

# SC09 Post-event assessment: Meetings

Report prepared by: Anne Heavey  
Report reviewed by: David Ritchie and John Urish  
December 15, 2009

I.	Executive Summary:.....	2
II.	Did the meetings at SC09 meet objectives? .....	2
	A. Build and nurture relationships.....	2
	B. Foster respect and highlight CD's role at Fermilab .....	2
	C. Encourage collaborative relationships.....	3
III.	Goal-outcome-benefit summaries of meetings.....	3
	A. Security .....	3
	1. "Cyber-security Science DOE Grass Roots" meeting .....	3
	2. Privacy-preserving sharing of network data.....	4
	B. Software.....	4
	1. ANI/Magellan kickoff .....	4
	2. JDEM Demonstrator .....	4
	3. ROSE compiler infrastructure .....	5
	4. Tech-X SBIR for message-passing standard DDS .....	5
	5. Kepler workflow .....	5
	6. TAU Performance Tools .....	6
	C. Hardware.....	6
	1. Mellanox Infiniband hardware.....	6
	2. AMD-containing hardware for USQCD cluster procurement.....	7
	3. Dell HPC-optimized cluster hardware for USQCD cluster procurement .	7
	4. Sun: Lustre and robot arm discussions .....	8
	5. DDN: SA10000 file system and S2A6620 storage system .....	8
	6. Intel's Larabee chip .....	9
	D. Open Science Grid.....	9
	1. Cloudera .....	9
	2. SCEC (Southern California Earthquake Center.....	9
	3. Tech-X: SBIRs of benefit to OSG and USCMS .....	9
	4. NeesCom.....	10
	E. SC09 planners from national labs .....	10
	F. Results from miscellaneous informal meetings .....	11
IV.	Technical session and tutorial take-aways.....	11
	1. Storage and Cloud Challenges.....	11
	2. Tutorial: Programming models for computers with CPU and GPU cores	11
	3. Workshop on Workflows in support of large-scale science.....	12
V.	CMS Centre .....	12

## **I. Executive Summary:**

At SC09, Fermilab staff welcomed scheduled meeting attendees and impromptu visitors into the booth, attended meetings in other booths and meeting areas, and participated in the technical sessions and tutorials. We report on roughly 20 of the meetings, identified as the most significant, some of which took place at other booths.

Fermilab staff learned about the status of several software packages and projects that are potentially useful to Fermilab researchers and/or opportunities for collaboration. These are largely noncommercial products from other research institutions. The staff also collected information about commercial hardware products that will inform future purchasing decisions.

Fermilab staff promoted the lab's scientific mission through informal conversations in the booth.

We conclude that the FNAL staff met the stated meeting-related objectives for SC09.

## **II. Did the meetings at SC09 meet objectives?**

### **A. Build and nurture relationships**

*Objective: Hold meetings with funding agency representatives, vendors, computing/scientific people from universities and other national labs, and interested public in order to build and nurture relationships*

FNAL attendees reported on a total of over 25 meetings, split fairly evenly between current or potential collaborators from other institutions and vendors. There were fewer, but still some, meetings with funding agency representatives.

**Answer:** YES, this objective was met.

### **B. Foster respect and highlight CD's role at Fermilab**

*Objective: Foster an interest, respect and appreciation for Fermilab's computing and scientific work among same groups as above, and highlight the role the Computing Division plays in Fermilab's mission.*

Visitors spoke informally with available booth staff, with CMS computing and scientific staff at the CMS centre, with educators from QuarkNet and with high energy physicists over the HD video link during "Ask a scientist" hours.

As the Computing Division's subject matter experts discussed their projects' software and hardware needs in meetings with SC09 attendees, they

communicated information at varying levels of detail about Fermilab’s leading-edge computing R&D, and how it serves the scientific experiments.

**Answer:** YES, this objective was met.

**C. Encourage collaborative relationships**

*Objective: Encourage collaborative relationships between Fermilab and other research institutions.*

The goal for several of the meetings was to form or strengthen software development collaboration efforts and to encourage new grid application teams to join the OSG. At the close of some of these meetings, participants identified steps for moving ahead collaboratively.

