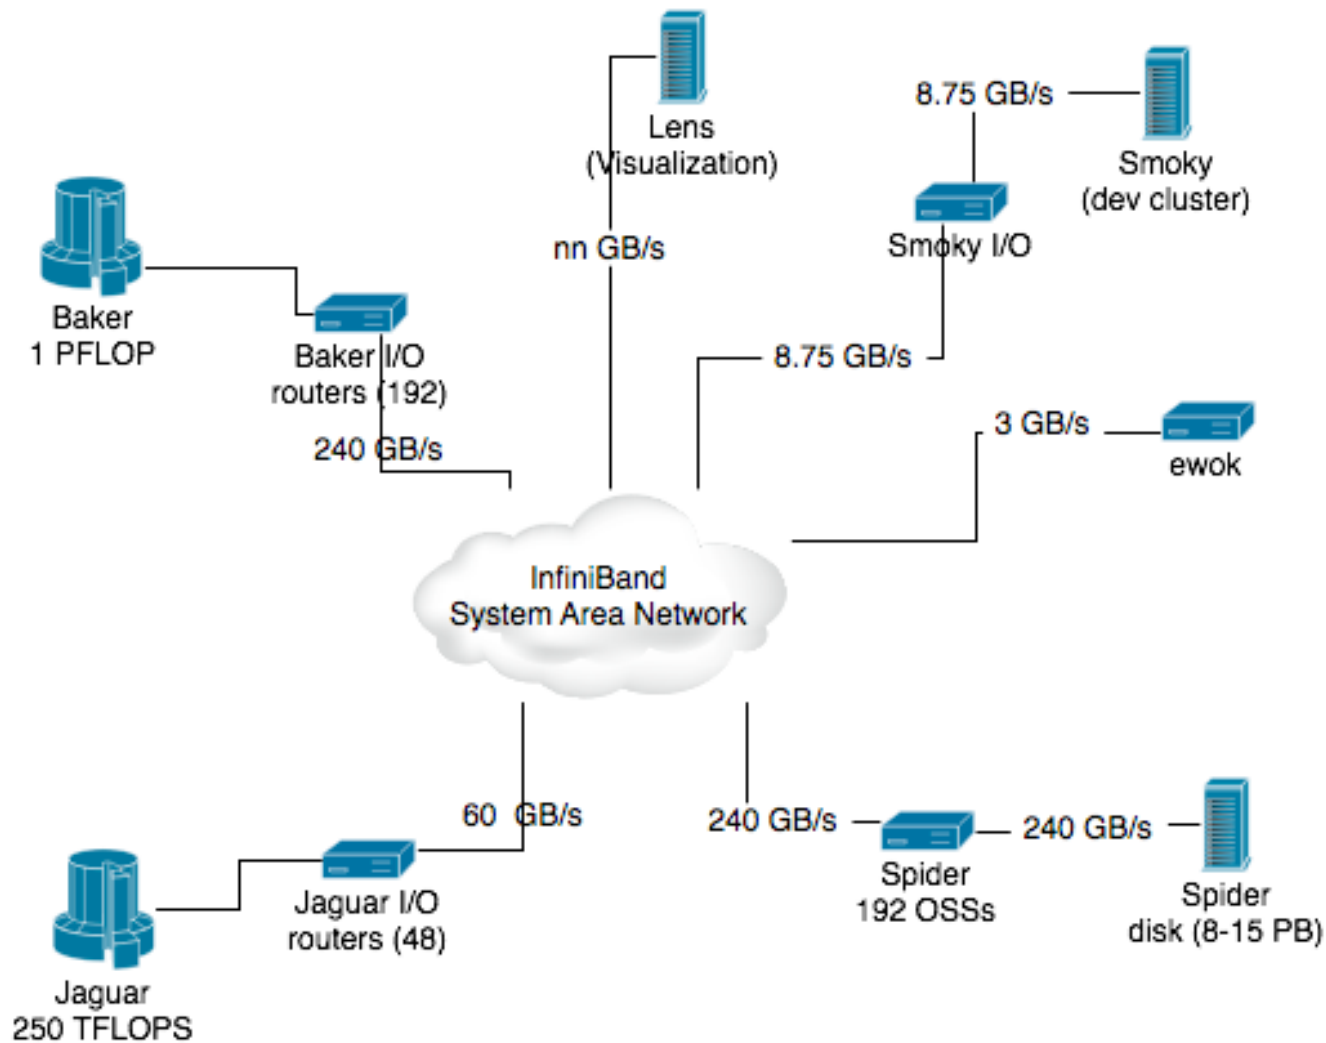
# Baker PF System and Lustre

- **SeaStar/Portals based (Like XT4 today)**
- **Spider will serve as THE file system for the ORNL Baker system**
- **Will rely heavily on LNET Routers**

# Spider - Lustre Infrastructure

- **Object Storage Servers**
  - 192 @ 1.25 GB/Sec = 240 GB/Sec
- **Object Storage Targets**
  - 6-7 per OSS @ 320 MB/Sec > 1.25 GB/Sec
- **Lustre Routers**
  - 192 @ 1.25 GB/Sec = 240 GB/Sec
- **MDS**
  - Maximize IOPs for a reasonable cost
  - Fat server (16 cores)
  - SAS/FC with very large cache is the near term solution


# Near term Spider Architecture (FY08)



# SION Status

- **A smaller scale test-bed is currently online**
  - Mix of 4x DDR and SDR InfiniBand
  - Performance is adequate for aggregate Spider performance
  - PCI-X riser cards on Cray XT machines a limiting factor
  - Prototype PCI-e risers have been tested in house
- **SION infrastructure has been specified**
  - 4x DDR InfiniBand
  - Large quantities of high quality optical cables (in production)
  - Small port counts per switch (288) is still a limiting factor
    - No, we don't need a Magnum, but > 500 ports would be nice
- **Full scale architecture will be online in FY08**

# Near Term SION Infrastructure



# Spider – Risk/Disadvantages

- **Becomes a single point of failure for the center**
- **Dealing with version skew between platforms**
- **Dealing with Quality of Service for different platforms (i.e. Viz slowing down simulation checkpoints)**
- **Router/congestion Issues**



# What Keeps Me Up at Night

- **Scale of Baker System**
  - 14k clients (100k task), 192 OSS, 1300 OSTs, 192 routers
- **Stability/resiliency – Timeouts**
- **Congestion – In the routers, in the network (both sides), at the OSSs**
- **Data Integrity – Over the network, on disk**

# What are we doing about it?

- **Working with both Sun and Cray to identify potentials issues and understand what tricks (both new and under development) can be brought to bear.**
  - Adaptive Timeouts
  - Network Request Scheduler
  - Improved routing
  - Ping changes
  - Improvements in LNET timeouts
  - Improving middleware layers (ADIO, HDF5, NetCDF) and working with Apps Teams to improve IO techniques

# Network Congestion

- **I/O infrastructure doesn't expose resource locality**
  - There is currently no analog of nearest neighbor communication that will save us
- **Multiple areas of congestion**
  - Infiniband SAN
  - SeaStar Torus
  - LNET routing doesn't expose locality
    - May take a very long route unnecessarily
- **Near term solution - Expose Locality**
- **Long term solution - Better Hardware?**

