Scientific Computing

As a means for generating simulations, scientific computing complements theory and experiment as a way to obtain scientific knowledge. But computing is more than the third leg of the discovery stool. Scientific computing also supports and enables the other two through data collection, reconstruction and analytics. It is and has always being an essential part of the Fermilab physics program.

Scientific computing activities at Fermilab are carried out across the two Divisions in the Office of the CIO, who work closely with the other laboratory organizations to deliver world-class scientific computing services, operations and software engineering support to Fermilab-based experiments and projects, CMS and the high-energy physics community at large. We have continued to build out activities as planned for supporting software and computing needs for the neutrino experiments. Scientific Computing includes R&D activities required to maintain or advance capabilities necessary for the success of our future physics program, as Fermilab moves forward with the P5 plan.

As we execute the P5 plan, our scientific computing R&D program has to address some very important challenges. As computing architectures evolve, we must change the paradigms for how we construct our algorithms, write our codes and organize our analysis workflows. We also need to continue making effective use of High Performance Computing (HPC) for large computational problems. In addition, while new technologies, such as cloud computing, provide new, potentially cost effective options, for deploying computing resources in the future, they require us to develop new services for on-demand reliable resource allocation. Finally, we must continue to provide high level of support to the ongoing Fermilab experiments and projects that have different timelines and priorities.

In order to meet these challenges and continue to serve the needs of our user community we have to leverage effort and resources: we develop and maintain core expertise, toolkits and services that are used to build and provide operational support to applications with a variety of requirements. We work in partnership with individual programs and projects to create (plan, architect, develop, deploy, operate) the applications that fit their needs. Also, we seek and utilize outside partnerships and collaboration to pursue necessary R&D on new techniques and technologies.
To help maximize leveraging of expertise and best-of-class tools from partnerships with individual projects we have recently reorganized scientific computing in the Scientific Computing Division (SCD) and aligned our activities across three crosscut areas.

- One is development, integration and research, in which we develop or package the software products that run on our facilities and develop the necessary architecture.
- The second is facilities, where we operate the services that run these products.
- The third is science operations and workflows, through which we tailor applications of the facility services to our experiments and projects and assist with operations.

The scientific computing contributions to the five science-based strategic themes of Fermilab's program are detailed in the main document for this review, "Fermilab Status and Strategic Plan". Here we summarize the drivers for scientific computing and the main elements of our strategic and tactical plans for this area, and discuss in detail the scientific computing R&D efforts.

Scientific Computing Drivers

Our activities are focused in the following five areas, with specific objectives and drivers:

Support the **CMS science program,** by
- hosting and operating the top performing CMS Tier-1 facility and the LPC computing facility and providing the workflow management and monitoring tools necessary for the operations,
- providing storage and data management services for raw, reconstructed and simulation data,
- developing and supporting the analysis framework.

Support the diverse FNAL **neutrino and muon experimental program,** in all aspects of its computing needs, by providing
- facility with performance and capabilities at the same high level as the Tier-1,
- common tools, services, and operations for all aspects of data collection, storage, and physics analysis,
- detector and physics simulation expertise.

Support selected **cosmic frontier experiments** as per P5 and FNAL priorities, with emphasis on workflow, workflow management, data storage, DAQ, and R&D in areas where our competencies align with programmatic needs (such as work on Magnetic Kinetic Inductance Detectors --MKIDS)

Enhance understanding for and optimize performance of **current and future FNAL accelerators,** by leveraging our HPC modeling competencies. Applications focus on Proton Improvement Plan (PIP) and PIP-II, with emphasis on large scale beam-dynamics simulations that include collective effects, to provide guidance in the Booster, Main Injector and Recycler operations, and any future design efforts. Our state-of-the-art modeling capabilities allow quantitative beam loss characterization and study of beam instabilities, which are essential for the success of the PIP and PIP-II programs.

