

Scientific Collaborations for Extreme-Scale Science

Workshop Report

December 6–7, 2011
Gaithersburg MD



Scientific Collaborations for Extreme-Scale Science

Workshop Report

December 6–7, 2011, Gaithersburg MD

Tom Abel, Manfred Auer, David Bader, Lothar Bauerdick, Brian Bockelman, Scott Brim, Richard Brower, Francesco De Carlo, Lali Chatterjee, Jackie Chen, Terence Critchlow, Silvia Crivelli, Eli Dart, Anshu Dubey, Michael Ernst, Ian Foster, Dan Fraser, Rob Gardner, Al Geist, Ian Gorton, Steven Gottlieb, Martin Greenwald, Dan Gunter, Bill Johnston, Jeff Kantor, Kate Keahey, Scott Klasky, Jerome Lauret, Kian-Tat Lim, Miron Livny, Stefano Marchesini, Robert McGreevy, Bronson Messer, James Osborn, Ray Osborn, Michael Papka, Elena Peterson, Don Petravick, Norbert Podhorszki, Ruth Pordes, Jeff Porter, Robert Ryne, Nagiza Samatova, David Schissel, Bill Shelton, Galen Shipman, Bill Tang, Brian Tierney, Thomas Uram, Jack Wells, Bobby Whitten, Dean Williams

Co-chairs: Richard Mount (SLAC) and David Skinner (LBNL)

Contents

Executive Summary.....	1
General Findings.....	2
General Recommendations.....	2
Specific Findings and Recommendations.....	3
Technologies for Teams.....	3
Technologies for Data.....	3
Technologies for Processing.....	4
Focusing the SC Strategy for the Support of Collaboration.....	4
1. Introduction.....	6
Potential Impact.....	7
2. Vision of the Future.....	8
Current State: Primitive collaboration systems hinder discovery	8
Desired Future State: A digital laboratory system that accelerates discovery.....	9
3. Specific Opportunities to Enable Science.....	10
Emerging Data-Intensive Science.....	10
Reconstruction Techniques	10
Visualization	12
Modeling, Analysis, and Theory	12
Simulation-Driven Science	13
Combustion	13
Fusion Energy Science.....	14
Climate Prediction	16
The Materials Genome	17
Data-Intensive Science.....	18
4. Mapping to Technology and Computer Science.....	20
Unifying Abstractions.....	20
Overview of Technologies	21
Technologies for Teams	21
Management of Collaboration.....	23
Collaborative Exploitation of Facilities.....	23
People and Roles in Collaborative Environments.....	24
Technologies for Data	26
Integrated Data Operations	26
Integrated Data Operations at Extreme Volumes	27
Integrated Data Operations on Heterogeneous Multidisciplinary Data.....	27
Provenance.....	28
Curation.....	28
Visualization	29
Technologies for Processing.....	29
Location Transparency of Processing	29
Resource Sharing and Co-scheduling	30
Monitoring Debugging and Forensics	30
Synergistic Services	31
5. Characterizing Technical Capabilities	33

6. Non-Technical Issues	35
Sociological Barriers to Collaboration	35
Removing Barriers to Collaboration	36
Gaps in the Office of Science Strategy for the Support of Collaboration	36
Focusing the SC Strategy for the Support of Collaboration	37
Appendix A: Science Narratives	38
Accelerator Science.....	38
Astronomy — LSST	39
Astrophysics Simulations	40
Biosciences	42
Experimental Magnetic Fusion Energy	43
Environmental Molecular Sciences Laboratory: A Long-Term View of Data Sharing,	
Analysis and Collaboration	45
Lattice Field Theory for High Energy and Nuclear Physics	46
Appendix B: Previous ASCR Collaboratories Projects.....	48
Open Science Grid.....	48
The National Fusion Collaboratory Project	48
The Earth System Grid	50
GridFTP and Globus Online.....	52
Grids and Virtual Organizations.....	54
Grid Security Infrastructure: Single Sign-on for Collaborative Science.....	55
DOE Public Key Infrastructure	55
Access Grid	56
References.....	58

Executive Summary

The U.S. Department of Energy (DOE) Office of Science is dedicated to world-leading research at the limits of achievable scale. Through its unique experimental facilities including the Advanced Photon Source, Linac Coherent Light Source, and Spallation Neutron Source, its collaborative exploitation of unique world facilities including ITER and the Large Hadron Collider, and its Leadership Class Computing Facilities, the Office of Science vigorously explores the extreme.

