

Data Management and Its Role in Delivering Science at DOE BES User Facilities – Past, Present, and Future

Stephen D. Miller, Kenneth W. Herwig, Shelly Ren, Sudharshan S. Vazhkudai

Oak Ridge National Laboratory

Pete R. Jemian

Argonne National Laboratory

Steffen Luitz, Andrei A. Salnikov, Igor Gaponenko

SLAC National Accelerator Laboratory

Thomas Proffen, Paul Lewis

Los Alamos National Laboratory

Mark L. Green

Tech-X Corporation

millersd@ornl.gov

Abstract. The primary mission of user facilities operated by Basic Energy Sciences under the Department of Energy is to produce data for users in support of open science and basic research [1]. We trace back almost 30 years of history across selected user facilities illustrating the evolution of facility data management practices and how these practices have related to performing scientific research. The facilities cover multiple techniques such as X-ray and neutron scattering, imaging and tomography sciences. Over time, detector and data acquisition technologies have dramatically increased the ability to produce prolific volumes of data challenging the traditional paradigm of users taking data home upon completion of their experiments to process and publish their results. During this time, computing capacity has also increased dramatically, though the size of the data has grown significantly faster than the capacity of one's laptop to manage and process this new facility produced data. Trends indicate that this will continue to be the case for yet some time. Thus users face a quandary for how to manage today's data complexity and size as these may exceed the computing resources users have available to themselves. This same quandary can also stifle collaboration and sharing. Realizing this, some facilities are already providing web portal access to data and computing thereby providing users access to resources they need [2]. Portal based computing is now driving researchers to think about how to use the data collected at multiple facilities in an integrated way to perform their research, and also how to collaborate and share data. In the future, inter-facility data management systems will enable next tier cross-instrument-cross facility scientific research fuelled by smart applications residing upon user computer resources. We can learn from the medical imaging community that has been working since the early 1990's to integrate data from across multiple modalities to achieve better diagnoses [3] – similarly, data fusion across BES facilities will lead to new scientific discoveries.

1. Introduction

The US Department of Energy Office of Basic Energy Sciences operates nine synchrotron radiation and high-flux neutron sources. In addition, it supports five nanoscale science research centers and another three electron beam centers [1]. In 2008, these user facilities enabled over 10,000 researchers to perform experiments which in turn engaged additional thousands of collaborators to also analyze data from these experiments. Common to these facilities is the production of experiment data, and the need for users to access, process, and analyze these data, resulting in peer-reviewed science publications.

This paper examines data management methodologies and their relationship to producing scientific results across selected DOE BES user facilities including neutron scattering facilities Intense Pulsed Neutron Source (IPNS - now closed [4]), High-Flux Isotope Reactor (HFIR) [5], Spallation Neutron Source (SNS) [6], and Los Alamos Neutron Science Center (LANSCE) [7], and X-ray facilities Advanced Photon Source (APS) [8] and Linac Coherent Light Source (LCLS) [9]. This combination accounts for about 67% of the total annual BES operating budget for user facilities [10]. The maturity of these facilities is also significant as it varies from retired, mature, emerging, and new user facilities. Examining these facilities provides a glimpse into the past and present, and a look to the future for the role of data management in accomplishing their mission. Facilities are examined individually, trends are identified, and a model for the future is put forward. Of particular interest are the volume of data produced by typical experiments, user access to data and computing resources, and the role of the facility in enabling scientific research by providing intelligent data management.

The premise of this paper is that facility produced data are valuable. User facilities with beam lines are rate limited as to the number of experiments that can be performed annually according to a number of factors. The unique scientific capabilities of these facilities are a scarce resource with demand exceeding availability typically by factors of two to three. Access is obtained via a peer-reviewed process unique to each facility but which is uniformly extremely competitive. Those who are successful in obtaining beam time greatly value the resource and often depend upon it for their individual science programs. We can also calculate a metric for the value of the data produced. In round numbers, one can take the facility operating budget divided by the number of beam lines and the number of operating days to determine a \$/instrument/day operating cost. Knowing the average data rate per instrument per day, one can then determine the \$/TB to produce data. This cost varies widely from thousands to millions of dollars per Terabyte of data produced by many of the instruments at these facilities. Given the high cost of producing this data, it is reasonable to consider the investment in a reasonable data management strategy.

Currently much of this data particularly at X-ray sources are ultimately discarded based upon facility data retention practices which are often, understandably, based upon available storage resources for storing data. When a facility discards data, the burden falls upon users to manage (store and maintain its provenance), process, and analyze their data. Add in the complexity for doing so with data obtained at multiple instruments across multiple facilities over many visits and one can quickly see that users face significant data management and access bottlenecks along the road to scientific discovery.