**Answer:** YES, this objective was met.

**III. Goal-outcome-benefit summaries of meetings**

The SC09 team asked attendees to report on significant meetings they held and/or attended; i.e. meetings that produced or led towards positive outcomes and/or benefits. Therefore, this section includes information for many but not all the meetings in which FNAL staff participated at SC09.

Fermilab staff are initially identified by their group within CD; subsequently, only as FNAL.

**A. Security**

**1. “Cyber-security Science DOE Grass Roots” meeting**

Attendees	G. Ghinita (Computer security team), K. Chadwick (FermiGrid), R. Pordes (Comm and Outreach, OSG) attended larger meeting (not at FNAL booth)
Goals	<ul style="list-style-type: none"> <li>• DOE high priority research topic: find a unified model for cyber-security to replace the current “detect-and-patch” approach. This will involve large-scale modeling of the Internet (up to 10<sup>5</sup> nodes).</li> <li>• DOE secondary priority: strike balance between protection and the impact that security measures have on end users’ productivity</li> <li>• DOE secondary priority: find metrics to quantify security</li> </ul>
Outcomes	Meeting attendees will hold bi-weekly telecons to discuss cyber-sec topics.
Benefits	Potentially allow FNAL to work with DOE to establish and implement more effective and user-friendly cyber-security solutions.

## 2. Privacy-preserving sharing of network data

Background: Deb Agarwal of LBNL is leader of the Data Intensive systems group and recipient of DOE cyber-security grant – with UCD – to devise effective intrusion detection systems that automatically detect cyber attacks based on pattern of network communication and analysis of application logs. Agarwal’s project seeks to devise techniques that allow intrusion detection on top of anonymized data.

Attendees	G. Ghinita (FNAL), Deb Agarwal (LBNL)
Goals	To discuss and identify possible strategies and privacy paradigms for privacy-preserving data sharing. Also to look at impact that anonymization has on detection accuracy and runtime performance.
Outcomes	Agreement to consider as first step k-anonymization algorithms (these implement certain syntactical constraints on output) and permutation-based approaches. First steps (to be taken by LBNL) will try to apply some existing algorithms to data and check the amount of distortion.
Benefits	Potential for fruitful FNAL-LBNL collaboration within funded scope.

## B. Software

### 1. ANI/Magellan kickoff

Advanced Network Initiative and Magellan cloud computing

Attendees	M. Crawford (Data Movement and Storage), P. Demar (WAN and Network Research), R. Pordes, and others (FNAL), Thomas Ndousse, Vince Dattoria and Susan Turnbull (DOE/ASCR), people from ANL, ORNL and several universities
Goals	<ul style="list-style-type: none"> <li>• To push forward our case for getting the 100Gb/s ANI to FNAL (goal met to 50% level)</li> <li>• To discover funding opportunities (met 75%)</li> <li>• To meet the principals and learn the plans (met 90%)</li> </ul>
Outcomes	Ndousse said that he wants certain work to be done by us, although he did not steer funding to go with it. FNAL aligned plans for possible storage research facility project here (which would be funded if approved) with Magellan activities.
Benefits	Potential funding for storage research facility

### 2. JDEM Demonstrator

Attendees	J. Kowalkowski (Computing Enabling Technologies) and others (FNAL), Deb Agarwal from LBL, who controls funding for JDEM. ( <a href="http://jdem.gsfc.nasa.gov/">http://jdem.gsfc.nasa.gov/</a> )
Goals	Discuss the JDEM Demonstrator system and the work goals for the year.
Outcomes	Produced a few revisions to put in project definition over the next month.