Extreme-scale science is not simply about facilities. It inevitably also requires the exponential acceleration of progress that can be achieved when many differing intellects attack challenges collaboratively. The Scientific Collaborations for Extreme-Scale Science Workshop, held December 6–7, 2011 in Gaithersburg, Maryland, focused on a strategic vision of how collaboration can be enabled, and on the research and development that will turn the vision into reality.

The workshop brought together scientists from Office of Science programs, scientists with major facilities responsibilities, and leading computer scientists from DOE programs and beyond. The enthusiasm for the challenge of unleashing collaboration was palpable, persisting through a highly participatory report-writing phase in the following weeks.

Collaboration increasingly pervades DOE science and thus touches on essentially all aspects of communication, computation, and data analysis. The workshop could neither define collaboration narrowly and ignore all wider issues, nor address the whole gamut of issues, most of which must remain domain-science or cyberinfrastructure responsibilities. The workshop's scientific organizing committee made an initial broad selection of collaboration-related issues that were then discussed at the workshop and ranked in the participatory report-writing phase.

Since the time of Galileo, science has relied on the printed publication as the principal vehicle for sharing information and recording achievement. Today data and software play increasingly important roles, and mechanisms other than the printed word serve to communicate. A vital component of an exponentially collaborative future will be a largely automated record of the provenance of data and ideas that will demonstrate that openness trumps secrecy as a route to career success.

The collaborative future must be characterized by *discovery* — all resources are easy to find; *connectivity* — no resource is an island; *portability* — resources widely and transparently usable; and *centrality* — resources efficiently, and thus likely centrally, supported. The Office of Science Collaboratories program must be driven by the needs of science, and guided by these four principles.

Following these principles, this report identifies key areas of research and development and ranks them for their likely benefit to science, their specificity to collaboration, and their exclusivity to the DOE mission. The report also characterizes the scale and nature of the effort needed for progress in each area.

The workshop left no doubt that a vigorous and coherent program of research, development, and support for collaboration will accelerate extreme-scale science. Mechanisms will be needed to ensure both focus and agility as science and the wider technological world undergo the evolutions and revolutions that the next ten years will bring.

General Findings

The workshop participants — from a cross section of DOE science programs, facilities, and computer science — communicated a broad enthusiasm for a coordinated and sustained approach to the removal of impediments to collaboration. Associated general findings include:

- GF1.** Unimpeded collaboration accelerates and empowers extreme-scale science: problems that were intractable to individual scientists become soluble; an infectious sense of intellectual excitement is created; science at extreme scale becomes attainable.
- GF2.** Impediments to collaboration for extreme-scale science can be readily identified, for example:
 - needlessly disparate ways of accessing facilities and information
 - high costs (complexity, insecurity, unreliability) of crossing administrative and community boundaries in a distributed collaborative environment
 - difficulty of discovering the knowledge and expertise relevant to possible collaboration.
- GF3.** Removing the impediments and empowering collaboration for extreme-scale science requires not only advances in computer science that will lead to novel frameworks and methods, but also effective technology transfer through adoption of best existing technologies and the removal of more banal impediments through concerted action across the Office of Science.

