Jefferson Lab Facilities

Chip Watson
Jie Chen, Ying Chen, Balint Joo

Outline

- Compute Resources: 1000 nodes!!!
- Storage: 15 terabytes (and growing)
- Batch System: Torque+Maui
- User Support, Staff increases, Web Reports, ...

Compute Resources

- Newest:

  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 (reboot hung node from home)
  - IB 4x cards, 17-18 nodes per leaf switch, core switch built from 5 of 24 port switches (modular and fault tolerant)

  - up to 2.5 GFlops / node DWF, 2.3 staggered
  - Single job **>600 GFlops, $0.8 / Mflops**
  - Testing now, operational May 1

6n cluster

Move to dual core

- As part of the SciDAC project, JLab evaluated dual core Pentium D in the Fall as an alternative to the pion clusters single core:
  - same bus speed, slower clock speed at constant price
  - 1 + 1 MB cache vs 2 MB cache
- Naively, one would expect no gain for large problems (memory bandwidth bound)

  In fact, a significant performance boost found
- No software changes required:

  just run 2 processes per node
- Additional gains expected from multi-threading

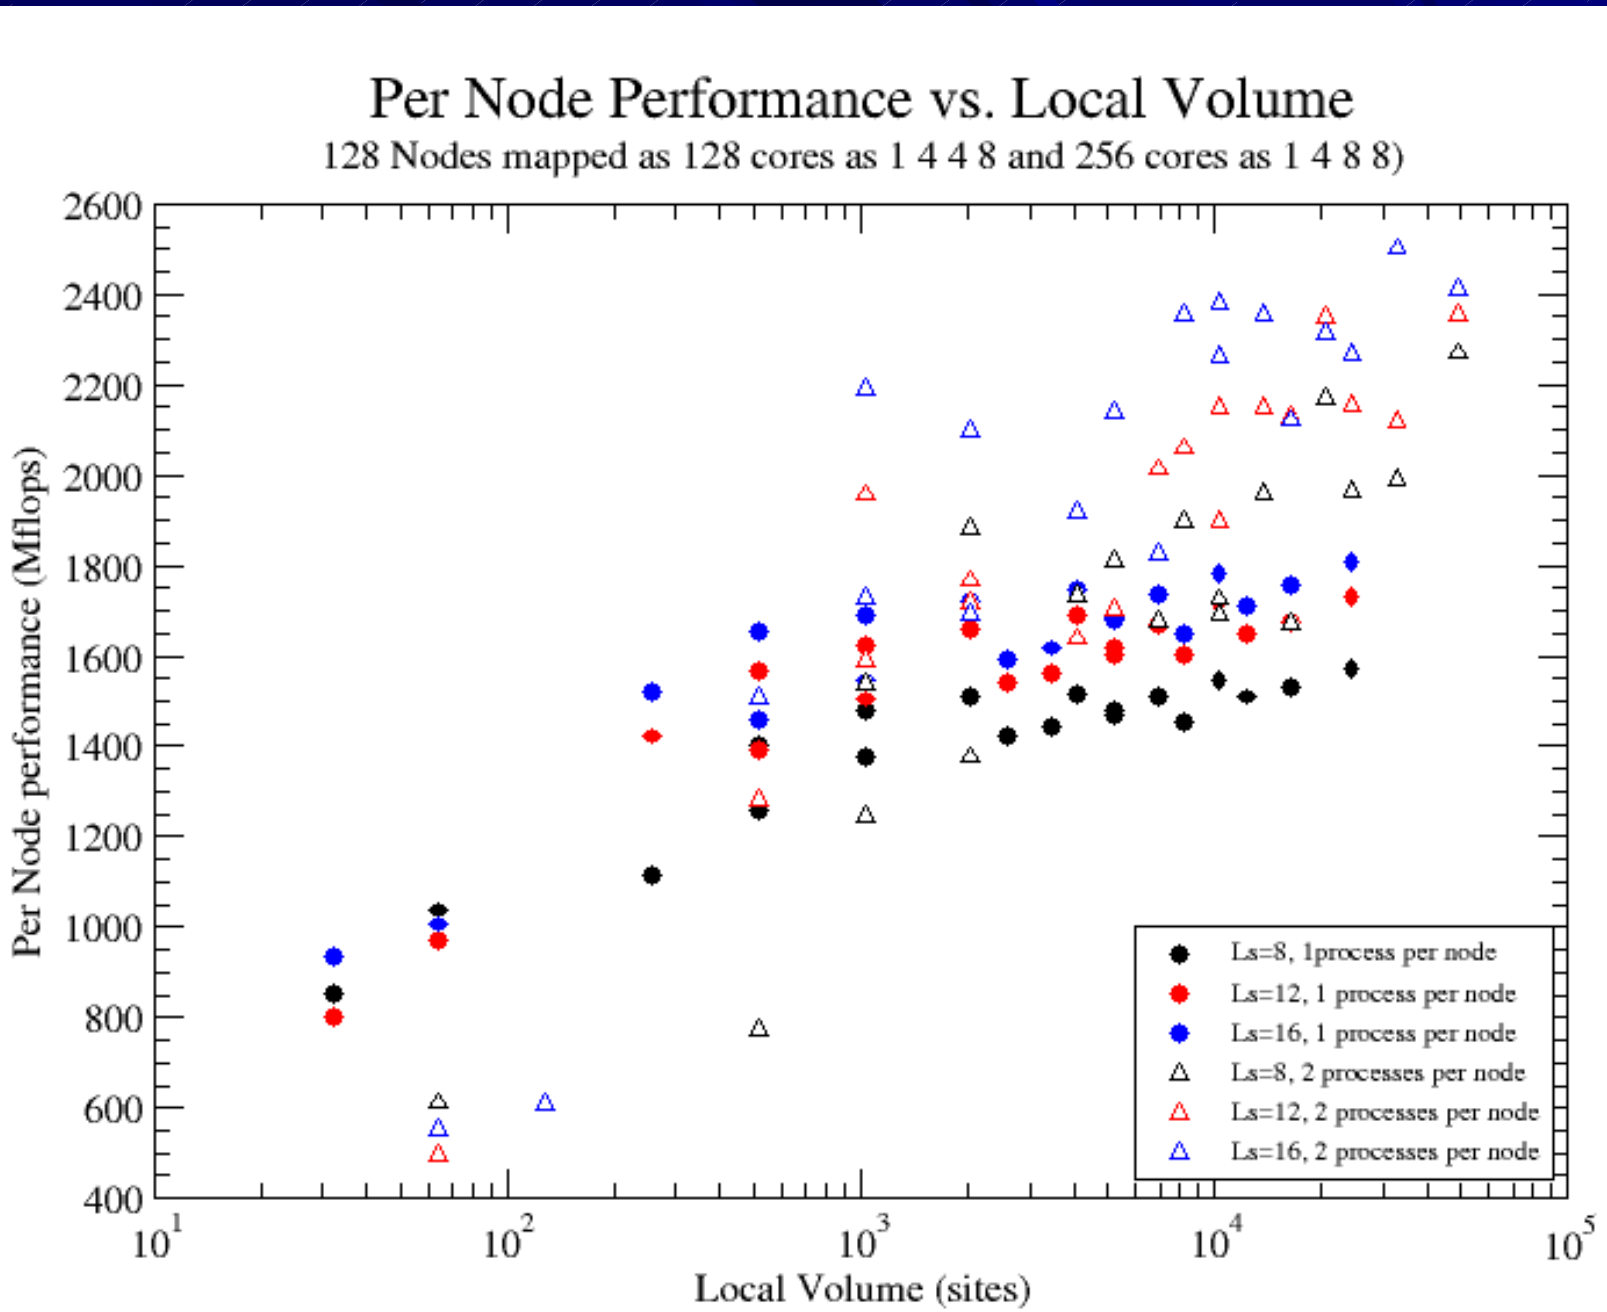
Per Node Performance vs. Local Volume
128 Nodes mapped as 128 cores as 1 4 4 8 and 256 cores as 1 4 8 8)
Per Node performance (Mflops)
Local Volume (sites)
2600
2400
2200
2000
1800
1600
1400
1200
1000
800
600
400
10^1
10^2
10^3
10^4
10^5
Ls=8, 1process per node
Ls=12, 1 process per node
Ls=16, 1 process per node
Ls=8, 2 processes per node
Ls=12, 2 processes per node
Ls=16, 2 processes per node