Benefits	We were able to begin to come to an agreement on what Fermilab would be doing for JDEM over the next year.
----------	--

### 3. ROSE compiler infrastructure

Background: ROSE is an open source compiler infrastructure to build source-to-source program transformation and analysis tools for large-scale Fortran 77/95/2003, C, C++, OpenMP, and UPC applications. ([www.rosecompiler.org](http://www.rosecompiler.org))

Attendees	M. Paterno (Computing Enabling Technologies) and Dan Quinlan (LLNL)
Goals	<ul style="list-style-type: none"> <li>Determine whether the compiler technology of ROSE is of interest for group's goals, including the parallelization of existing code, quality analysis of code, and the development of scientific data processing frameworks.</li> <li>Determine specific tools provided that are of interest.</li> </ul>
Outcomes	It appears promising. Follow-up with the ROSE development team will be necessary if our initial investigations bear out the conclusions from our meeting at SC09.
Benefits	(not explicitly stated)

### 4. Tech-X SBIR for message-passing standard DDS

Background: Tech-X is providing FNAL with DDS experience (a message passing standard) and code to exercise DDS in a way that is interesting to us. DDS is the system we are evaluating for various uses around the lab, including DAQ work, and workflow reliability and monitoring.

Attendees	J. Kowalkowski, M. Paterno et al (FNAL), with Sveta Shasharina (Tech-X)
Goals	<ul style="list-style-type: none"> <li>To discover the status of the project and walk through the code that was provided.</li> <li>Inform Tech-X of FNAL's expectations.</li> </ul>
Outcomes	A follow-up meeting is scheduled to discuss implementation details.
Benefits	Able to discuss face-to-face the progress and sort out a few rough edges in our collaboration efforts. (Benefits of DDS not explicitly stated.)

### 5. Kepler workflow

Attendees	J. Kowalkowski, M. Paterno, et al (FNAL), with Ilkay Altintas (SDSC)
Goals	Discuss progress of Kepler package with regards to issues of interest to group, namely functional workflow specifications, provenance interface, recovery through rules or simple logical expressions, and a framework that appears to be extensible (by FNAL). Discover if there are opportunities to work together.
Outcomes	The goal was met. Jim et al will have a meeting in December with

	Altintas to discuss working together probably on the messaging and reliability aspects of workflow. Jim's group will look for a simple application to give Kepler another try, including use of their provenance interface and distributed computing components.
Benefits	To be able to walk through many aspects that we are interested in about Kepler very rapidly. (Benefits of Kepler not explicitly stated.)

## 6. TAU Performance Tools

Attendees	J. Kowalkowski, M. Paterno et al (FNAL), with Sameer Shende of University of Oregon
Goals	To meet Sameer, with whom they expect to work to tailor the TAU tools in the future. Also, to run TAU (in Sameer's presence) on a body of code on our machine to learn how to operate particular features of the tools, and to verify that we installed TAU properly and can use it to make a good set of performance measurements.
Outcomes	<p>Jim et al will likely need some further training/consulting on the use of TAU as they figure out what more they want out of it. They are equipped now to make use of it to some level.</p> <p>TAU can be used as is for some of the work the group does to improve the utilization of their computing resource. After they use it for a while, they will probably want to add things to it or use the data it generates behind-the-scenes in different ways.</p>
Benefits	The TAU tools report many application execution performance numbers that Jim Kowalkowski and his group are interested in. These tools work well.

## C. Hardware

### 1. Mellanox Infiniband hardware

Attendees	A. Singh, D. Holmgren (High performance parallel computing facilities), representatives from Mellanox (Gene Crossley, Brandon Hathaway, Marc Sultzbaugh) and JLab (Chip Watson)
Goals	<p>To understand relevant current and upcoming Mellanox Infiniband hardware that may be used in the upcoming USQCD cluster procurement that will be performed by Fermilab and housed in GCC-C.</p> <p>To learn specific details and dates of availability for new products that will affect performance on lattice QCD codes.</p>
Outcome	The goals were met. We learned that MPI collectives will be optimized in this hardware, and that it will be available at the time of our purchase. We discussed at length the use of Mellanox hybrid switches that can bridge Infiniband, 10 gigE, and fibre channel. Mellanox committed to making cluster resources available to

	Fermilab for benchmarking prior to our RFP.
Benefits	The information learned about the upcoming “ConnectX2” Infiniband silicon is very important for the design of the new cluster. This information is also relevant to other storage needs at Fermilab.

