

CD FY08 Tactical Plan Status

**FY08 Tactical Plan Status Report for
SCF/QCD**

James N. Simone and Amitoj Singh

With

Bakul Banerjee, Bill Boroski, Don Holmgren, Jim Kowalkowski, Luciano Piccoli

2008-06-17

FY08 Tactical Plan for SCF/QCD

- **Relevant Strategic Plan(s):**
 - FY08 Strategic Plan for Lattice QCD [2295-v2]
 - Also (briefly) covers the Cosmological Computing Initiative and RF Cavity Simulations support for TD
- **Tactical Plan Leader: James N. Simone**
- **Organizational Unit home**
 - SCF/QCD
- **Tactical Plan Goals**
 - Support DOE Office of Science LQCD Computing Project (“LQCD.OMB300”)
 - Support DOE Office of Science Scientific Discovery Through Advanced Computing Lattice QCD Computing Project (“LQCD.SciDAC-2”)
 - Provide architectural and operational guidance for high performance computing initiatives such as Cosmology Computing.

Activities Summary: FTEs

Level 0 Activity:

Lattice QCD

% of FY Complete:

67%

Personnel Usage (FTEs)

Tactical Plan Level 1 Activity	Allocation		Actual YTD			Current FY08 Forecast
	FTE-yrs	FTE-mos	FTE-yrs (Ave/mo.)	FTE-mos	% Consumed YTD	
Lattice QCD						
Lattice QCD	2.50	30.0	2.29	18.36	61	91%
SciDAC Lattice QCD	2.65	31.8	2.73	21.84	67	100%
LQCD FNAL Project	2.45	29.4	2.08	16.60	56	100%
Total	7.60	91.2	7.10	56.80	62	
Lattice QCD [and Experimental Astrophysics]						
COSMO-Computing	[0.5 hire]	[6.00]	1.25	9.99	167	0.5 actual rate 4mos 200%
Total	[0.5]	[6.00]	1.25	9.99	167	

Luciano Piccoli (DAQ/Controls) is contributing 0.40 FTE (avg) to SciDAC Lattice QCD

Activities Summary: M&S (Internal Funding)

Level 0 Activity:

Lattice QCD

percent of FY Complete:

67%

Operating & Equipment M&S

CD Internal Funding

Tactical Plan Level 1 Activity	Operations M&S				Equipment M&S			
	FY Obligation Budget	YTD Obligations + RIPS	% Spent	Current FY08 Forecast	FY Obligation Budget	YTD Obligations + RIPS	% Spent	Current FY08 Forecast
Lattice QCD								
Facility ops and support	29078	4300	15	100%	0	0	---	
Scientific research (travel)	2475	0	0	90%	0	0	---	
Dept. infrastructure	1320	9860	750		0	0		
Total	32873	14160	43%		0	0		

Dept. Infrastructure obligations:

- Travel to SC08 and to LQCD-II CD0 review each for one member of QCD
- OpenClovis training ~5.5K

Activities Summary: M&S (External Funding)

Level 0 Activity: **Lattice QCD** % of FY Complete: **67%**

Operating & Equipment M&S

CD External Funding

Tactical Plan Level 1 Activity	Operations M&S				Equipment M&S			
	FY Obligation Budget	YTD Obligations + RIPS	% Spent	Current FY08 Forecast	FY Obligation Budget	YTD Obligations + RIPS	% Spent	Current FY08 Forecast
Lattice QCD								
LQCD FNAL Project	18000	2500	14		1489900	1452500	97	100%
SciDAC Lattice QCD	12000	2800	23		0	0	---	
Total	30000	5300	18%		1489900	1452500		

- M&S == travel
- EQ == “J/psi” cluster purchase + new servers and storage systems