R&D for new tools and services: evolution of computing architectures, technologies and the transition of the Fermilab scientific program to that of a fully international Laboratory call for major software re-engineering to maintain and extend capabilities as we move in the future.

Our objective is to develop the competencies and technology that allow us to upgrade our facilities and applications to the capacity and capabilities necessary for the future CMS, IF and broader scientific programs. Our program focuses on emerging new technologies such as multicore/many-core, co-processor technologies, the challenge of reduced memory/core footprint of new architectures, and the emergence of the cloud concept as a scientific computing tool. We target selected high impact/relevance areas such as analysis frameworks, Geant4, data acquisition and online controls, and accelerator modeling.

Tactical approach

Supporting the current program and preparing for the future requires significant investment in scientific computing. Leveraging resources and expertise is essential for continuing to provide state of the art reliable facilities and computing services to the community. For example, we leverage expertise and tools from our very successful CMS program and integrate tools and services we have developed for the IF program, in a coherent computing "ecosystem" that we make available to all of our users. We then work with individual experiments and projects to develop and deploy their specific applications, using components from the "ecosystem".

Although Scientific Computing has the responsibility and the lead in this area, we work closely with all lab organizations to support the entirety of Fermilab's program. We also have strong ties with other institutions and participate or lead many collaborative efforts on computing, especially on R&D on new architectures in order to transform major HEP toolkits. This is especially important, as we aim to maintain competencies and infrastructure capabilities to help the lab position itself to play a significant role in next-generation experiments, upgrades, and develop future programs.

Scientific Computing services are an integral part to the success of the scientific program as specified by P5. Scientific Computing at Fermilab works with each experiment to discuss an end to end computing model to best fit experiment's computing needs and exploit the available resources. Discussions and planning start as early as possible and can be used to help with preparations for planning documents etc. We are full partners with the scientific community with each scientist as a full member on at least one experiment, providing continuous feedback both on planning and operations. Additionally, each experiment has an assigned liaison to ensure effective communication between scientific computing and the experiment or scientific project and to ensure the computing needs are met. Liaisons help to engage the experiment with computing expertise that can help at all stages from architecture through design and deployment and into operations. 30+ staff in computing, or over 15%, are performing liaison duties.

The directorate provides oversight and governance to laboratory IT systems through the IT Executive Council. One of the three underlying branches oversees scientific computing: the Scientific Computing Portfolio Management Team (SCPMT). Requests for scientific computing resources are made through the formal annual review process by the SCPMT. SCD helps the experiments prepare for the SCPMT review and runs a smaller working group that meets regularly with experiments throughout the year to address any allocation changes and requests for additional services. These resource levels and services are agreed to and documented in the experiment's Computing Sector (CS) Technical Scope of Work (TSW).

In order improve the services that we are delivering to the scientific community, scientific computing is committed to the overall IT strategy of ITIL best practices. The Scientific Collaboration Tools service was certified in ISO20K this year and is the second SCD service to be onboarded. This coming year we will on-board all services to incident and change management.

Organization and Activities

Fermilab recently took the next step in better aligning Scientific Computing along functional areas, to eliminate any leftover "stovepiping" for specific applications or experiments and increase resource utilization effectiveness, to better serve the entirety of our scientific program. Also, in addition to the usual line management oversight, we have introduced coordinators for each one of the three major areas to provide tighter coordination across activities, defined by overall service deliverables, and increase efficiency and quality of service. The activities and portfolio of these high-level, cross-cutting, activity areas across Scientific Computing are discussed below.

Development, Integration and Research

The purpose of research, development and integration projects are to enhance or improve aspects of computing for experiments to positively impact their ability to arrive at physics results. Scientific Computing includes many activities along these lines, encompassing the full stack of software for HEP.