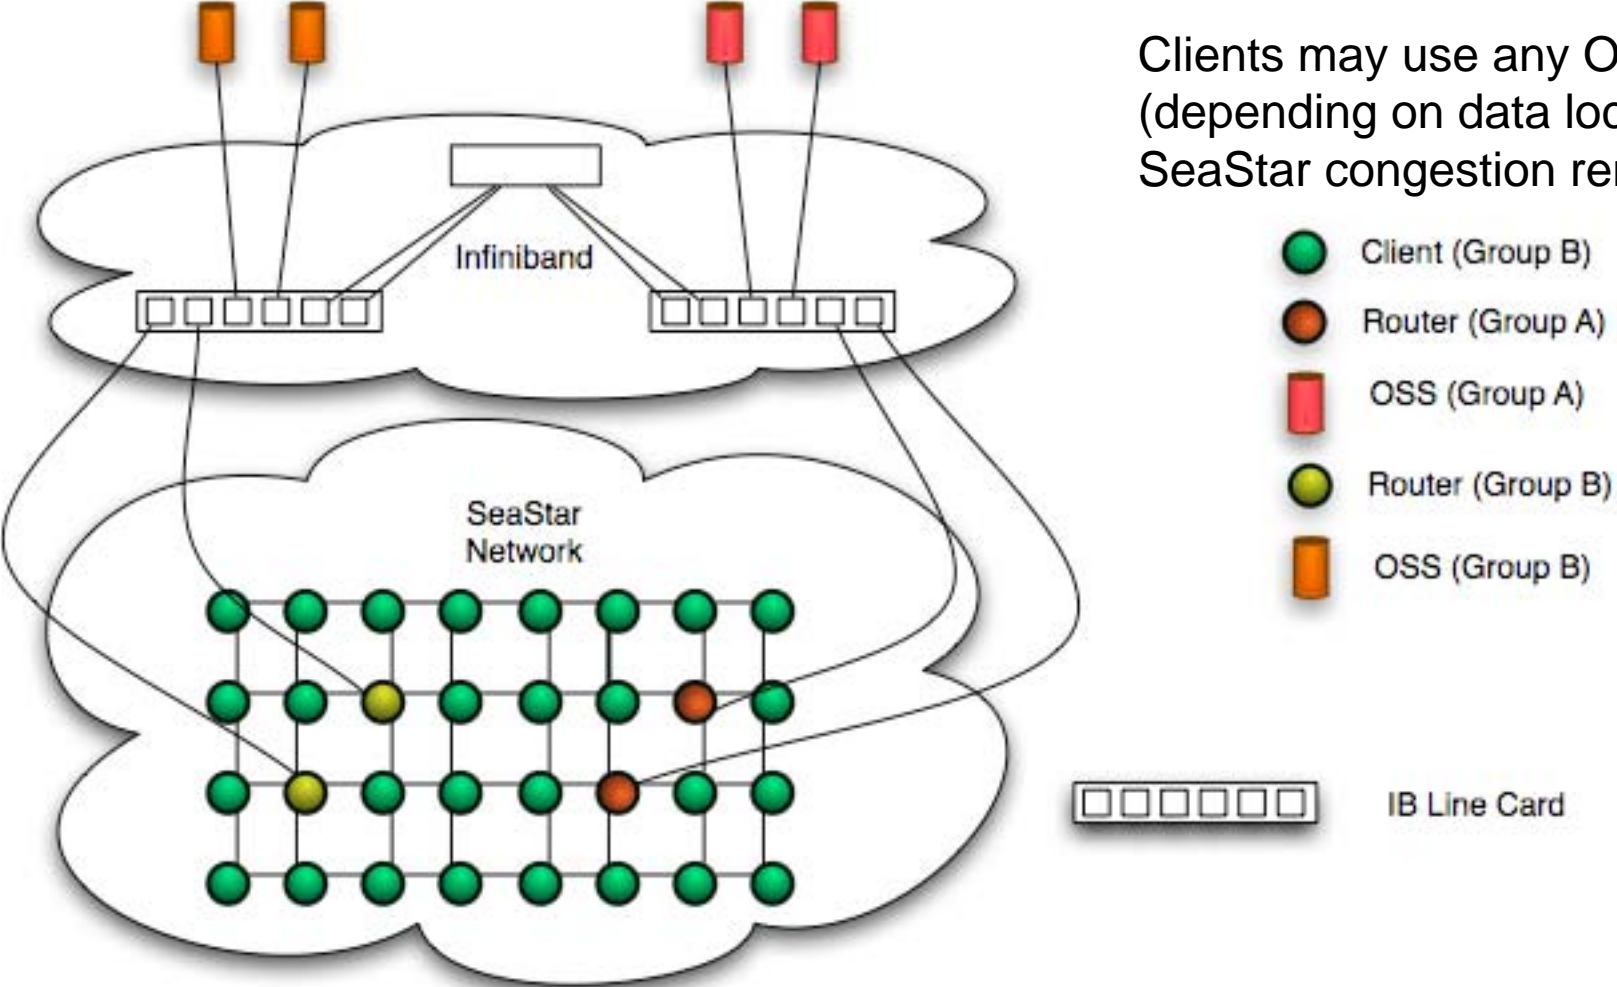
# Design To Minimize Contention

- **Pair routers and object storage servers on the same line card (crossbar)**
  - So long as routers only talk to OSSs on the same line card contention in the fat-tree is eliminated
- **Place routers strategically within the Torus**
  - In some use cases routers (or groups of routers) can be thought of as a replicated resource
  - Assign clients to routers as to minimize contention

# Intelligent LNET Routing

Clients prefer specific routers to these OSSes - minimizes IB congestion

Clients may use any OSSes (depending on data location) - SeaStar congestion remains



# Parallel file utilities

NCCS researchers are developing tools to help NCCS users. One example is a suite of tools to address the day-to-day needs of users for file management. These include parallel-capable tools to copy, compress/decompress, and archive (tar) data. These tools will leverage the parallel file systems by distributing these operations across many worker clients.

## 1. SPDCP – Staging Parallel Distributed CP

- Allows an arbitrary tree or set of trees to be copied from one location to another, propagating LUSTRE striping information
- Balances the load of the copy, exploiting parallelism in I/O at the file level and the data level

## 2. Parallel Utilities Framework

Allows the user to fully leverage the capabilities of the parallel file system.

- Takes an arbitrary LINUX command and executes it in parallel across the Cray XT
- Can execute "sed," "bzip2" compression/decompression, and many other LINUX commands in parallel
- Potential to scale to 1,000 processors; everything is parallel (including I/O) with minimal synchronization requirements

