
Jefferson Lab

2006 LQCD Analysis Cluster

Chip Watson

Jefferson Lab, High Performance Computing

Analysis Cluster Goals

- ▶ **Additional Capacity for Data Analysis**
 - ▶ Clusters were oversubscribed, additional capacity needed primarily for data analysis
 - ▶ Goal: 400 GFlops, to be deployed as early in the year as possible (funding constrained)
- ▶ **Extend Infiniband Experience to Jefferson Lab**
 - ▶ Anticipating that the FY 2007 machine might be sited at Jefferson Lab, the project wanted Jefferson Lab staff to gain the necessary experience on Infiniband

Strategy

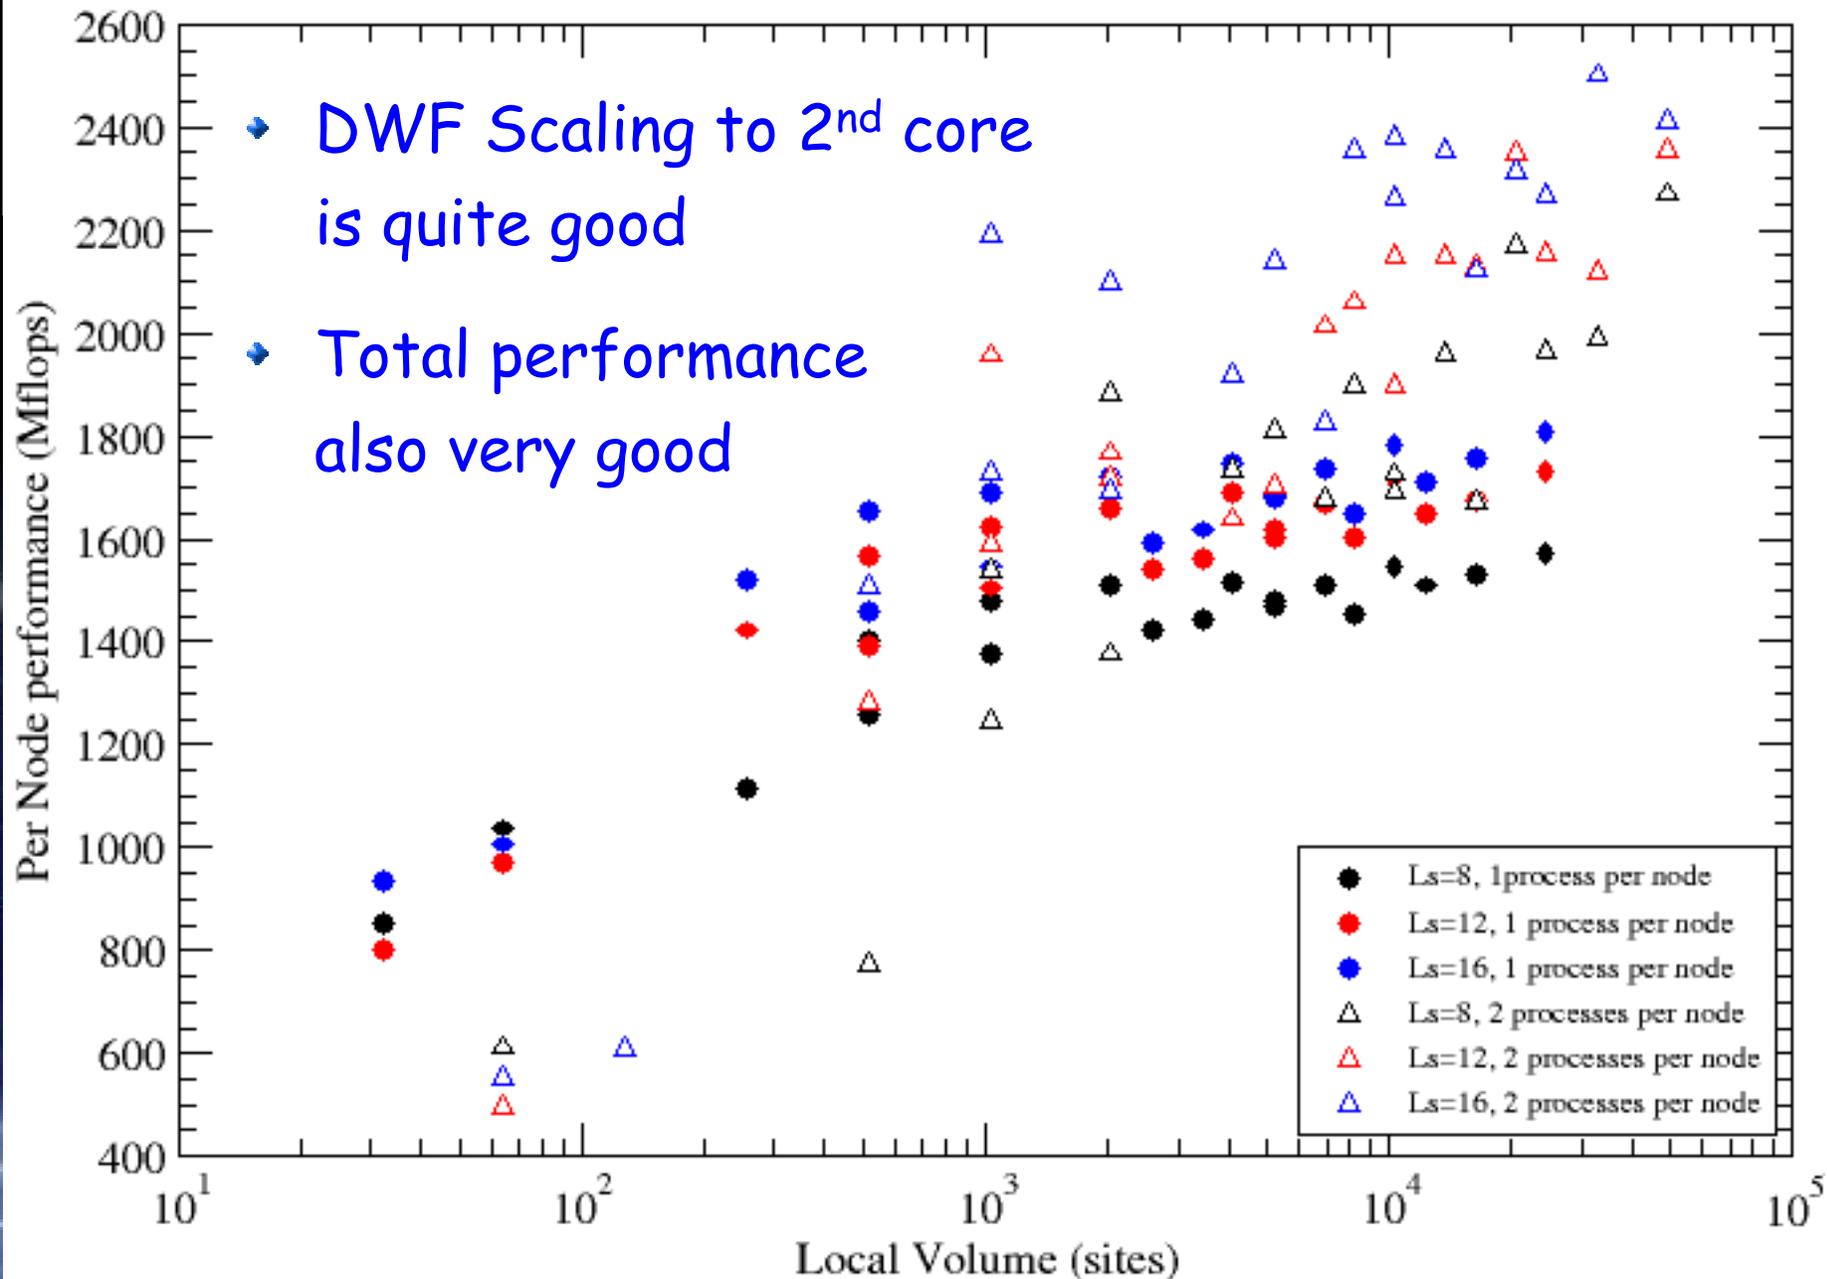
- ◆ As presented last year, the plan was to replicate the Fermilab “pion” cluster (at reduced scale) at JLab, combining funds from the last year of SciDAC (prototype hardware) and the first year of the LQCD Computing project.
- ◆ One promising new alternative (Intel dual core Pentium-D) was evaluated, leading to a decision to make the first step in the direction of dual core CPU's with this cluster...