Steps to change this process have been taken by the data management team for HFIR and SNS which has produced a web based portal that manages experiment data on behalf of the users and also provides users access to computing resources. Thus far, all experiment data produced for SNS have been archived and are available via the portal as well as secure shell access. Eventually the composite size of the repository will grow beyond what the facility can maintain itself, and this point may only be only a few years away even though the facility is expected to have a 40 year lifespan. The challenge increases significantly for X-ray facilities which typically produce multi-TB data sets per experiment. Yet even these data volumes look quite tractable when compared to the data volumes produced by the High Energy Physics (HEP) community which has data volumes approaching the Exascale [21]. To facilitate international research, the HEP community has developed an extensive data storage and analysis infrastructure including over 140 computing centers in 34 countries [23].

Another premise of this paper is to enable user access to data on the same material obtained with multiple techniques (at multiple facilities) will facilitate a new mode of scientific discovery. This leads to the desire for a common approach to data management across these user facilities. There are already a significant number of cross-over users between X-ray and neutron scattering in the fields of powder and single crystal diffraction, small angle scattering, reflectometry, and imaging and tomography. Ultimately, the success of this approach also depends on development of more sophisticated analysis methods which allow easy integration of these multiple data sets collected in a variety of modalities but all seeking to address a common scientific goal.

To this end, work is already underway between APS and SNS to facilitate inter-facility data caching where selected data are to be made available via a storage lease for a period of time. This work is being done among the facility data management personnel working in conjunction with the CEDPS SciDAC team [22]. There are many challenges to face in providing such a resource including user authentication among users and user facilities, cyber security, and cyber infrastructure development, and human labor to develop the necessary software and systems. It is hoped that demonstrating the benefits and how they outweigh the costs will help bring more resources to this work perhaps leveraging more of the SciDAC resources and teams.

2. Neutron Scattering User Facilities

Neutron scattering user facilities generally fall into two broad categories – reactor-based providing a continuous source of neutrons, or accelerator-based generating pulses of neutrons. HFIR is reactor-based, and IPNS, SNS, and LANSCE are pulsed sources. Neutron scattering techniques first utilized continuous neutron sources and were found useful for uncovering both the structure and dynamics at the atomic and molecular levels. The Nobel Prize for physics was awarded to Clifford G. Shull and Bertram N. Brockhouse in 1994 “for pioneering contributions to the development of neutron scattering techniques for studies of condensed matter” [11].

Utilizing an event data format is taking root at the neutron scattering facilities where detected neutrons are stored in a list where the individual event detection records the detected position along with the timing information for when an event is detected. Additional metadata (sample conditions, dynamic instrument parameters) are also archived in a time-stamped list and can be correlated with the data post-collection. The event list can then be (re)histogrammed as desired allowing researchers to sift through the raw data to construct data sets of interest. For example, parametric studies over temperature, pressure, and/or magnetic field changes etc. can be rebinned over desired parameter spans in order to reach a desired statistical data quality or desired “error bar” range. Thus the event data format is ideal for supporting pump-probe type experiments, which is also of interest for X-ray experiments – significant new scientific discoveries are forecast utilizing the pump-probe technique.

The neutron scattering community is quite fortunate in that up to this point, facilities have managed to collect and store all of the data that has been produced. Though this is the case now, the future holds challenges to continue this pace. Several new instruments will have the capability to produce data sets of similar size as those more typical of X-ray instruments. Thus, a new strategy will be needed for neutron scattering facility based data management.

2.1. IPNS

As of January 2008, IPNS began to undertake decommissioning and decontamination activities which included ramping down and terminating the user program [4]. IPNS had been operational for more than 26 years and had collected approximately 520GB of data across 14 instruments over the lifetime of the facility. Users typically went away with their experiment data or had anonymous ftp access to the data. Metadata were sparse, thus finding data via ftp could be challenging.

2.2. HFIR

The HFIR facility has been refurbished in 2006-07 and currently has 15 instruments operating via an 85 MW reactor neutron source. During the lifetime of the facility, the reactor has had over 420 run-

cycles to produce neutrons where a run cycle is approximately 3 weeks in duration. Since 2007, HFIR data for 5 of the instruments has been archived totalling around 2GB of storage space. Experiment and processed data are generally manageable in size and instrument users typically leave taking their experiment data with them. However, HFIR has plans to operate a tomography beamline which will produce hundreds of GB to TB size experiment data sets, causing data management for this instrument to be rethought.