- <u>Frameworks and Tools for physics analysis and reconstruction</u>
  Physicists want to write science code and not worry about the low level computing infrastructure, such as file i/o, dynamic library loading, data delivery and handling, module scheduling, etc. A *framework* provides these services, allowing physicists to write their science code in modules. We have

led the development of the framework for CMS, called *CMSSSW*, for many years. We also develop a framework based on CMSSW, called *Art*, but tailored to the unique requirements of neutrino and precision muon experiments. Art is written generically to be used by many experiments and is in production at NOvA, Mu2e, Muon g-2, and many LArTPC neutrino experiments such as MicroBoone, LBNF, LarIAT, etc. Management of such widely used code so deep in an experiment's software stack can be a challenge, but the Art stakeholders organization is very successful. Decisions on the evolution of Art's feature set are arrived at by consensus among the stakeholder experiments. While the CMSSW and Art frameworks serve different communities, the fact that Art started from CMSSW suggests that they have large functional overlap. To optimize efficiency in framework support and development, we are working to identify and develop common functions and infrastructure so that the evolution of these frameworks occurs together, guided by common principles.

We have also been developing a higher level layer on top of Art, called *LarSoft*, that directly supports the development of analysis and reconstruction algorithms for LArTPC neutrino experiments. And the SCD is also writing a framework for cosmology experiments called *Cosmosis* used by DES and is the basis for developing a framework and workflow engine for LSST.

Scientific Computing activities include research into how these frameworks and other science codes evolve in this era of multiprocessing and accelerator co-processors (GPGPU's, Intel PHI, etc). The lab recently awarded the SCD an LDRD to explore use of a pattern-recognition automata processor, a truly new, cutting edge, and novel architecture. Furthermore, integrating these frameworks with "Big Data" methodology and tools used in industry is an active area of study.

Nearly every physics analysis in recent years involves the ROOT data analysis application. Fermilab Scientific Computing contributes extensively to its development, especially in the area of i/o and ensuring that ROOT continues to meet the needs of CMS and Fermilab's neutrino and muon experiments.

Finally, SCD physicists are deeply involved in physics analyses at experiments and contribute to development of their experiment's science code and physics results. These physicists also aid in developing and deploying scientific computing tools and services. Many of the scientists, especially those on neutrino or muon experiments, are a CS liaison to their experiment.

- <u>Real time Software and Engineering</u>
  A modern DAQ system that is as experiment generic as possible is under development by SCD physicists and engineers. Called *ArtDAQ*, it leverages the

Art low level infrastructure as a framework for data capture and event building. It is used by several experiments running now, soon to run, or in the farther future: NOvA (for enhanced triggers), DarkSide-50, LArIAT, and Mu2e. The SCD was also awarded an LDRD to research and develop an "off-the-shelf" DAQ, using some of the plethora of new inexpensive devices designed for the "internet of things" community as opposed to the traditional approach of designing and building more expensive custom DAQ components. The engineers also operate "PREP", a pool of and repair for standard off the shelf electronic instrumentation for high-energy physics research.

- Simulations

  Simulations are crucial for arriving at physics results. Fermilab Scientific Computing is involved in many aspects of research and development of simulation tools. For example, we have key collaborators in the development of the event generators GENIE (physics of interaction involving neutrinos) and Pythia (hard scattering and event evolution for the LHC experiments).

  Geant4 is ubiquitous in HEP experiments for simulating the interaction of particles with the experimental apparatus and surrounding material, and the accuracy and speed of this application is under continuous improvement. The FNAL SCD is an anchoring member of the Geant4 collaboration and is developing, with CERN and other partners, Geant-V, a next generation of the application that takes advantage of emerging computer architectures and techniques.

- Workflow and production

  As described below, Fermilab's computing facility provides the computing resources necessary for physics results. Research, development, and integration are aimed to improve the efficiency, usability, and reliability of the facility itself as well as to improve the experiments use of the facility. Experiments interact with the facility with the integrated set of tools and services called *FIFE*, described in this document. *FIFE* also allows experiments to use resources at remote sites as well.