SciDAC Prototyping

- Under the SciDAC project, the collaboration has for the last 5 years been evaluating a series of commodity technologies for optimized clusters for LQCD.
- As a continuation of this work, the Pentium-D was evaluated in the Fall of 2005.
 - Expectations were that dual core scaling would be poor, as it has been for the last 4 years with dual CPU Xeons.
 - We were pleasantly surprised!

Per Node Performance vs. Local Volume

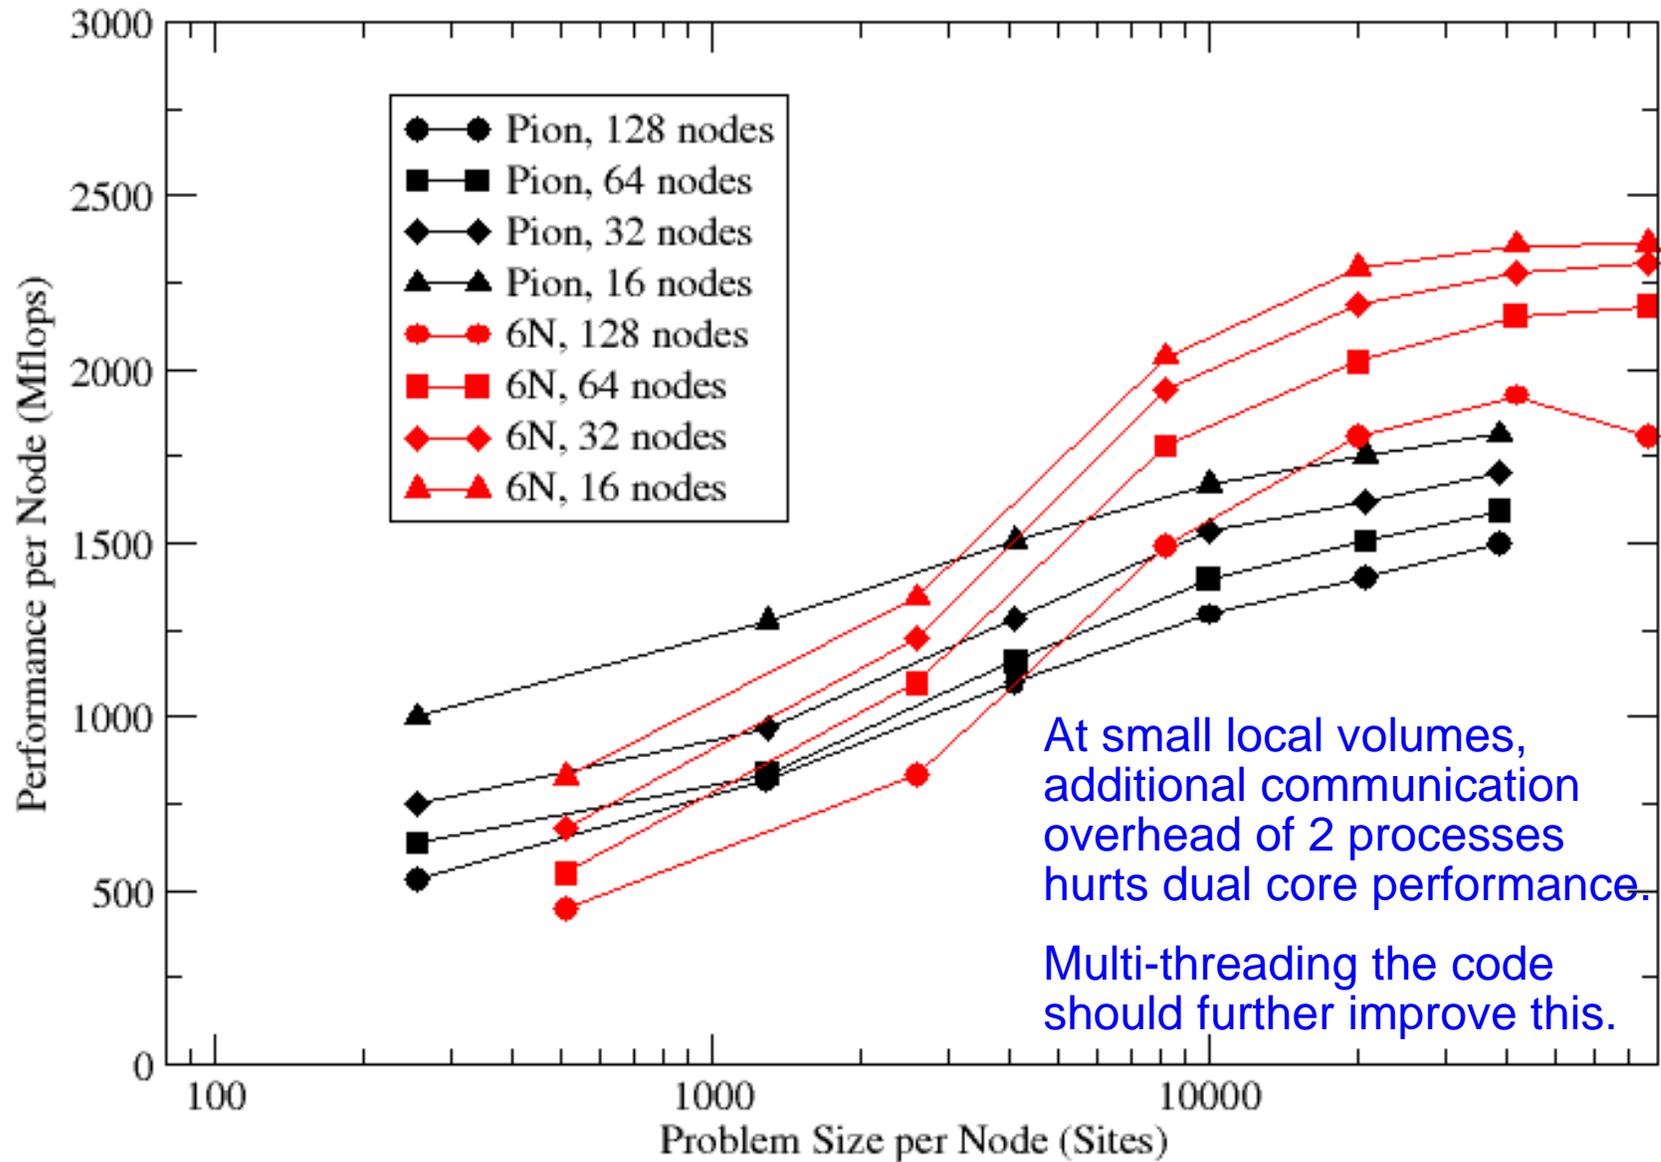
128 Nodes mapped as 128 cores as 1 4 4 8 and 256 cores as 1 4 8 8)