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

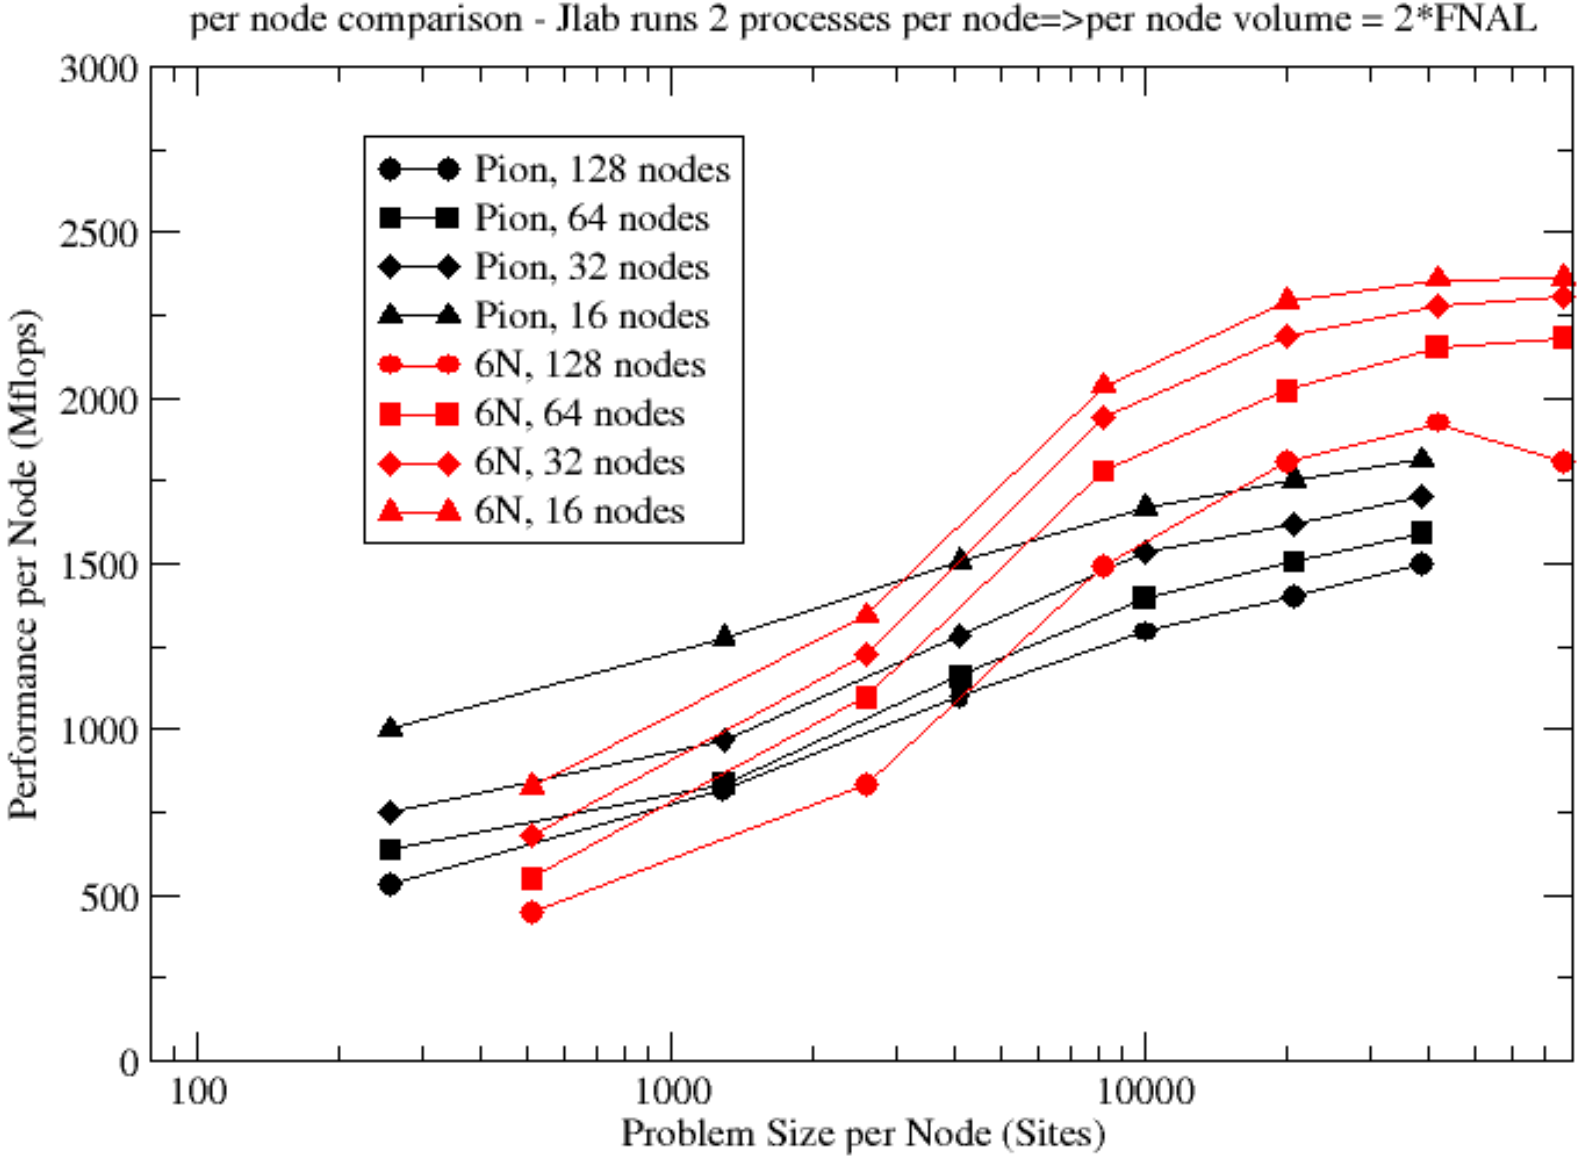

New Computer Room

- 7,500 sq ft
  - large enough for 4,000+ nodes (1U)
- 400 KVA UPS
  - 6n cluster uses about 50 kva
- 180 tons A/C
  - supports heat load of ~ 500 Kwatts
- UPS to be upgraded in 2008 (as needed)
- A/C to be upgraded in 2008-2009 (as needed)

gigE Mesh Clusters

- 384 node 4g cluster
  - Dell 2850
  - 7d gigE mesh: 6x2x2x2x2x2x2
  - usually 3 partitions configured as 4x4x8
  - ½ GB memory, 667 fsb
  - 40 GB disk, SCSI
  - IPMI
- 256 node 3g cluster
  - Supermicro / whitebox
  - 7d mesh, usually 2 partitions of 4x4x8
  - ¼ GB memory, 533 fsb
  - 30 GB disk, IDE