  The SCD develops many of the internal services and tools that run on the Fermilab facility, including Enstore (archive system), dCache (disk caching), glideInWMS (low level job coordination and management), SAM (data file and metadata management), and WMAgent (workflow manager for CMS). Monitoring the facility, production jobs and user jobs is crucial, and to that end the SCD develops "FIFEMon". Development of *jobsub*, a FIFE component, is important as the hub for submission of jobs to the facility and the Open Science Grid.

  Research on cloud computing facilities and dynamic provisioning of resources are important for optimizing and operating high throughput computing.

Database applications have a large and varied role in experiments, from archiving beamline information and tracking the experiment's runtime configuration to bookkeeping for physics analyses. We develop database applications to those ends as well as research the evolving landscape. We develop tools for collaboration including the Electronic Control Room Logbook.

- Consultations
  Collectively, the Scientific Computing organizations at Fermilab house an enormous store of knowledge and expertise on scientific computing and regularly consults with experiments on a variety of topics, including C++, build systems, data and job management, data flow, best programming practices, assistance with simulations and Geant4, and DAQ and controls hardware and software. Training and education are being developed and delivered such as a summer C++ class and the *art* workbook that contains many relevant examples for the benefit new users. As mentioned, SCD scientists and CS liaisons are deeply embedded in every Fermilab experiment, and bring to the experiment their computing knowledge as well as access to expertise, tools, and computing projects - from the planning and design phase through to production operations.

Scientific Facility

Fermilab provides a world-class computing facility for the support of its scientific programs. The facility contains resources for data storage, processing and analysis, connections to the data acquisition systems of Fermilab based experiments, and the networks to interconnect these both locally and globally. The facility provides the services that allow scientific users to manage their data and submit jobs to process that data. The facility also provides the compute resources needed for analysis, simulations, and theoretical calculations. The Computing Divisions operate this facility, monitor its performance, interact with the experimental community to determine capacity, availability, and capability requirements, procure resources as necessary, and research current and future computing architectures, operating systems, and management tools.

The people within computing who operate the facility possess a wide range of skills, including detailed knowledge of processor architectures, tape and disk storage hardware architectures, network architectures, the Linux operating system (Fermilab is home to *Scientific Linux*) and computer security. The facility staff operates the services necessary for a distributed computing environment, including batch, grid, authentication, authorization, virtualization, and cloud functions. The staff designs, procures, installs, configures, and operates a massive high-performance data storage system. The skills of the facility personnel are as valuable as the physical resources.

The core of the data storage facility consists of seven 10,000 slot robotic tape libraries. The current active utilization is 54 petabytes, with an additional 10 petabytes of available tape media. As tape technology advances, there is a program to migrate data from older to newer formats in order to ensure future availability and preservation. The facility currently holds an additional 33 petabytes of data that has been so migrated but is retained as a secondary copy. With the current tape technology (T10000-D), the tape library complex has a potential capacity of over one-half exabyte.

Access to the data on tape relies on a complex of over 100 enterprise class tape drives. Average data transfer to/from tape is currently over 1.5 petabyte per month (50 terabytes per day), and reached 6 petabytes per month during the previous LHC running period.

Although the tape system is highly capable, the bulk of the data access is from a complex of buffer and project disk systems with a total capacity exceeding 26 petabytes. About 50% of this is dedicated to CMS Tier-1 managed production activities, 20% to Fermilab-based (general purpose) experiment production and analysis activities, and the remainder to user (including CMS LPC) analysis needs. The disk system serves as the front-end to the tape system, and is supported by the data management services and tools necessary for efficient local and global distribution of content. The disk complex and associated network provides aggregate read rates averaging over 100 Gbits/s for the CMS dCache and EOS instances, and over 10 Gbits/s for the general purpose dCache instance. Write rates, as expected from the caching nature of the disk systems, are 1-10% of the read rates and predominantly reflect the addition of new raw data or analysis production results.