➤ DWF Scaling to 2nd core is quite good
 ➤ Total performance also very good

Comparison of MILC Benchmark on JLab 6N and FNAL Pion Clusters

per node comparison - Jlab runs 2 processes per node=>per node volume = 2*FNAL



Decision Process

- Option 1: stay with "pion" for deployment
 - Pro: known platform, lower risk
 - Con: give up some potential performance gain
- Option 2: switch to Pentium-D for deployment
 - Pro: 22% increase in performance (at fixed cost)
+ start on experience with dual core
 - Con: slight risk in using new chip, motherboard

Note: same HCA; can use 2 cores in multi-process mode so no software changes anticipated initially

Decision Process & Plan

- ♦ Change Control
 - ♦ Cost neutral
 - ♦ node cost 10% higher, but we can buy fewer nodes
 - ♦ Slight increased risk
 - ♦ 20% improvement in performance and price / performance
- ♦ Large Scale Prototype First
 - ♦ Use JLab SciDAC funds, plus JLab theory group funds to build a prototype of scale 128 nodes
 - ♦ Once operational, use project funds to add an additional 128+ nodes

Timeline

Sept 2005 - procured Infiniband fabric for 130 nodes (end of year funds)

Oct 2005 - procured two Pentium-D nodes for testing

November 2005 - Two node performance data very encouraging (Streams, Chroma DWF application, MILC staggered); decision taken to use SciDAC + base funds to procure 128+ node cluster for large scale test of dual core platform (approved by SciDAC oversight committee)

December 2005 - Procurement of 140 nodes, half from SciDAC funds, half from JLab theory group (this project decides to await full scale tests)

Jan/Feb 2006 - Cluster installed at JLab. Based upon good performance (and behavior), Change Request submitted for LQCD Computing Project Change Control Board recommending to procure dual core instead of "pion" nodes

Feb 2006 - Change Request approved, additional 140 nodes ordered

March/April 2006 - Machine commissioned (tested with several versions of IB and Linux O/S)

May 1, 2006 - Start of production running

Cluster Details

- ▶ 280 node 2006 Infiniband cluster - 6n
 - ▶ Dell 850
 - ▶ 3.0 GHz Pentium-D dual core
 - ▶ 1 GByte DDR2-667 memory (800 MHz fsb)
 - ▶ 80 GB SATA disk
 - ▶ IPMI for node monitoring, control
(can reboot hung node from home)
 - ▶ Mellanox PCI-Express 4x "memfree" cards (same as FNAL)
 - ▶ 24 port Flextronics switches (same as FNAL)
 - ▶ 35 nodes / rack, 17 or 18 nodes per leaf switch (2 per rack)
 - ▶ spine switch built from 5 of 24 port switches (modular and fault tolerant)

6n Cluster

- ◆ 2:1 oversubscription with smart layout, else 3.5:1
- ◆ up to 2.5 GFlops / node DWF, 2.3 staggered
- ◆ Single job on combined resources achieved **> 600 GFlops, \$0.8 / Mflops**
- ◆ Operational since May 1



Results

- ◆ 480 GFlops added to project analysis capacity
 - ◆ Project procurement: 140 nodes, \$250K plus spares, installation
 - ◆ 2.3 GFlops / node (average, sustained) => \$0.80 / MFlops
 - ◆ SciDAC (70 nodes, ~0.16 TFlops) + Project (140 nodes, 0.32 TFlops) yields 480 GFlops managed by this project
- ◆ Two staff members (Balint Joo and Jie Chen) have now gained proficiency with Infiniband
 - ◆ Balint Joo manages the fabric (upgrade firmware, run diagnostics)
 - ◆ Balint identified an error in latest IB-2 stack
 - ◆ Jie Chen localized it to bad memory management in the eager send protocol when shared received queue is enabled (default for PCI-Express platforms); we have now disabled this mode
 - ◆ Bug report submitted, bug fix received last week