2.3. LANSCE

The Lujan Neutron Scattering Center within LANSCE utilizes a pulsed source and has a complement of 15 instruments. It maintains a data archive of approximately 4 TB that includes all neutron scattering data collected since its 1986 inception. The size of experimental data sets has grown significantly over the past few years. Raw data collected for a typical user experiment is on the order of 2 GB, but can vary widely depending on the instrument and the experiment particulars. Generally, initial data reduction is performed on site and users depart with a reduced data set. Access to both the raw and reduced data is available through the instrument scientist. Ever tightening computer security regulations over the past several years has made it difficult to provide users with direct offsite data access.

An initial proof of concept effort has been undertaken to include data from the Lujan NPDF instrument into the Neutron Science Web Portal in which approximately 2.5GB of data have been made available via the portal. To access this data, users are required to get an ORNL external computing user (XCAMS) ID and request portal access.

2.4. SNS

The SNS facility took its first data in April 28, 2006 with instruments collecting data via the user program in September of 2007 with first user data collected on the BASIS backscattering spectrometer. Currently there are 7 instruments in the user program with another 4 to enter commissioning soon in 2009. These instruments are located in target station 1 which will have a complement of approximately 23 instruments when fully built out, and an additional target station 2 is being planned to house a similar complement of instruments. Thus far data from every experiment has been archived for the SNS instruments providing excellent opportunities to examine data statistics. The total SNS archive consists of approximately 2.8TB with another 1TB of user contributed data (shared) data. Experiment data rates vary from around 200MB/ day to around 4.7GB/day per instrument with an average of around 1.3GB/day/instrument with a standard deviation of around 1.6GB/day. Thus on average assuming an operating budget of approximately \$180M and about 200 operating days over a full complement of 23 instruments the value of data at about \$32M/TB. Compare this with around \$2K/TB in a high performance data storage system, which is over 4 orders of magnitude difference.

Within the next year, SNS will begin commissioning a diffractometer which will have the ability to produce around 200GB/hour resulting in tens of TB per experiment. In addition, both SNS and HFIR have plans to support tomography instruments which could routinely produce hundreds of GB to tens of TB per experiment. Techniques such as inelastic single crystal experiments will produce multi-TB sized data sets as well. In addition, pump-probe experiments utilizing an event data format have the potential to produce very large data sets perhaps approaching 1TB per experiment in addition to the raw experiment data.

3. X-ray User Facilities

BES supports 6 synchrotron radiation light sources with techniques for spectroscopy, scattering, and imaging. Here we focus on APS which is a mature, flagship facility supporting thousands of users producing hundreds of publication per year, and LCLS which is a new highly-anticipated facility expecting first beam in September of 2009. APS provides a full suite of instruments addressing all of the above techniques where LCLS produces ultrafast pulses of X-rays enabling imaging of single

molecules and stop-motion of atoms and molecules in motion. The LCLS coherent hard X-ray beam will be delivered in a form of 120 Hz pulses (30 Hz in Fall 2009) with 100 fsec timing. LCLS instruments will support imaging, spectroscopy, and pump-probe measurement techniques for both imaging and spectroscopy.

3.1. APS

APS has almost 60 beamlines with 1-3 instruments per beamline and rapidly produces copious amounts of data. Typical experiments will produce between a few KB to 100GB, while imaging experiments (such as tomography) produce more, on the order of 1TB to 10TB of data per experiment. Data rates for some instruments, such as X-ray Photon Correlation Spectroscopy and 3D X-ray Diffraction Microscopy will approach 300 MB/s on a continuous basis. Data management across beam lines varies considerably but can be best described as central to each instrument. A variety of data formats are in use, each chosen by the instrument support team based on the needs of the particular community. Some beam lines utilize facility-defined formats while other beam lines, based upon expediency, just utilize company proprietary data formats as provided by the detector or data acquisition system. Adoption of international standard formats such as NeXus is slow due to many factors. Data storage and archiving policies are local to each instrument. For instruments with small data volumes, data is retained in perpetuity. For high-volume instruments such as imaging or tomography, the amount of local storage is constrained by cost in dollars, time, and physical storage space for the media. For these instruments, on-site storage resources are limited and thus users are responsible for storing and managing their own data. The facility will provide users a two week window of opportunity to copy their data before it is deleted from its servers.

APS is beginning to address the need for a facility approach to the coordination of metadata associated with the primary data from each experiment. Collection of metadata can facilitate many aspects of downstream data processing (reduction, analysis, access, and archival) and provide the possibility of data fusion. The metadata collection process begins at the facility with the user's experiment proposal and continues through publication.