Service Activity: LQCD/OMB300/Operations

- Operate the three LQCD.OMB300 clusters (QCD, Pion, Kaon) housed at LCC, Fermilab.
- Original Goals Related to this Activity
 - 1. Deliver 3.58 TFlops-yrs computing by Sept 30, 2008.
 - This is the aggregate TFlops-yrs for all three LQCD.OMB300 clusters.
- Adjusted Goals Related to this Activity
 - 1. Deliver 3.58 TFlops-yrs computing by Sept 30, 2008.
 - Project has delivered 2.58 TFlops-yrs till date and is on pace (2.387 TFlops-yrs, or 108% of pace).
- Key Metrics
 1. Uptime, utilization and successful job completion rate.
 2. <http://kaon2.fnal.gov/cluster/usage.html>
 3. <http://whcdf03.fnal.gov:10440/index.html>
- Milestones

As per external 2007 DOE USQCD progress review recommendations,

 - developed and deployed usqcd.org/fnal web-site on Feb 20, 2008.
 - implemented mechanics to track job failure rates
- Issues and Risks

Facility issues at LCC could result in the failure to meet LQCD/OMB300 milestone to deliver 3.58 TFlops-yrs.

Project Activity: LQCD/OMB300/Deployment

- Design and initiate procurement of 4.2 Tflop/s FY08 J/Psi cluster.
- Original Goals Related to this Activity
 - 1. Benchmark data for candidate J/Psi cluster.
 - Interact with computer vendors.
 - 2. Passing formal external review of preliminary design.
 - Yearly DOE USQCD progress review at BNL, May 13-14, 2008.
- Adjusted Goals Related to this Activity
 - 1. Benchmark data for candidate J/Psi cluster.
 - Lattice QCD performs and scales best on Intel Nehalem architecture but this will not be in GA till Jan 2009.
 - 2. Passing formal external review of preliminary design
 - Reviewers have passed the preliminary J/Psi cluster design.
- Key Milestones
 - After extensive benchmarking and attending the Intel HPC roundtable (May 2008), it is clear that the FY08 half of the J/Psi cluster will be an AMD Barcelona based system.
 - At the May external DOE review, the oral preliminary reviewers report endorsed the J/Psi design and procurement strategy.
- Issues and Risks
 - 1. If the GCC computer room C is not completed on time (current completion date of July 30, 2008), we might slip a top level OMB300 milestone.
 - 2. Unforeseen performance issues with a new cluster could delay promised release-to-production date.

Service Activity: LQCD OMB300 Project Management

- Provide project management for the DOE SC LQCD Computing Project
- **FY08 Objectives**
 1. Respond to any issues raised by OMB pass-back of the BY09 Exhibit 300.
Done – no issues raised by OMB.
 2. Respond to all data call requests from the DOE SC or OMB.
Done – all requests completed and submitted on time.
 3. Complete and submit quarterly reports to the Federal Project Manager according to timeline defined by DOE SC.
Done – FY08Q1, Q2, Q3 reports submitted on time; all perf. metrics met; have consistently received green scores.
 4. Organize the FY08 Annual DOE Progress Review.
Done – review held at BNL on May 13-14, 2008.
 5. Report on project progress and status to external review committee and receive favorable assessment.
In progress – Project status presented to committee; received favorable comments at review closeout; awaiting written report.
- **Key Metrics**
 1. Percent complete on FY08 deliverables
All performance deliverables on track; delivered Tflops-yrs through April: 7.00 (actual) vs 6.98 (goal).
 2. Budget performance on FY08 deliverables
Steady-state spend rate below linear baseline forecast through April (46% spent with 58% of year consumed). BNL QCDOC requiring significantly less support than anticipated; FNAL effort will ramp up to support FY08 cluster deployment.
 3. **Progress on specific scientific projects**
On track; specific performance details summarized in FY08-Q3 report.
- **Issues and Risks (specific to this activity, includes budget impact)**
 1. Procurement and deployment of FY08 cluster at FNAL is major project deliverable. RFI is complete; preparation of RFP and requisition in progress. Critical to keep this moving through system.

Project Activity: SciDAC Lattice QCD

DOE funded project to foster lattice QCD software development and investigations.