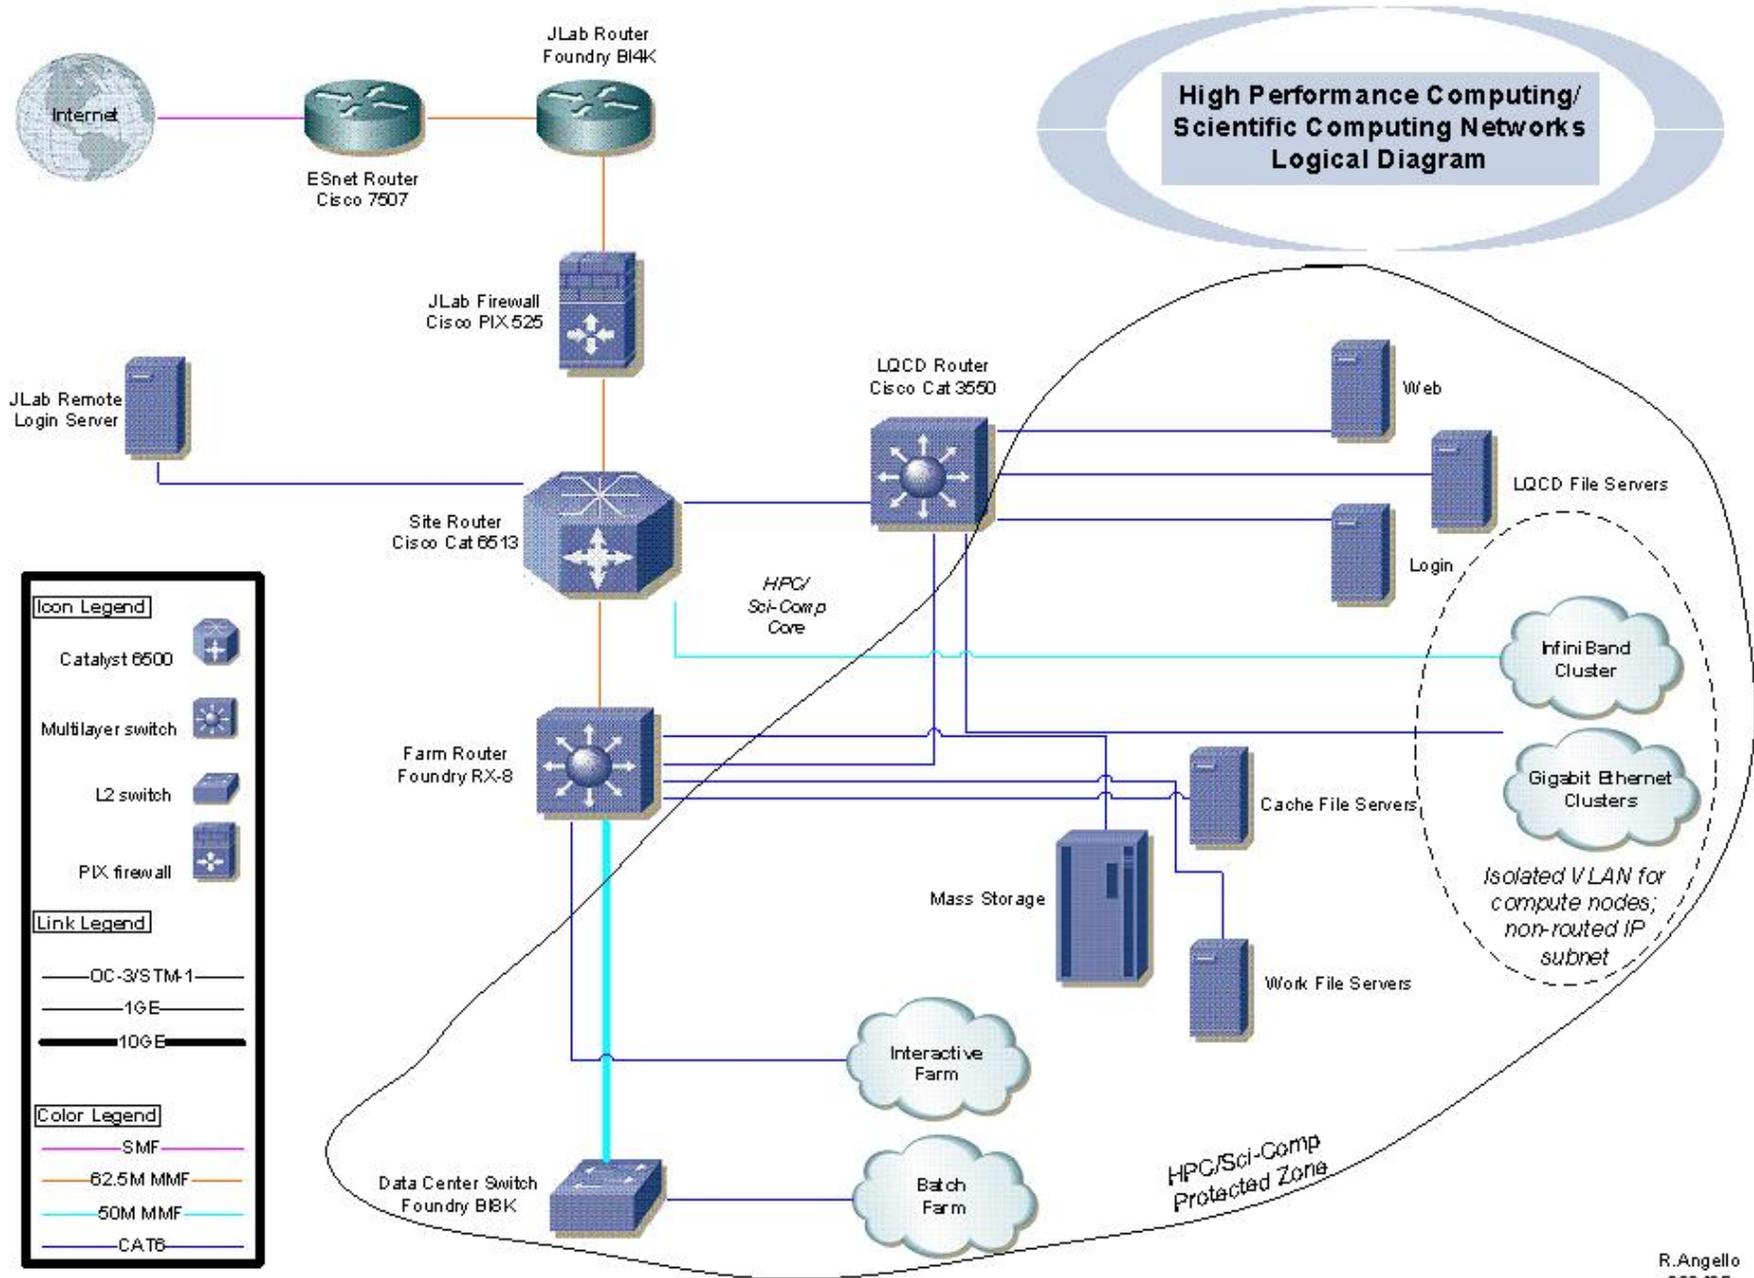
Summary

- JLab cluster "6n" is now in production
 - Price / performance \$0.80 / Mflops
- JLab is now running over 1000 nodes, with infiniband in the mix
- JLab is ready to host the FY 2007 cluster - lots of space!



(Backup Slides)

Architecture Diagram



Operational Data

- Approximately 50 users, 20% active
- Most jobs on gigE mesh clusters are 128 node (4x4x8)
- Uptime for the year was below 90%, due to installation of new wing (& computer room); currently doing > 90%.
- All first year allocations have now been delivered (except for one small project that didn't use their time); most projects will end the first "year" with 15% extra CPU cycles since this resource was not formally allocated for the May-June friendly running period.
- 60 terabytes now in JLab silo (< 5% of resource); project pays only for tape, not slots nor share of robotics
- Many detailed views of the computing systems, batch system and resource usage available via the web, <http://lqcd.jlab.org/>