# Contacts

**Shane Canon**  
**ORNL**  
**(864) 574-2028**  
**canonrs@ornl.gov**



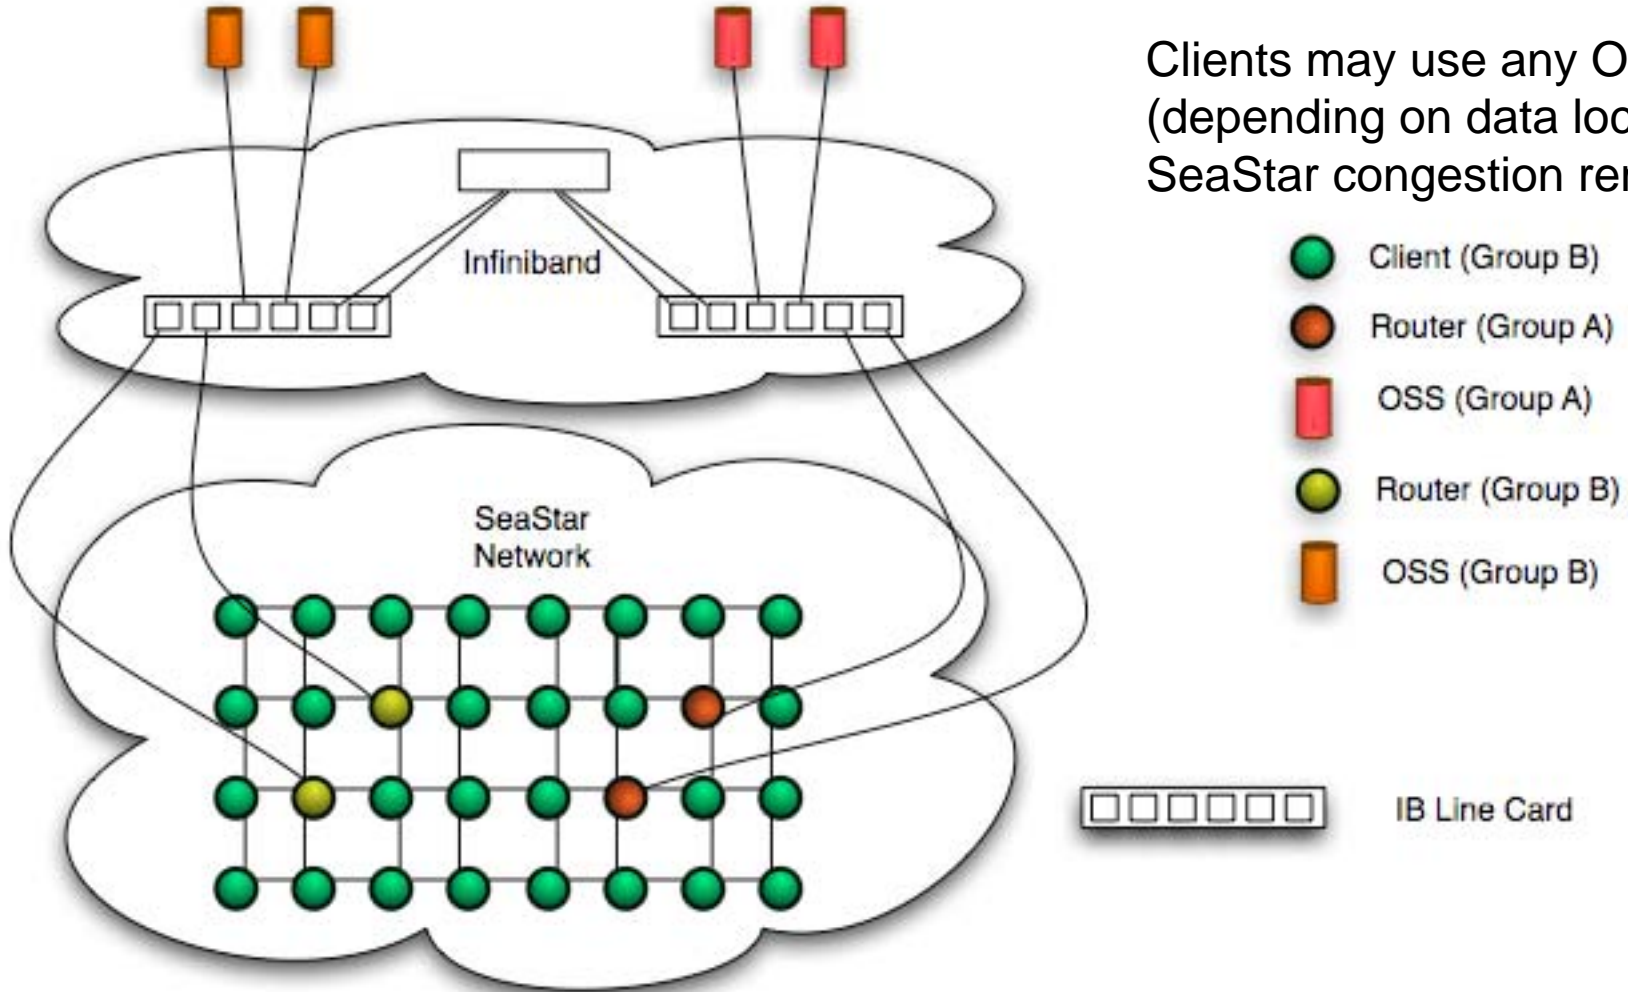
# Backup Slides



# Intelligent LNET Routing

Clients prefer specific routers to these OSSes - minimizes IB congestion

Clients may use any OSSes (depending on data location) - SeaStar congestion remains




# Future Plans for Spider



**Baker**



**Jaguar**  
**Cray XT3/XT4**



**Data Analysis  
& Visualization**

**ESnet, USN,  
TeraGrid,  
Internet2, NLR**



**HPSS**

## 2007

- 1.6 PB
- 30 GB/s (aggregate)

## 2008

- 8-12 PB
- 240 GB/s (aggregate)

# Intelligent LNET Routing

Clients prefer specific routers to these OSSes - minimizes IB congestion

Assign clients to specific OSSes - minimizes SeaStar Congestion