Real-time access to high-performance computing is necessary to handle the data processing for advanced instruments such as the XPCS noted above. Communication of the data through the network and access to the data and computing capabilities must be built upon standardized security infrastructure and high-volume data communication protocols that spans beyond the APS to include the breadth of the DOE facilities and user communities. Standardized infrastructure to manage the HPC capabilities will enable the user community to focus on the production of science from the mass of accumulated data.

3.2. LCLS

There will be 6 sets of versatile high data bandwidth instruments installed in two hatches of the LCLS experimental area. Some instruments will be capable of producing up to tens of Gbytes/sec of data in peak. In the final implementation of the system up to two of those instruments can be used simultaneously. This will result in multi-terabyte data volumes to be handled on a daily basis. The data rate will ramp up over the next 3 years. The first experiments begin their data collection in September 2009 and will produce up to 1 TB of data per day with nearly 100 TB of data to be stored before the end of the year 2009. In three years the amount of data to be stored per day will increase up to 15 TB from only one of the instruments, and that would correspond to nearly 2-3 PB of data per year. All stored data will be preserved at SLAC Computer Center for the duration of at least 1 year. All stored data will have to be delivered to the corresponding experimental groups for further analysis.

Unprecedented amounts and complexity of data expected at LCLS which poses a number of challenging problems to be addressed:

- Implementing a high-bandwidth data delivery pipeline for sending hundreds of MB/sec of data from the detectors to the long-term storage and archival system

- Providing a reliable and secure storage mechanism for many petabytes of data for a period of at least one year (most likely – much longer)
- Building an efficient data access mechanism to the stored data. Such mechanism should allow parallel access to the data from an analysis farm of thousands of computing cores.
- Building a data exportation system to send terabytes of data to off SLAC consumers. The system would need to cope with the same aggregate data rate the data are delivered from instruments.

The architecture of the data management system which is currently under development is based on iRODS [17] – the next generation data management infrastructure introducing the highly customizable rule-based management model. The main file format for stored data will be HDF5 [18]. In addition there will be a number of specialized relational databases augmenting the bulk data.

Users will interact with the data storage system primarily via the Web portal which is currently under development. The portal will be based on contemporary dynamic Web 2.0 technologies in order to deliver enhanced user experience with the system.

The off-site data exportation will employ high-performance multi-stream data transfer protocols such as BBCP [19] which would allow utilizing 30 Gbit/sec SLAC WAN connection with ESNNet [20]. Building a distributed data exportation system involving a number of regional centers is also planned for later stages of the LCLS experiments. Such a distributed architecture will facilitate efficient data delivery to remote clusters of users. The development will require substantial cooperation and coordination efforts with other BES laboratories and with international counterparts.

4. Future Cross-Cutting Science Opportunities

In looking at methods for scientific discovery, today's facilities appear as islands with the instrument appearing as isles among these islands. As a result, scientific research is performed in relative isolation. To illustrate, today there are a plethora of data file formats among these various instruments, and this alone is a significant inhibitor of software re-use and interoperability, and collaboration. Another serious inhibitor is the challenge to simply move data between facilities. Even more basic is the lack of robust authentication systems to support facility users.

Assume for the moment that we have overcome these obstacles and let us look at what we could stand to achieve.

- Data mining spanning across both neutron and X-ray data sets facilitating research for both experiment teams and non-experimentalist researchers.
- Creation of new cross-technique reference databases for researchers.
- Common software to work with data for visualization and basic analysis.
- Definition of experiments spanning across facilities such as advanced pump-probe experiments utilizing both neutron and X-ray sources – novel scientific discoveries await this capability.
- Seamless co-refinement of data from X-ray and neutron single crystal and powder diffraction instruments.
- Better enable the study of biological materials via integrated SANS and SAX techniques.
- Combination of imaging and tomography utilizing neutrons and X-rays in order to better exploit the measurement properties of each technique.
- Incorporation of high performance computing at the beamline to provide better sample alignment, experiment steering, and experiment simulations.
- Development of a common support infrastructure across the user facilities better enabling cross-facility support and integration.
- The ability to readily collaborate with academic and international partners.
- A common or federated to approach user authentication and data access controls which comply with federally-mandated computer security.

5. The Future of Facility Data Management

The BES user facilities are well poised to take advantage of software and systems developed via the SciDAC program [12]. With the advent of user facilities producing “bigger data”, the need for data movement and data management tools grows larger. However it appears that software produced by the SciDAC Collaboratories may be applicable for helping to address some of the facility data management issues[13]. There is also good intersection between facility needs and the resources provided by a number of SciDAC Institutes [14] and Centers [15] as well such as visualization, data storage, scientific data management, networking, computing, and simulation.