4g, 3g mesh clusters

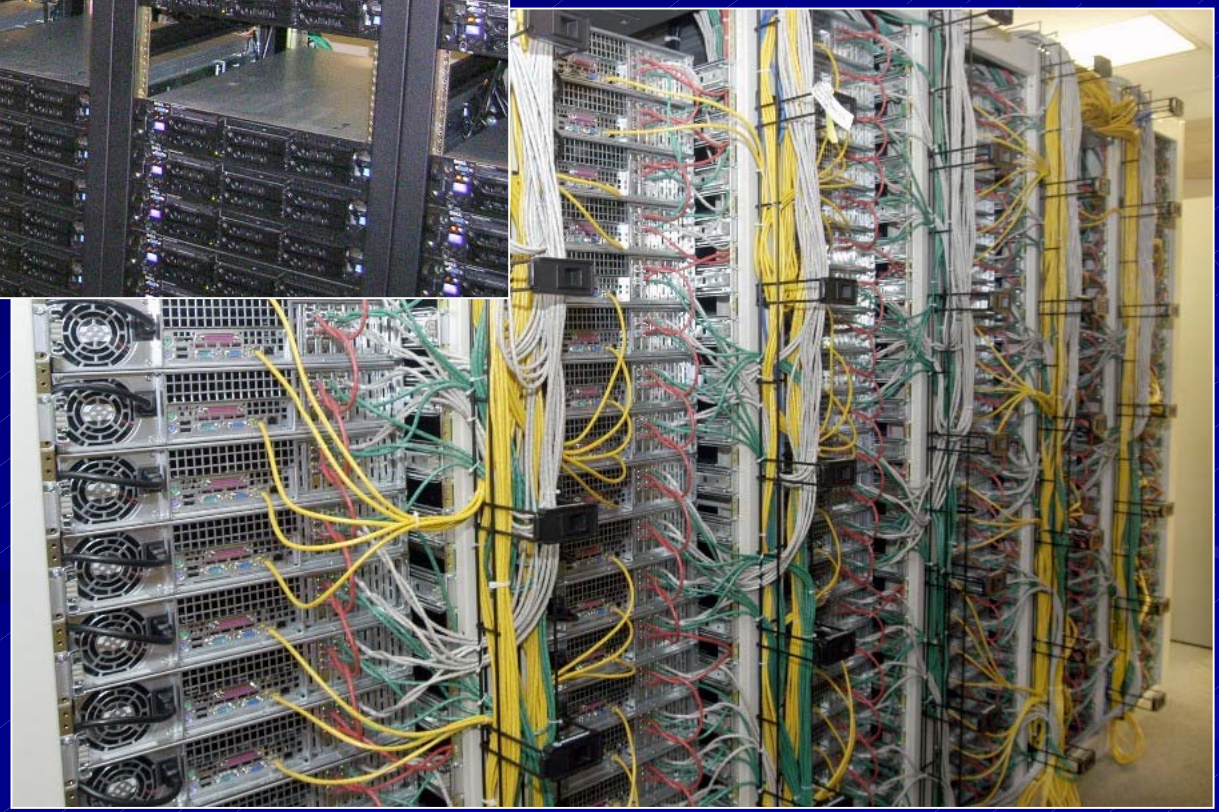

Other Compute Resources

- 64 node 2m cluster, myrinet,
  - being de-commissioned
- ~20 gigE test nodes
- 3 interactive nodes
  - drop back to 2 once 2m is de-commissioned
  - dual processor Xeon, 3 GHz, 800 fsb

Tools

- Primarily open source:
  - gcc, make, bison, editors, etc.
- Some Licensed software:
  - Intel C++
  - Soon to add: F90 (user requested), VTune, ...