## **2. AMD-containing hardware for USQCD cluster procurement**

Attendees	D. Holmgren (FNAL), representatives from AMD (Ron Schooler, Boris Cownie, Annie Flaig, Chris Cowger), Koi Computers (Fanny Ho), and JLab (Chip Watson).
Goals	To understand, via a non-disclosure presentation, the relevant AMD and AMD-containing hardware (processors, chipsets, motherboards) that will be available at the time of FNAL’s upcoming USQCD cluster procurement to be housed in GCC-C. To learn specific details about new processors (memory channels, floating point execution units) that will affect performance on lattice QCD codes.
Outcomes	Goal was largely met: Because the NDA between Fermilab and AMD had not yet been executed at the time of the meeting (it was executed two days later), some specific details about performance projections and chip speeds were withheld. Learned a great deal about the upcoming Magny-Cours processor family that will help inform our upcoming RFI and RFP. I agreed to provide a lattice QCD benchmark suite in December to the technical contact (Boris Cownie) that will give both Fermilab and USQCD valuable information about the performance of this new processor family; AMD agreed to run this benchmark suite and provide results. AMD also agreed to work with Koi to provide early samples to Fermilab for hands-on testing.
Benefits	Information will help inform our upcoming RFI and RFP for cluster procurement

## **3. Dell HPC-optimized cluster hardware for USQCD cluster procurement**

Attendees	J. Simone (HPPC facilities), A. Singh and D. Holmgren (FNAL) met with representatives from Dell (Claudine Conway, Mickey Henry, Michael Riley, Garima Kochhar, Mike Wilmington).
Goals	To understand relevant current and upcoming Dell AMD- and Intel-based cluster hardware that may be used in the upcoming USQCD cluster procurement.
Outcomes	The goals were met. We had the opportunity to do a hands-on

	inspection of Dell HPC-optimized cluster hardware. Dell agreed to send a pre-production server (dual motherboard, dual socket) to Fermilab for evaluation; this server will be available to all interested parties at the lab.
Benefits	Information will help inform our upcoming RFI and RFP for cluster procurement

#### 4. Sun: Lustre and robot arm discussions

Attendees	G. Oleynik (Data movement and storage), M. Crawford (FNAL) met with Miriam Wagner (Sun), delegation leader; Lustre developers Hua Huang and Nathaniel Rutman
Goals	Get attention onto our renewed robot arm failures -- 50% met Get confirmation of their work on LTO-4 drive problems -- 100% met Make headway with Lustre support & HSM integration -- 75% met
Outcomes	Sun "will look into" robot arms Sun is standing by LTO-4 maintenance/replacement commitment, despite IBM's (the mfr.) claim of no problems found. General availability of the Lustre HSM interface is not expected for some time, though it is sufficiently developed for us to be an "alpha" (beta?) user. No Lustre support for free with tape libraries. Early access to new source code can be had with a maintenance contract.
Benefits	A new HSM feature of Luster (to allow Lustre file systems to be a component of a taped back tiered storage system). This could provide the framework to integrate Lustre with enstore.

#### 5. DDN: SA10000 file system and S2A6620 storage system

Attendees	G. Oleynik, M. Crawford (FNAL) and Rosen, McKenna, Busch and others from DDN
Goals	To learn more about their storage system features: SATASure data integrity checking/correction on reads (raid 6 parity checked on all block reads) which goes beyond normal raid scrubbing, high density, and the integration of Lustre into their controllers.
Outcomes	DDN has promised an evaluation S2A6620 to CMS since last November and CMS still hasn't seen one. We will follow up with DDN shortly.
Benefits	Information will help us in long-term planning of disk storage, in particular, DDN's integration of Lustre file system in their controllers.

## 6. Intel's Larabee chip

Background: Intel's Larabee chip has a heterogeneous architecture that includes specialized processors in addition to general purpose CPU cores.

Attendees	A. Singh and J. Simone (FNAL) met with Intel personnel
Goals	To discuss their Larabee product as it relates to programming heterogeneous processors
Outcomes	We learned about the "Ct" language that Intel is developing to ease the task of programming for heterogeneous processors.
Benefits	Discussions helped clarify outstanding questions as Fermilab works through the process of signing an NDA with Intel concerning Larabee and related architectures.