General Recommendations

- GR1.** Explore the creation of an Institute for Extreme-Scale Collaborative Science to play a coordinating role in advancing and promoting the adoption of scientific collaboration methods and technologies. The institute should promote communications among scientists, facilities, and computer science researchers aimed at removing barriers and maximizing the coupling of research to the collaborative science that it should empower. It is recommended that:
 - a. The institute should be a distributed entity.
 - b. The institute management and oversight should involve science programs as well as computer science and therefore should be jointly funded by the Office of Advanced Scientific Research (ASCR) and the SC science programs.
 - c. The institute should serve as a central information exchange point for the facilitation of collaboration, including relevant computer science research, software development, and deployments.
 - d. The institute should foster collaborative projects, especially those involving interdisciplinary efforts aimed at exploring or sharing technologies for extreme scale collaboration, by providing matchmaking where appropriate plus advice and support.
 - e. The institute should also be charged with organizing stakeholder-driven meetings and focus groups to reassess the evolving needs and priorities in the support of extreme-scale collaborative science.

- GR2.** Solicit proposals for computer science research in areas identified in this report as having high potential benefit to science, specificity to extreme-scale collaborative science, and little likelihood of being addressed by the wider world. Collaborative computer science/computational science proposals would be particularly welcome. The computer science research would be appropriately funded by ASCR.
- GR3.** Solicit proposals for development, consolidation, and support of software and services needed by extreme-scale collaborative science. In addition to ASCR, the DOE science programs should have a significant role in funding this program to ensure that it will be focused on the areas of most value to their science mission.

Specific Findings and Recommendations

Technologies for Teams

- 1. Finding:** Collaborative extreme-scale science requires a range of services to manage the collaboration itself. These may include machine-accessible lists of collaborators with their roles, authorities and duties, project planning tools, effort and financial reporting services, meeting scheduling and management services. The lack of uniformity and indifferent quality of many of the existing services is a major impediment to collaborative science.
- 1. Recommendation:** Develop Office of Science guidelines for centrally supported services, with special focus on the Software-as-a-Service mode of delivery. The central support and any essential research and development may be provided by ASCR in concert with science programs or by commercial services where appropriate.
- 2. Finding:** Extreme-scale science can be a unique laboratory for new ideas benefitting the science itself and even the wider world — the almost too obvious example of this being the invention of the World Wide Web. The need for research-driven ideas and the transformational effects they bring should not be ignored.
- 2. Recommendation:** Conduct research into new methods for delivering innovative services to meet unique needs of collaborative extreme-scale science, that meet needs for consistency, economies of scale, and sustainability.

Technologies for Data

- 3. Finding:** Integrated data operations — unimpeded discovery and collaborative exploration of all relevant data — is a dream that is currently out of reach of most scientists engaged in extreme-scale research. The dream seems particularly distant for interdisciplinary collaborative science. Where the dream can be realized, the benefits to extreme-scale science will be revolutionary.
- 3. Recommendation:** Solicit proposals in the area of Integrated Data Operations for research into groundbreaking new tools, and for the consolidation and generalization of promising existing approaches.
- 4. Finding:** A rigorous treatment of provenance is pivotal in extreme-scale science. Provenance allows complex analyses to be validated, allows credit to be rigorously attributed without

the latency of passing through journal publications, and brings many practical benefits such as avoiding the need to instantiate and preserve every derived dataset. Current provenance-tracking technologies have had little impact on extreme-scale science.

- 4. Recommendation:** Solicit proposals focused on moving existing provenance technologies into extreme-scale science, thus exposing the extent to which both fundamental research and generalization/support are required. Research and development to fill clearly identified gaps should also be encouraged.

Technologies for Processing

- 5. Finding:** Extreme-scale science requires multi-platform and multi-facility processes, but is rendered slow and labor intensive by the need to adapt to multiple architectures and policies. Location transparency of processing resources and seamless co-scheduling of both processing and experimental facilities are the dream.
- 5. Recommendation:** The first vital step in rendering differences invisible is to formalize how resource characteristics and policies are to be captured and made available to the automated workflow management systems that will achieve the goal of transparency. Research and development into this formalization and its implementation should be given priority.
- 6. Finding:** The ability to monitor and debug the software for extreme-scale science is a key factor in its reliability and usability. Because software stacks are connected both horizontally across networks and vertically at endpoints, reliable determination of bottlenecks and root causes requires integrated end-to-end observation of behavior. The challenges of analyzing this increasing volume of monitoring data are similar to data analysis challenges for scientific experiments, and require similar levels of intellectual effort. These tools must be collaborative and able to deal with both highly distributed computations and highly distributed collaboration members.
- 6. Recommendation:** The software components for proposed research projects should include, as part of the design, a plan for how monitoring information will be produced and consumed. Ideally, the solicitation would avoid unnecessary diversity by providing guidance as to preferred existing (and widely-used) formats and methodologies. The program should also solicit research on scalable data reduction and analysis methods for the growing volume of end-to-end monitoring information.