As a strategic goal,, an Active Archive Facility (AAF) is being developed to enable, with cost recovery, external scientific collaborations access to the Fermilab tape data storage resources. A Work For Others agreement has been established with one university computing center, and this agreement is being used as a template for agreements with several other organizations. The AAF is seen as a valuable service to scientific efforts that are otherwise without access to a facility that provides massive capacity yet ease of access through well exercised and optimized data transfer tools. The monitoring and accounting tools needed for the AAF additionally provide benefits to the storage facility operational support of many distinct internal customers.

Resources In the Facility

The facility provides a complex of approximately 2,500 nodes with 50,000 CPU cores for High Throughput Computing (HTC) simulation and analysis jobs. Roughly 50% of this quantity serves CMS Tier-1 and LPC functions, 15% is for legacy (and in the process of being retired/repurposed) Tevatron experiments, and the remaining 35% (General Purpose) supports the Fermilab based experimental program. Most of the compute capacity, roughly 30,000 job slots, is grid enabled and can thus serve multiple scientific experiments and programs. The General Purpose grid facility has noted usage by approximately 25 distinct experiments or projects over the previous year.

The facility also operates a set of High Performance Computing (HPC) clusters, predominantly for the US Lattice QCD (LQCD) program. Three clusters, with a total of 25,000 CPU cores and Infiniband interconnects, are operated for LQCD. Two of these are also partially configured with GPUs. There are two additional smaller (totaling 2,500 CPU cores) HPC clusters operated to support accelerator modeling, cosmology simulations, and general purpose HPC activity.

In addition to the large number of compute nodes for bulk data processing, the facility provides the infrastructure design, procurement, configuration and system management for higher level services. These include database servers, configuration and monitoring systems, interactive systems, dynamically instantiated (virtual or private cloud) development and integration testing systems, authentication and authorization systems, software build and release validation systems (and supporting software), software and code distribution systems (cvmfs), and easily configured and instantiated interactive systems. The facility operations expertise is also leveraged to aid in procurement, configuration, and system management of the Fermilab-based experiment DAQ and control room systems.

The data and processing systems are enabled by high speed networks that interconnect the three data centers. A fabric of connections provides 40 to 80 Gbits/s links between distribution switches in the data centers. Individual compute and server nodes have 1 to 10 Gbit/s links to the access layer switches, which in turn have multiple 10 Gbit/s links to the distribution switches. WAN connectivity is supplied by ESNET with a 100 Gbit/s connection (soon to be upgraded with an additional 100 Gbit/s link) with additional secondary multiple 10 Gbit/s paths.

The Open Science Grid (OSG) is treated as an extension of the facility. That is, the "facility" is considered to be the target of requests for processing resources, and those resources are provided from local Fermilab systems or by submitting work opportunistically to the greater OSG. Similarly, the Fermilab facility can accept work from the external OSG Virtual Organizations (VOs) to opportunistically utilize idle Fermilab resources. Opportunistic (non-Fermilab OSG VOs) use of Fermilab resources averaged 3400 job slots (slightly over 10% of capacity) last year, equivalent to approximately 27M CPU-hours. Fermilab-based experiments and projects conversely utilized approximately 1M CPU-hours at OSG sites external to Fermilab. Fermilab use of OSG is expected to rise as the workflows of more experiments become compatible with the OSG.

Extending to the Cloud

The Fermilab facility is in the process of being extended to utilize commercial external (cloud) compute and storage resources. This provides the scientific computing user community with a single view of the Fermilab funded facility, but enables the facility to make the economic decision as to whether to provide the actual resources locally, request the resources opportunistically from the OSG or other federated facilities, or to purchase the resources from a cloud provider. Any business decision is guided by the science objectives as reviewed by the SCPMT process.

The facility program includes a survey, analysis, and detailed study of new computing hardware and techniques. A significant effort is associated with provisioning of the HPC hardware, including performance benchmarking of GPU and other accelerators as well as the fabric interconnects necessary for an HPC cluster. Other investigations study storage (tape, disk, SSD) hardware and software.