Storage

- 5 file servers, 15 terabytes RAID
  - Additional 5+TB server to be added 2Q2006
- /home NFS mounted on all compute nodes
  - backed up by computer center
- /cache NOT mounted, accessed via rcp
  - not backed up; auto-migrate to silo
  - replacement of rcp (script) hides knowledge of where particular file is located (4 servers)

    rcp /cache/project/abc .
- local disks on compute nodes give exceedingly high parallel bandwidth for temporary files

Storage (2)

- Storage Resource Manager symantics
  - user managed storage (pros & cons)
  - policy based management (user controlled)
  - pin / unpin
  - permanent / volatile
  - auto migrate of permanent files to silo (large files) or mirror machine (small files)
  - auto delete of oldest not pinned files
  - 1 Petabyte silo, ~10% usage limit (new silo in FY2007 or FY2008)

Storage Challenges

- Running NFS with 1000 clients is challenging with commodity servers
  - most recently encountered problems popped up when new 6n nodes were built with SELINUX=1 (default on latest RedHat)
- Occassional data corruption
  - being diagnosed, might be a failing RAID controller (only occurs on one server);
  - corruption caught by checksum – data validation in Chroma or external tool
  - need better diagnostic tools

Storage Future

- Put file servers onto Infiniband
  - increase aggregate bandwidth several fold
- Plan, as part of SciDAC-2, to investigate other storage systems, incuding dCache
- Parallel file systems (tbd)
- New silo, dedicated drive(s) for LQCD
  - Lab will buy silo, we will buy tapes, probably one drive
- SRM version 2 migration
- ILDG support...

ILDG

- International Lattice Data Grid
  - Grid-of-grids, linking multiple collaborations
  - Plan to go operational June 2006!
  - Major pieces:
    - Metadata XML schema (standardized descriptions)
    - Middleware
- Web services based architecture
  - Metadata Catalog
  - Replica Catalog
  - File access
    - SRM – Storage Resource Manager
    - file servers: gridftp, http, ... (multi-protocol)
  - Membership (authorization services) – tbd

Batch System

- Torque version of Open PBS
- MAUI scheduler (as of Nov 1)
  - SciDAC project based fair share
  - mesh queue for the 5 gigE partitions (640 nodes) (most jobs are 128 nodes)
  - ib queue for Infiniband nodes (280 nodes)
  - test queue for extra gigE nodes, single nodes
  - (myrinet not on Maui)
  - user priorities (within their own jobs)

Batch System Challenges

- MAUI has hundreds of parameters
  - we are now using ~20
  - still gaining experience / understanding
  - fair share not exact:
    - MAUI uses fixed time window, not sliding window; large window gives correct long range behavior but poor short range behavior (biggest users dominate near window start); short window gives inaccurate long range behavior
    - 2 jobs starting concurrently can grab more than fair share (doesn't count against you until it completes)

User Support

- Lean staffing, getting better
  - Balint Joo added in Sept 2005 (long visa delay)
  - Hiring additional sysadmin to be shared with computer center (+ ½ FTE)
  - Will add another FTE in FY2007 to support next large cluster
- Trouble tickets:
  - Soon to release: LQCD web interface to JLab trouble ticket system
  - currently, email list accessible from main web page
  - good response on work days, poor on weekends
  - other plans in development for greater shift coverage

User Support (2)

- Web Interfaces
  - JavaFaces allows rapid creation of new views
  - Data sources:
    - Maui completed jobs database
    - Cluster monitoring (load, memory, node batch state)
- Standard User Environment (future)
  - Eventually make FNAL, JLab, BNL appear “the same” to users
    - file system layout, env variables, batch, ...

Jlab Lattice Portal - Mozilla

File Edit View Go Bookmarks Tools Window Help

http://lqcd.jlab.org/ Search

Google Jefferson Lab LQCD.jlab www.LQCD HPC Admin BBC NEWS Local weather HPC Wire Chinese-English JLab Insider

JLab Lattice Home

Status

Cluster Queues
2m mesh

Utilization
2m mesh

Maui stat/info
Fairshare Project

Cache Disks

Servers

Computer Room

News Archive

Documentation

User Environment
USQCD Standard
JLab Specific

Software
QMP QDP++
Chroma

Facilities
Clusters
Disk Cache

Links
Contacts

Active Job View | Completed Job View | Node Batch Status | System Status | Resource Status