Focusing the SC Strategy for the Support of Collaboration

- 7. Finding:** The ASCR Collaboratories program and some existing DOE large-scale collaborative science projects have developed a range of tools supporting collaborative science, ranging from widely used utilities (e.g., GridFTP) to high-function but currently application-specific utilities (e.g., the ATLAS Distributed Computing System).
- 7. Recommendation:** Establish coordination across the Office of Science to ensure that the benefits to extreme-scale science of funding research into new approaches are correctly balanced against the benefits of supporting existing widely used tools or generalizing apparently domain-specific tools.

8. Finding: Accelerator-based experimental facilities serving emerging data-intensive science are a target of opportunity to foster collaborative research. These facilities are just now identifying their immediate and long-term needs for the support for collaborative research. In particular, data management and visualization techniques and learning from more mature collaborative sciences should be researched for this particular scientific domain.

8. Recommendation: Research opportunities should support joint activities that will allow emerging data-intensive collaborative science to engage with more established data-intensive collaborative science disciplines. This will promote an understanding of where generalizations of existing successful approaches might be the best way forward for a range of collaborative sciences.

1. Introduction

A DOE workshop on Scientific Collaborations for Extreme-Scale Science was held December 6–7, 2011 in Gaithersburg, Maryland. The goals of the workshop were to describe the critical collaboration needs that are emerging in extreme-scale science and to identify the research gaps that must be filled for this vision of collaborative science to be met in the coming decade.

DOE manages the world's largest collection of scientific user facilities and the largest supercomputers for open science research. Collaboration is increasingly vital to the effective use of these resources and the process of scientific discovery. These facilities were used in 22 of the Nobel Prizes awarded in the past decade. What new types of scientific collaboration in the coming decade will continue this rate of scientific discovery?

This workshop focused on the expected forms of extreme-scale collaboration, including:

- Remote use of facilities, particularly use of multiple heterogeneous facilities.
- Small to large teams ranging from individuals to tens or hundreds of researchers, particularly diverse groups working together, many for the first time.
- Volumes of data precluding easy movement across the Internet, particularly data in multiple formats located at distributed facilities.
- Complex workflows involving science instruments, computing and data resources, analysis and visualization.

Enabling extreme-scale science is a major ASCR priority. The 1000-fold increase in computational capabilities expected over the next decade and a similar increase in the data volume from experimental facilities will change the way scientific discoveries are made. Multidisciplinary teams of collaborators and analysis of large, distributed datasets will become more common.

The complexity of scientific discovery is growing, not just with data volume but also from the use of multiple facilities. It is not uncommon today to have a material studied in both a neutron source and a light source, and for analysis of the combined data to yield an experimental result that is then explored further in a computational simulation. The final science outcome involves light source experts, neutron source experts, computational experts, mathematicians, and the scientists interested in this particular material.

Likewise in biology, we witness an explosion in knowledge fueled by a decade of high-throughput approaches including genomics, transcriptomics, proteomics, structural genomics, and metabolomics, which are now being combined with the ever increasing power of multimodal, multiscale spatial mapping and imaging approaches. At an ever increasing pace, the combination of “omics” and imaging will revolutionize our mechanistic understanding and therefore modeling of biological systems. Clearly such large, diverse groups of people need to be able to collaborate effectively for scientific breakthroughs to occur.