Original Goals

3. Automated workflow for LQCD
 - Implement preliminary design
4. Cluster reliability
 - Implement preliminary design
5. 64-bit (Opteron) optimizations
 - Linear algebra libs and validity checks
6. Native Infiniband QMP
 - Preliminary implementation
7. Common runtime environment
 - deployment

Adjusted Goals

2. Automated workflow for LQCD
 - Implement preliminary design for configuration generation
- Cluster reliability
 - Implement sensor scripts and database, preliminary design
4. 64-bit (Opteron) optimizations
 - 64-bit SSE for some linear algebra operations
5. Native Infiniband QMP
 - Deferred: QMP/MPI is already pretty good
6. Common runtime environment
 - Dropped as goal by USQCD software coord. committee

• Key Milestones

- Configuration generation workflow prototype implemented in ruby.
- Reliability sensor scripts implemented; 10mos of monitoring in relational dB
- <http://whcdf03.fnal.gov/exp/WorkflowProject.html>
- <http://whcdf03.fnal.gov/exp/ClusterReliabilityProject.html>
- Metrics: manpower delivered 2.7 FTE/mo (avg)

• Issues and Risks

1. Delivered FTE's at risk from budget constraints and manpower reassignments in CD.

Project Activity: Cosmo Computing

Development and operations of computing for cosmology.

- **Original Goals Related to this Activity**

1. Integrate hardware ordered in FY07 into the Computational Cosmology Cluster.
2. Release Computational Cosmology Cluster to production use.

- **Adjusted Goals Related to this Activity**

1. Integrate hardware ordered in FY07 into the Computational Cosmology Cluster.
2. Release Computational Cosmology Cluster to production use.
3. Implement LQCD-like metrics.
4. Assist in the purchase of components to expand the CC cluster.

- **Key Metrics**

1. Metrics are not currently tracked, and an important goal is to port the LQCD assessment framework to the CC cluster

- **Milestones**

- CC Cluster released to users late January, 2008

- **Issues and Risks**

1. On-going problems with AMD quad-core hardware
2. Job mix requires much more I/O than we were told – critical to deploy a parallel file system (Lustre?) in the coming months to handle the load

Activity: TD RF Cavity Simulations

This activity was not part of the 08 Tactical Plan

Provide MPI computing capability to TD engineers doing accelerator related RF cavity design.

- **Original Goals Related to this Activity**

1. Understand TD's FY08 computing requirements.
2. Provide TD engineers an MPI cluster so that they can meet their deadlines.

- **Adjusted Goals Related to this Activity**

1. Understand TD's FY08-09 requirements.
2. Expend minimum effort needed to keep MPI cluster running.

- **Key Metrics**

1. Usage since 10/23/07: 1020 jobs, 252K node-hours, 49% utilization
2. RFSIMS cluster usage <http://fnpc275.fnal.gov/>

- **Milestones**

- RFSIMS cluster commissioned 2007-10-23

- **Issues and Risks**

1. Activity not budgeted in FY08 and subject to available manpower.

Tactical Plan Issues and Risks

- Action Items
 - Need to plan for the transition from the Cosmos Computing Initiative to a future Cosmos Computing Facility.
 - Strategy for handling requests for high performance computing capability: one-off requests, initiatives, o(yr) projects, facilities.
- Issues and Risks
 - We are running very lean on effort
 - We must meet our OMB300 and SciDAC commitments
 - Deploy J/psi cluster: GCC-C ready on time
 - No extended LCC facility downtime
 - Sufficient FTEs


Tactical Plan Status Summary

- **Status Summary**
 - FTEs delivered mostly on track
 - QCD clusters have had few major issues the past few months.
 - Some base effort shifted over to Cosmos computing.
 - Bulk of LQCD project deployment effort will come after cancellation of furloughs.
 - SciDAC effort on track (temporary reassignment of L. Piccoli to QCD).
 - M&S spending:
 - Expect bump in base spending on deployment of the new J/psi cluster.
 - LQCD Project EQ Requisition for J/psi cluster purchase is in process.
 - Service and project activities
 - On target to reach goals and milestones for QCD project and SciDAC-II.

End of status report

Additional supporting slides follow...