## D. Open Science Grid

### 1. Cloudera

Attendees	R. Pordes (FNAL), Jeff Hammerbacher (Cloudera Company)
Goals	To discuss exchanging support for testing of new software on OSG.
Outcomes	Agreement in progress between Nebraska, OSG and Cloudera
Benefits	For OSG: Support for Cloudera at no \$ cost For FNAL: Increased profile with Hadoop development.support organizations (The Apache Hadoop project develops open-source software for reliable, scalable, distributed computing. -- <a href="http://hadoop.apache.org">hadoop.apache.org</a> )

### 2. SCEC (Southern California Earthquake Center

SCED runs across TeraGrid and OSG, XD-OSG collaboration

Attendees	R. Pordes, (FNAL), Dan Katz, Richard Moore, Ewa Deelman – UofChicago, SDSC, ISI
Goals	to facilitate collaboration on SCEC, earthquake engineering application
Outcomes	parties agreed that UofC and ISI groups will work together rather than continue separately
Benefits	OSG: Reduced risk that the project (assumed to mean the cross-grid SCEC application) will fail FNAL: Reduced risk that OSG will fail

### 3. Tech-X: SBIRs of benefit to OSG and USCMS

Attendees	R. Pordes (FNAL), Mark Green (Tech-X)
Goals	Discuss SBIRs of mutual benefit to OSG and US CMS
Outcomes	Paul Sheldon and Ruth wrote letters of support to DOE about the SBIR for Virtual Machines and Cloud support.
Benefits	CMS: Increased understanding of how to run applications using VMs

	and Clouds Fermilab (CMS): No cost increase (relative to what, not specified)
--	--

#### 4. NeesCom

Attendees	R. Pordes, Rudi Eigerman (NeesCom)
Goals	discuss next step in NEES contributions
Outcomes	Increased likelihood of collaboration with NEES; delivery of responsibilities of NEES Community collaboration (to?) Sub-committee lead (Pordes)
Benefits	(not stated)

R. Pordes listed four additional OSG meetings without providing details.

#### E. SC09 planners from national labs

Attendees	A. Heavey, D. Ritchie and J. Urish (Comm & Outreach), Daniel P? (ORNL), Cheryl Drugan, Beth Cerny and Chel Lancaster (ANL), Melinda Lee (SLAC)
Goals	To share information and ideas about how different labs approach this yearly event, how they set and achieve objectives for it, and how they measure their success.
Outcomes	<p>We exchanged information on amount of people and time each lab spends on planning the booth, how many people each lab sends, how each advertises ahead of time, manages the booth meeting space during the convention, and how planned presentations go.</p> <p>We discovered that we (FNAL) had fewer complaints than others about staff cluttering up the booth with personal items and usurping meeting areas. Labs that had back-to-back presentations didn't find them terribly well-attended, and our "meeting area" plan worked comparatively well. For pre-show advertising next year, ANL agreed to share its email list (we might want to edit it to hit targeted communities), and we may follow ORNL's idea of sending out postcards (they send out 1600 two weeks ahead).</p> <p>We did not discuss enough about booth objectives and metrics for success.</p> <p>We planned a phone meeting in the March 2010 time frame as follow-up.</p>
Benefits	Potentially enable more efficient and effective SC planning and execution

## F. Results from miscellaneous informal meetings

These notes come from informal conversations with attendees. The notes are representative, not comprehensive.

- (A. Tiradani) Talk about OSG interoperability with TeraGrid and with EU grids, about gLite and Condor. He got contact information to learn about how to apply these to OSG.
- (A. Tiradani) A meeting with John Paul Navarro (TeraGrid) was productive: “maybe what they do, we could do in OSG”.
- (K. Chadwick) Potential to bring in KOI as vendor for FNAL as a “white box” vendor of “headless” servers (what we call worker nodes).
- (K. Chadwick) Discussion with RedHat about compiling issues
- (K. Chadwick) Infiniband — alternative (to ethernet) interconnect for computers. Price needs to come down — will happen as cost of switches goes down, silicon production improves, etc.