Collaboration is a broad concept. It can be as simple as a teleconference or as complex as multiple facilities simultaneously exploring the same study, with each facility in a feedback loop steering the other to accelerate scientific discovery. An example of the latter is data from an experiment at the Spallation Neutron Source (SNS) being analyzed in real time on the Jaguar supercomputer, with results fed back to SNS to modify the running experiment to maximize the value of the data relative to what is being looked for. In the future, such real-time feedback combining computing and experimental facilities will become more common, even required, for science breakthroughs.

Collaboration can also occur serially. For example, experiments at the Advanced Photon Source to understand the structure of a particular protein are planned months in advance by researchers. During the actual beam time, several team members collaborate while monitoring the experiment, and after data are collected, team members work together to analyze and understand the results.

Collaboration can also be an important part of exploiting the capabilities of DOE's leadership computers. For example, at a leadership computing facility a multidisciplinary team of computer scientists, mathematicians, and material scientists collaborate to develop a code that can run on an exascale system to simulate an advanced lithium-air battery. Additional collaboration tools are needed to guide the team to understanding the data from the simulations. In both examples, team collaboration is a key to success, but each used a different set of collaboration tools.

This workshop focused on identifying the collaboration needs of such extreme-scale science teams and key areas in which research is needed to meet these needs.

Potential Impact

The potential impact of the research described in this report is to accelerate and expand discovery in an era of exascale science. Democratization of collaborative research will give access to new realms of knowledge and innovation. The extreme-scale collaboration research will:

- Facilitate multidisciplinary research activities among small to large teams across diverse computational environments.
- Provide researchers with managed access to distributed scientific facilities and team efforts wherever they are geographically located.
- Provide effective access to and sharing of science instruments, computing resources, and scientific data and information.
- Provide protected, secure, and effective scientific collaboration technologies in a world of increasing cyber attacks.

The rest of this report is structured as follows. The next section describes a vision for the future and the role collaboration tools play in this vision, particularly how they enable investigation across experimental, simulation, and computationally driven science, bringing increased transparency to the whole cycle of creative discovery and innovation. Subsequent sections discuss unifying concepts in collaboration technologies, such as discovery and transparency; map capabilities (and the technologies that implement those technologies) to teams, data, and processing; and describe the synergistic services upon which these collaboration services depend. The technologies exhibiting gaps with respect to science needs are ranked in terms of the benefit to science of closing the gaps, their specificity to collaboration, and other properties. Finally the non-technical barriers to collaboration for extreme-scale science are discussed.

2. Vision of the Future

We foresee an environment for science in which relevant information is easy to find, easy to use, and where scientists freely share information, confident that their contributions are being recorded and will be acknowledged. “Information” in this context includes ideas, code, data, and relationships; “easy” implies no unnecessary effort or toil; and “use” may involve computation, transmission, or visualization. In the context of extreme-scale collaborative science, this vision includes seamless use of facilities and resources across the DOE computing complex, from leadership-class exascale systems, through multi-user facilities and local community facilities, to individual investigator laptops and smart devices.

We further assert that transformative progress towards this simple ideal is possible in the next ten years if DOE chooses to make the required R&D a priority.

The Office of Science’s \$5B/year budget supports 26,000 investigators at 300 academic institutions and all DOE laboratories; 27,000 researchers use its scientific user facilities. These investigators and facility users are all information workers, consuming and producing large quantities of digital data. The performance of these information workers must be of the greatest concern to DOE and the nation. Yet despite the central role of digital data in DOE research, the methods used to manage these data and to support the information and collaboration processes that underpin DOE research are often surprisingly primitive. We must do far better if we are to address 21st century energy challenges in a timely and effective manner. Improving on the current state is a worthy grand challenge for DOE that, if successfully pursued, has the potential to transform DOE research and facilities.

Collaborative science is in urgent need of new technologies that will strain the limits of our ingenuity. It is also in urgent need of the management, consensus building, and incentives that will remove utterly needless barriers to collaboration such as differing procedures for making proposals to facilities.

Current State: Primitive collaboration systems hinder discovery

Consider the disturbing extent to which the following caricature represents current reality:

The journal article (a 17th century innovation) is a primary information exchange method. The paper notebook (first century?) is widely used to document research. The workstation file system or portable drive (modern analogs of the filing cabinet, 19th century) is a primary data storage method; and due to poor indexing, much data stored on those devices is inaccessible to even its original producer.