Jlab 'mesh' Cluster Completed Jobs Status in 2 day(s)

User	Total finished Job	Success Job	Failed Job	Failed Ratio	System Utilization
all	153	150	3	2.0%	71.5%
edwards	1	1	0	0.0%	0.2%
engel	1	1	0	0.0%	0.1%
fhannon	2	2	0	0.0%	17.1%
flemingg	36	33	3	8.3%	37.5%
kostas	13	13	0	0.0%	0.4%
silas	97	97	0	0.0%	16.0%
ychen	3	3	0	0.0%	0.1%

Job Id	Job Name	User	Account	Exit	Nodes	CPU Time	Wall Time	Start Time	End Time	Queue
13981	pp.csh	silas	NPLQCD	0	128	00:00:31	05:39:02	Apr-05 15:31	Apr-05 21:11	mesh
13980	pp.csh	silas	NPLQCD	0	128	00:00:20	04:30:08	Apr-05 14:24	Apr-05 18:54	mesh
13979	pp.csh	silas	NPLQCD	0	128	00:00:20	04:30:26	Apr-05 14:20	Apr-05 18:51	mesh
13978	pp.csh	silas	NPLQCD	0	128	00:00:22	04:29:45	Apr-05 14:19	Apr-05 18:49	mesh
13974	pp.csh	silas	NPLQCD	0	128	00:00:32	05:38:14	Apr-05 09:53	Apr-05 15:31	mesh
13976	pp.csh	silas	NPLQCD	0	128	00:00:20	04:27:54	Apr-05 09:56	Apr-05 14:24	mesh
13986	3c.csh	flemingg	HASTE	0	128	00:00:01	00:00:36	Apr-05 14:18	Apr-05 14:19	mesh
13977	pp.csh	silas	NPLQCD	0	128	00:00:20	04:21:54	Apr-05 09:57	Apr-05 14:19	mesh
13975	pp.csh	silas	NPLQCD	0	128	00:00:21	04:23:42	Apr-05 09:54	Apr-05 14:18	mesh
13972	test-128-4g	ychen	DEFAULT	0	128	02:18:32	00:05:00	Apr-05 09:13	Apr-05 09:18	mesh
13970	test-128-4g	ychen	DEFAULT	0	128	02:18:39	00:05:02	Apr-05 09:12	Apr-05 09:17	mesh
13971	test-128-4g	ychen	DEFAULT	0	128	02:18:33	00:05:05	Apr-05 09:12	Apr-05 09:17	mesh
13968	pp.csh	silas	NPLQCD	0	128	00:00:01	00:01:14	Apr-05 06:31	Apr-05 06:32	mesh
13966	pp.csh	silas	NPLQCD	0	128	00:00:02	00:01:14	Apr-05 06:24	Apr-05 06:25	mesh
13929	RsdCG_1p00e-9.c	flemingg	HASTE	0	128	717:38:53	05:46:21	Apr-04 23:51	Apr-05 05:38	mesh
13930	RsdCG_1p00e-9.c	flemingg	HASTE	0	128	546:53:03	04:18:35	Apr-05 00:00	Apr-05 04:19	mesh
13965	pp.csh	silas	NPLQCD	0	128	00:00:00	00:01:03	Apr-05 02:52	Apr-05 02:53	mesh
13964	pp.csh	silas	NPLQCD	0	128	00:00:01	00:01:13	Apr-05 02:51	Apr-05 02:52	mesh

Maintained by Webmaster

Done

SciDAC Project Status

- Most projects are on track for consumption of allocations
- By the end of the first year running period (June 2006) JLab will have delivered more node hours than was scheduled
  - 14 months instead of the 13 months required by the allocations (>12 due to 2m de-commissioning)
  - does NOT include additional running coming from the new 6n cluster (2+ months of friendly user running)