The user facilities have several significant additional requirements necessary to establish dynamic and versatile scientific research communities. These can be roughly divided into two categories: hardware infrastructure, and software systems – yet these both must be combined to perform synergistically. At the foundation is needed the infrastructure to provide data storage capacity at some level serving the user facilities. This includes the ability to capture metadata to support data cataloguing incorporated into a “facility node” (or “network node”) which provide some level of computing, data storage, and networking capacity independent of the operation of the facility. Then a user facility network we are calling the User Facility Network (UFNet) needs to be established to enable networking across the facilities interfacing to the network nodes at each facility. UFNet would operate synergistically and leverage ESNNet to provide inter-facility networking capacity and capabilities with federated user authentication to support scientific community growth. These *network nodes* would provide the abstraction layer between facility operations and the infrastructure supporting the user to access data and computing resources. Facilities need to have what we are calling “*nodeware*” to interface their facility data collection processes, archival systems and computing resources with their *network node* and UFNet will provide the middleware infrastructure connecting the nodes together.

At the upper level is the need for the users to work smarter providing them the ability to easily leverage the above described infrastructure. Users need not be concerned about computing-related issues in order to pursue their research. Work is already underway via a company called the Tech-X Corporation to via SBIR efforts to produce an “*Orbiter*” thin- and thick client application [16] residing on a user computer which provides programmatic access to remote data and computing resources provided by a common Service Oriented Architecture (SOA) hosted within the HFIR/SNS network. Users are also provided the opportunity to work autonomously or to define their own research teams via the definition of *ad hoc* virtual organizations.

There are significant benefits to this approach:

1. Facilities provide a uniform interface to data and computing resources thus better enabling their reach into research communities.
2. *Significant* return on investment for the SciDAC project
3. Extensible to incorporate scientific research communities external to DOE
4. The user can perform scientific research holistically by creating data portfolios from among the user facilities as necessary
5. Users can work better both collaboratively and autonomously.
6. As we learn from medical research, advancement requires multi-technique data integration. Enabling integrated scientific research *is* the next tier of scientific discovery.

References

- [1] <http://www.sc.doe.gov/bes/BESfacilities.htm>
- [2] “The SNS/HFIR Web Portal System – How Can it Help Me?” Stephen D. Miller Invited Talk I3.1 2009 International Conference on Neutron Scattering.
http://www.mrs.org/s_mrs/bin.asp?CID=19203&DID=239713&DOC=FILE.PDF
- [3] Medical Imaging, Analysis of Multimodality 2D/3D Images Kuhn, M. And Beolchi, L. January 1, 1995. **ISBN-13:** 978-9051992106
- [4] IPNS: <http://www.pns.anl.gov/>

- [5] HFIR: http://neutrons.ornl.gov/hfir/hfir_facilities.shtml
- [6] SNS: http://neutrons.ornl.gov/facilities/facilities_sns.shtml
- [7] LANSCE: <http://lansce.lanl.gov/>
- [8] APS: <http://www.aps.anl.gov/>
- [9] LCLS: <http://lcls.slac.stanford.edu/>
- [10] http://www.science.doe.gov/bes/archives/budget/BES_FY2009budget.pdf
- [11] http://nobelprize.org/nobel_prizes/physics/laureates/1994/press.html
- [12] SciDAC: <http://www.scidac.gov/>
- [13] SciDAC Collaboratories: <http://www.scidac.gov/collabs.html>
- [14] SciDAC Institutes: <http://www.scidac.gov/institutes.html>
- [15] SciDAC Centers: <http://www.scidac.gov/aboutET.html>
- [16] “Doing Your Science While You’re in Orbit”. Mark L. Green, Poster I4.7 2009 International Conference on Neutron Scattering.
http://www.mrs.org/s_mrs/bin.asp?CID=19203&DID=239713&DOC=FILE.PDF
- [17] “IRODS – Rule-Oriented Data Systems”: <https://www.irods.org/>
- [18] HDF5: <http://www.hdfgroup.org/HDF5/>
- [19] BBCP – multi-stream file transfer protocol for WAN: <http://www.slac.stanford.edu/~abh/bbcp/>
- [20] ESNet – Energy Sciences Network: <http://www.es.net/>
- [21] <http://www.science.doe.gov/sbir/solicitations/FY%202009/34.HEP.Computing.htm>
- [22] <http://www.cedps-scidac.org/>
- [23] <http://lcg.web.cern.ch/LCG/>