## IV. Technical session and tutorial take-aways

Only two attendees reported on a total of three sessions. We did not obtain a count of how many and which sessions the FNAL group attended collectively.

### 1. Storage and Cloud Challenges

Attendees	G. Oleynik attended presentation by Henry Newman of Instrumental, Inc.
Outcomes	Speaker reviewed the implications of the expected bit error rates of various storage devices and data transport hardware with Petascale+ volumes of data in terms of the volume of data one expects to be corrupted.
Benefits	It was useful to see all of this information collected in one place and the transparencies are good reference material.

### 2. Tutorial: Programming models for computers with CPU and GPU cores

Attendees	J. Simone attended two tutorials "S04: High Performance Computing with CUDA" (all day) <a href="http://scyourway.nacse.org/conference/view/tut151">[http://scyourway.nacse.org/conference/view/tut151]</a> and "M13: OpenCL: A Standard Platform for Programming Heterogeneous Parallel Computers" (half day) <a href="http://scyourway.nacse.org/conference/view/tut149">[http://scyourway.nacse.org/conference/view/tut149]</a> .
Goals	To gain more knowledge about programming GPU and heterogeneous systems.
Outcomes	The common theme of both tutorials was “programming models for a computer having both general purpose CPU cores as well as specialized processors such as Graphics Processing Units (GPUs)”.

	<p>The CUDA tutorial described the language Nvidia has developed for programming GPU systems such as their Tesla product as well as the millions of Nvidia GPUs already deployed in desktop/laptop graphics cards (e.g. all recent Macs). The second tutorial on OpenCL described a similar language which is somewhat more vendor neutral. Videos and recordings of the tutorial on Cuda programming are available online:  <a href="http://news.nvidia.com:8080/t/114573/13091975/5217/0/">http://news.nvidia.com:8080/t/114573/13091975/5217/0/</a></p>
Benefits	<p>These two tutorials were particularly timely since the HPC group has already deployed Tesla GPU systems in both a production system for Lattice QCD as well as a software development system for CD use.</p>

### 3. Workshop on Workflows in support of large-scale science

Attendees	<p>J. Simone attended the morning session of the "4th Workshop on Workflows in Support of Large-Scale Science (WORKS09)" <a href="http://scyourway.nacse.org/conference/view/wksp110">[http://scyourway.nacse.org/conference/view/wksp110]</a>, a workshop on scientific workflow systems. The conference details are here: <a href="http://www.isi.edu/works09/">[http://www.isi.edu/works09/]</a>.</p>
Goals	
Outcomes	<p>Of particular note are the presentations:</p> <ul style="list-style-type: none"> <li>• Johan Montagnat, Benjamin Isnard, Tristan Glatard, Ketan Maheshwari and Mireille Blay Fornarino, "A data-driven workflow language for grids based on array programming principles"</li> <li>• Daniel Zinn, Shawn Bowers, Timothy Mc Phillips and Bertram Ludäescher, "Scientific Workflow Design with Data Assembly Lines"</li> <li>• Matthew J. Sottile, Geoffrey C. Hulette and Allen D. Malony, "Workflow representation and runtime based on lazy functional streams"</li> </ul>
Benefits	<p>The workshop session "Workflow Representation" was of particular interest to Fermilab developers working on a workflow system for Lattice QCD. This led to separate meetings (at the Fermilab booth) with Ewa Deelman of the Pegasus project and Ilkay Altintas of the Kepler project (also cited in this document).</p>

## V. CMS Centre

The CMS centre attracted a lot of attention. No formal meetings were held, but many impromptu discussions took place. P. Gartung estimates about 15 visitors/day. He spoke at length with at least five visitors/day. He was one of three people manning the booth, off and on.

Staffers informed visitors about the LHC, CMS and Fermilab's role in these projects. They answered questions on how CMS does computing and analysis – how data is moved and how jobs find the data -- and found that people are surprised that MPI isn't used. They also answered questions about "the black hole" and about competition from the new accelerator.