Science Operations and Workflows

Scientific Computing provides the essential services needed for experiments to access very large data sets across a distributed computing infrastructure - both local and remote - in a robust and scalable way. We develop, configure, and operate job submission and data management services that allow reconstruction and simulation processing as well as analysis to run on the facilities.

We provide world class services for CMS on the Tier-1 and LPC facilities with custom solutions for CMS. This started with LHC Run 1 and will continue through LHC Run 2 and beyond. For the smaller Fermilab based experiments and scientific projects, we provide a toolkit of common services under the FabrIc for Fermilab Experiments (FIFE) umbrella. FIFE has been able to capitalize on some of the CMS developed tools and we hope to bring them closer together over time.

FIFE integrates a suite of underlying software packages that enable a single interface to submit jobs both to the Fermilab facility, to the OSG and lately also to commercial clouds through the jobsub tool; move and catalog data files through the Sequential data Access via. Meta-data (SAM) system; and provide monitoring and accounting services. The fifemon monitoring suite provides a simple interface to view jobs both for the scientists running jobs as well as the operations support staff. To date, 14 experiments/collaborations have been onboarded in various degrees to FIFE, including, but not limited to, darkside50, LArIAT, LBNE 35T, MicroBooNE, MINERvA, MINOS+, Mu2e, Muon g-2, NOvA, and SeaQuest. Most of these experiments also use interactive facilities at Fermilab for their user's needs.

Production and Workflow Management

CMS is operating on a diverse and world-wide distributed computing infrastructure since the beginning of LHC data taking and even before. We are providing a complete workflow management suite called WMAgent on top of the glideinWMS submission infrastructure to run data processing and MC production workflows on the more than 60 CMS GRID sites and opportunistic resources. Average scales of 60-70k jobs in parallel are reached through a highly distributed service, which is also the basis for the Tier-0 processing infrastructure handling the processing of just recorded data and the analysis GRID tool of CMS. In addition to the basic services used on resources all over the world, we provide leadership for the operation of central production and processing workflows as well as world-wide transfer operations. We handle site readiness and share the load of running the Tier-0 making Fermilab one of the two centers for CMS computing operation in the world. User support is particularly important for a very large and geographically diverse collaboration. At Fermilab, we are operating an interactive clusters primarily for US users to conduct their analysis activities and submit jobs to local resources and the GRID.

Following CMS' example, this year we have embarked on a service to run intensity frontier production processing workflows or a few select experiments currently taking data at Fermilab using the successful model deployed for CMS. NOvA, minerva, and minos are all successfully onboarded for their most crucial workflows. More will be onboarded this coming year.

Cosmic Frontier Experiments

We participate in the planning, construction, operations, and data preservation phases of cosmic frontier experiments. The current experiments are the Dark Energy Survey (DES), Dark Energy Spectroscopic Instrument (DESI), and Large Synoptic Survey Telescope Dark Energy Science Collaboration (LSST-DESC).

DES is currently in operations. Specific scientific computing roles include support for a number of software and database systems, survey strategy and planning, calibrations, data backup, and data analysis systems. SCD scientists participate in DES scientific analyses including galaxy clusters, strong and weak lensing, and photometric redshifts as part of the dark energy program. A number of SCD scientists also have leadership roles in DES science working groups and DES data management. We have been also providing CPU cycles for DES data processing via FermiGrid.

DESI is in the design phase at present. Specific SCD responsibilities include development of an online database system and code for the fiber positioner placement. LSST is about to enter construction. SCD responsibilities here include working with the Dark Energy Science Collaboration (DESC) software infrastructure group to develop a framework for integration of project software systems with science analysis codes.

Distributed Computing Services

Both the CMS and FIFE workflow operations work as partners in the open science grid consortium to allow access to OSG sites across the US. In addition, many experiments will rely on OSG computing resources to meet their FY15 (and beyond) simulation capacity needs. We provide support for transparently starting experiment and user jobs at grid sites and clouds infrastructures via *glideinWMS*. SCD has developed and operates the *GRATIA* accounting services for the OSG. Vital in this distributed computing model is the need to protect resources from unauthorized access yet simultaneously allowing collaborations to share data. We provide consulting and solutions for security and trust management for Scientific Computing.