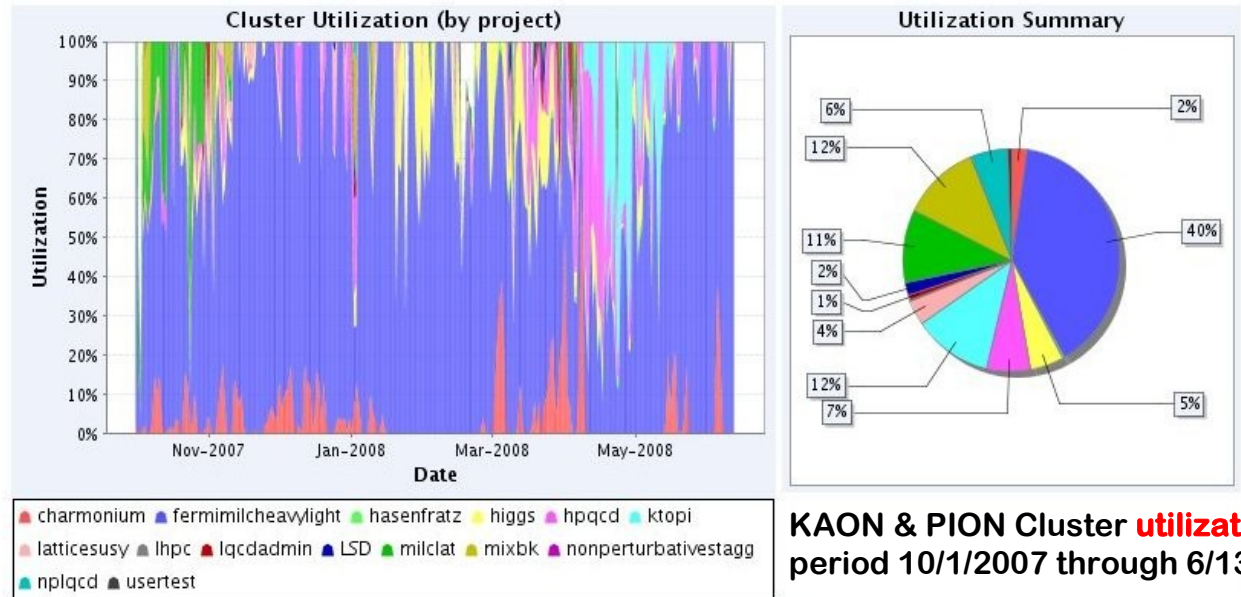
Service Activity: LQCD/OMB300/Operations



KAON yearly busy nodes plot (**uptime**)

Fermilab KAON & PION Cluster Utilization Chart

From: (yyyy-mm-dd) To: (yyyy-mm-dd) By:



KAON & PION Cluster **utilization** for period 10/1/2007 through 6/13/2008

Project Activity: SciDAC QCD Workflow: Goals

- Original
 - Complete evaluation of workflow systems ☺
 - Document simple workflow patterns used by LQCD ☹
 - Work with Askalon and Swift teams to create prototypes ☺
 - Select system to use ☹
 - Develop back and front end systems for workflow systems ☹
 - SuperComputing'07 poster ☺
 - Participate in SciDAC conference ☹
 - Have ConfGen workflow running by July ☺
- Adjusted
 - Detailed analysis of workflow use in ConfGen
 - Selected OpenWFU for prototyping
 - Delayed description of workflow patterns
 - Delayed work with Swift developers
 - Unable to participate in SciDAC conference aiming for posters at e-Science Conference and Lattice'08

Project Activity: SciDAC QCD Workflows: Status

- Workflow evaluation document
<http://whcdf03.fnal.gov/exp/attachments/WorkflowProject/WorkflowEvaluation.doc>
- Study of applicability of Askalon and Swift as workflow tools
 - Conclusion: current tools do not fulfill requirements:
<http://whcdf03.fnal.gov/exp/attachments/WorkflowProject/FunctionalRequirements.doc>
- Configuration Generation analysis document
- Development of ConfGen management system
 - Recording of data provenance and physics parameter
 - Uses OpenWFE (simple BPM engine), may be used with Swift or Askalon or any other workflow system
- Swift interactions are on hold (until we have a better understanding of our workflow needs)
- OpenWFE + Management System proof of principle is running on a small scale