Information sharing is clumsy at best, due to idiosyncratic schema, semantics, and data sharing conventions. Most computational results cannot be reproduced, due to their reliance on manually created and inadequately documented computational pipelines. Email (1970s), telephone (1870s), and airplane (1920s) remain the most widely used collaboration tools. Despite much progress in distributed authentication and authorization mechanisms, security concerns remain a frequent obstacle to collaboration, sometimes shutting down inter-laboratory communications for weeks at a time. And much modeling and simulation is performed using spreadsheets and proprietary packages, rather than the far more powerful methods developed by DOE and other computational researchers.

While modern tools such as Skype, wikis, and Google docs are becoming pervasive, and several scientific disciplines can claim they are very far from the caricature, it still rings alarmingly true for much of science. These primitive collaboration methods severely constrain researcher productivity and the pace of discovery and innovation.

Desired Future State:

A digital laboratory system that accelerates discovery

Let us imagine what a fully digital laboratory system, designed to accelerate discovery to 21st century speeds, would look like. In this system:

- All data, code, and documents system-wide would be accessible, discoverable, reusable, reproducible, and computable (subject, of course, to access control).
- Those same information products would be linked by a distributed knowledge base that permits automated navigation of content and connections.
- Advanced software and computational processes would be available on demand and used routinely by every researcher.
- Collaboration would occur within spaces that people want to use even when they are not collaborating.
- Intrinsic and proactive security mechanisms would encourage rather than discourage collaboration, while protecting against attacks.
- These capabilities would be as intuitive, flexible, and collaborative as the best modern consumer software. Imagine if research data and software were as easily accessible as movies from Netflix and applications from the Apple App Store.

Such a system would transform DOE research in many ways. For example, DOE researchers could rapidly design efficient, cheap, and manufacturable catalysts. Today catalyst design is largely an artisanal process. Prior data are obtained via manual literature review. Expert judgment suggests experimental strategies. Supercomputer simulations provide occasional input. And there is little or no sharing of simulation results. In contrast, a fully digital laboratory system would permit a rational design process, in which automated review of literature, experimental data, and simulation data suggest design strategies; user-guided heuristics guide exploration via large-scale ensemble computations; automated experiments evaluate promising candidates; experimental results inform further refinement; and publications are backed by full documentation of all relevant data, code, and computations. Another strategic technology that would be enabled by this system is materials research (e.g., batteries and solar cells), which is a cornerstone of the U.S. government's Advanced Manufacturing Partnership (AMP).

We could also make far more effective use of next-generation instruments at DOE light sources. Today such experiments are still largely a cottage industry. A researcher travels to the light source for an experiment. Data are collected, transferred to a USB drive, and carried back to their home institution, where reconstruction and evaluation are performed by a grad student. If problems are detected, the researcher must wait months to repeat the process. In contrast, a fully digital laboratory system would permit high-throughput experiments in which the researcher performs an experiment remotely; data are collected at 100+ MB/sec; real-time reconstruction, analysis, and quality assurance processes permit interactive control of experiments; coupled simulations are used to guide experiments; and both experimental data and a rich set of derived products are instantly available to a worldwide community.

3. Specific Opportunities to Enable Science

Emerging Data-Intensive Science

Ensuring that the scientific output from DOE light and neutron facilities keeps pace with their dramatic improvements in source flux and instrument data rates is a major challenge. In the case of synchrotron x-ray and neutron science, the range of science is so diverse and the size of collaborative teams so small that it has proved extremely difficult to marshal the resources needed to handle the data-intensive experiments now being performed every day. Nevertheless, a failure to address this challenge risks squandering much of the investment in improving the facilities themselves.