Scientific Database Applications
Database applications provide the necessary metadata to understand how to interpret the data. Scientific Computing activities include developing and operating a number of database applications for the scientific collaborations including hardware, detector conditions, run information, and accelerator data. These types of databases are currently in operations for MINOS+, NOvA, MINERvA, MicroBooNE, and DES.

Collaborative Software Development
In addition to providing tools to help users run their job, we also have a suite of services to help them work within their collaboration in particular with regard to software development. We operate source code repository systems in the form of GIT, SVN, and CVS, and a REDMINE twiki to manage and document the packages. We currently support over 180 projects in REDMINE. Collaborations have long used the facility services to acquire interactive machines for building and testing. This year we have introduced a Jenkins based build service that would allow for regular[eg, nightly] builds of experiment software including MacOS builds. Once software releases are complete, code needs to be available to run on worker nodes across the grid. We operate a CVMFS service hosting repositories for a number of experiments. A final piece of our collaboration tools is an electronic control room logbook is also included as one of the services within the collaboration tools and it is operated for 20 collaborations at the lab.

R&D projects and collaborations

In order to achieve our strategic and tactical goals we pursue R&D projects utilizing collaborations with both other labs and with universities:

- COMPASS remains at the forefront of the SciDAC projects contributions to accelerator simulation science. Fermilab maintains leadership positions and expects to build on the SCiDAC-2 and SciDAC-3 R&D with future projects in manycore, fine-grained parallelization, and use of the new advanced computing provided by the HPC centers – continuing the close collaboration and synergies with computer and domain scientists at UCLA, LBNL and ANL. (Spentzouris, leading-PI and overall coordinator).
- PDACS (Portal for Data Analysis Services for Cosmological Simulations), a collaboration with ANL, to demonstrate simulated data workflow processing through the Galaxy workflow engine, aiming to serve the community.
- CosmoSIS is an international effort (Fermilab, Manchester, Chicago) that aims to develop the framework and libraries necessary for parameter estimation using Markov Chain Monte Carlo for DES and LSST needs. This is a true community project, with standardized packaging and build system that enables community contributions.

- The Open Science Grid contributions - complementing the usage described above - include local staff providing OSG project management, security (including the Security Officer), user and facility support, accounting and monitoring services, and the Council chair.
- R&D in advanced networking using concepts such as Software Defined Networking (SDN), in-flight capture and application oriented network analysis using GPUs.
- R&D towards wide-area federated data management and access using technologies such as XROOTD and big-data concepts, building from the success of SAM and the US CMS AAA project, with initial use-cases for US CMS and NOvA.
- R&D towards future preservation of the data and analysis capabilities of the new round of experiments as well as CMS will build on the successful project for Run II. Efforts include collaborating as part of the Data and Software Preservation for Open Science as well as Inspire/invenio with the Core Computing Division.
- A multi-year CRADA with KISTI has provided the means for both institutions to research and develop services for private scientific clouds and enable access to commercial clouds for existing and new applications.
- Active leadership participation in the new cross-cutting HEP Computing Center of Excellence and international Software Foundation. (Roser, co-leader).
- HEP funds computing for the particle and nuclear physics lattice communities through the SciDAC program. This program is organized by the USQCD collaboration, where Fermilab plays a leading role. We deploy and operate large computing clusters for the use of the United States lattice gauge theory community, and our most recent software efforts focus on porting software infrastructure and algorithms to next-generation hardware from NVIDIA and Intel. Our lattice software group plans to work with the accelerator modeling group to develop a common approach to this next-generation hardware, further increasing leveraging of resources. (Mackenzie is the PI of the particle physics SciDAC grant.)