Project Activity: SciDAC Cluster Reliability: Goals

- Original
 - Chose a messaging/control system ☹
 - Test Clovis on a cluster ☹
 - Reliably record monitoring data ☺
 - Recognize and correct a few problems automatically ☹
 - Instrument MPI to record job completion ☺
- Adjusted
 - Complete architecture and analysis/design document
 - Include essential failure and correction scenarios
 - Define possible worker node states
 - Investigate Clovis, SNMP and ActiveMQ
 - Clovis on hold
 - Implement data handling strategy for monitoring information
 - Enhance monitoring framework to carry out actions
 - Implement web based configuration for monitoring system
 - Review prototype infrastructure (Monitoring DB)

Project Activity: SciDAC Cluster Reliability: Status

- OpenClovis class
 - To be used as messaging/control layer
- Monitoring Database (10 months of data)
 - Syslog-ng used for monitoring, heartbeat and job status
 - Web access to job and health information
- Collected failure and corrective action scenarios
- Added instrumentation to MPI to track job completion

Project Activity: Other SciDAC activities

- 64-bit (Opteron) and other optimizations, SSE maintainance
 - Assisted Balint Joo at JLab with conversion of SSE inline assembly macros to more portable compiler intrinsics for use in the Chroma application
 - Isolated and found a work-around for a GCC compiler bug/feature that led to incorrect SSE SU(3) algebra results in production MILC code
 - Still need to package and/or refine Cray-supplied 64-bit SSE SU(3) optimizations for Chroma and the SciDAC libraries
- Native Infiniband QMP implementation
 - On indefinite hold since MVAPICH/OpenMPI continue to improve
- Common runtime environment
 - SciDAC software committee considerably relaxed requirements; this work is done except for maintenance
 - BNL/JLab/FNAL implemented web documentation following the common standard agreed-to by the committee

Project Activity: Computational Cosmology: Goals

- Original
 - Integrate hardware ordered in FY07 into the Computational Cosmology Cluster 😊
 - Release Computational Cosmology Cluster to production use 😊
- Adjusted
 - Integrate hardware ordered in FY07 into the Computational Cosmology Cluster 😊
 - Release Computational Cosmology Cluster to production use 😊
 - Solve hardware reliability issues 😞
 - Implement LQCD-like metrics
 - Implement a parallel file system
 - Assist in the purchase of components to expand the CC cluster
 - Integrate new hardware into the CC cluster

Project Activity: Computational Cosmology

- Release CC Cluster to production use
 - Troubles with the AMD quad core processor (virtual memory bug) caused the delivery of cluster nodes to be delayed from October 2007 until January 2008
 - Engineering sample processors were used
 - The full cluster was released to production in late January
 - About 25% of the cluster nodes power off under heavy computational load, believed to be caused by overtemperature
 - AMD acknowledges the issue and will replace all processors with the new stepping
 - The engineering sample processors were 120 Watt parts; the replacement CPUs will be 90 Watt parts
- Implement LQCD-like metrics
 - We want to report uptime, utilization, job failure rates, and so forth
 - Procedures and scripts from LQCD can be used, but available manpower to date has been devoted to reliability issues, and recently, I/O issues

Project Activity: Computational Cosmology

- Implement parallel file system
 - We designed the cluster to accommodate the FNAL MPI codes (“ART”), which are computationally-intensive and not I/O intensive
 - Many of the external users run other code bases which are very I/O intensive
 - The burden on the existing NFS storage infrastructure often leads to unacceptable job performance
 - We believe the solution is to move storage to a set of dedicated nodes, and to use a parallel filesystem like Lustre or PVFS2 to simplify access
 - Will likely release a Lustre prototype to production in the next month
- Cluster Expansion
 - PPD has funds to expand the cluster by 60 dual-socket quad-core Opteron nodes + another 42 Tbyte SATAbeast
 - U.Chicago/KICP would add an additional SATAbeast
 - Expect new hardware by end of FY08