Today's operational paradigm, in which DOE experimental facilities provide the instrumentation and users analyze the data, was effective when data rates were low. However, new experimental techniques and detector technologies are enabling dramatic increases in the amount of data that can be collected from a single experiment. Techniques such as time-resolved microtomography, ptychography, x-ray photon correlation spectroscopy, and three-dimensional micro-diffraction can all produce three or more dimensional data at rates of 800 megabytes per second (MB/s) on a single instrument. The results from a successful week of experiments at a free electron laser can occupy students, postdocs, and scientists over a year inventing and experimenting with data reduction techniques to uncover new science. Similarly, instrumentation at spallation neutron sources typically includes large multi-detector arrays that collect data as real-time streams of many terabytes each day. Likewise in biology 3D volumes of cells, tissues, and microbial communities at molecular resolution collected by advanced electron microscopic approaches can reach several terabytes, and are treasure troves of information for a variety of biological disciplines.

DOE user facilities need to take a more concerned and global approach to software development and infrastructure integration at a level than has not been done before. For example, a significant fraction of the data taken at facilities today is examined not in real time but only months after data collection. This phenomenon is due mainly to the fact that the evaluation procedure is not automated. Often, small teams of scientists compete by developing better instruments as well as better analysis techniques. To build upon the research, improved algorithms should be benchmarked and inserted in common reduction pipelines. Moreover the software currently in use is a patchwork of unrelated programs that do not easily work together and that do not always address the specific local needs. Many programs currently in use have been developed by individual scientists, and are often not well designed, well maintained, or portable, and are frequently proprietary. Errors in the calibration of experimental parameters are often not quantified. There is a lack of tools to allow the user to process the entire data workflow from acquisition to publication, encompassing the ability to modify the experiment in real time after seeing the results of modeling, analysis, and visualization.

The integration of software tools and computing infrastructure involves reconstruction techniques, visualization, modeling, analysis, and theory.

Reconstruction Techniques

Inverse problems in photon science evolve as experimental techniques take advantage of many orders of magnitude higher rates. For example, an emerging technique in x-ray science called ptychography provides imaging, diffraction, and spectroscopy in one instrument. The ability to

process an entire dataset of about 10^{14} bytes ($\sim 10^2$ angles, 10^4 scan positions, 10^6 pixels, 10^2 energies) is an exciting challenge requiring developments in applied mathematics and computer science along with ongoing experimental efforts.

Even for established techniques such as tomography, active developments in applied math, statistics, and computer science have shown dramatic improvement in the solution to undersampled or noisy signals using sparse models (Figure 1) [32, 33, 45]. For imaging problems involving a known object and K unknown experimental parameters, the raw data can be parameterized (nonlinearly, in general) in a manifold signal model [8]. Adapting these models and optimizing solvers to DOE facilities would be a boon.




Figure 1. Compressive sensing (left) vs. standard reconstruction (right). Recent advances in reconstruction techniques such as tomography and MRI have been remarkable. Compressive sensing algorithms are able to overcome the Nyquist–Shannon condition and achieve higher signal-to-noise ratio with lower dose and a smaller number of views. (Image: M. Lustig, UC Berkeley)

In published papers, researchers should be encouraged to include algorithms and the full computational environment used to produce the results, such as the code, data, and whatever else is necessary for reproduction of the results. Building upon the research should be encouraged whenever possible.

One of the advantages of centralized support for collaborative data analysis would be cross-discipline fertilization. Advances in imaging and image analysis developed in the context of cosmology or material sciences would great benefits the analysis of biological multidimensional datasets, as the underlying challenges are similar, and hence existing approaches can be readily adapted.

Visualization

A pressing issue in scientific software is online visualization and processing of data as they are collected at DOE user facilities. Many tools exist for these tasks, but the pervasive lack of standard tools at the APS and SNS has been a persistent bottleneck, making it difficult for users to see changes in their data and identify potential problems with the data acquisition. Open source visualization software should be developed both in addition to and in support of commercial visualization products, such as IDL, IGOR, and MATLAB. Of particular interest would be a Web-based immersive visualization tool that would allow easy and intuitive navigation through a 3D volume, with the ease of a virtual reality game. Visualization ideally would be combined with user-guided as well as template-guided automated feature extraction, real-time annotation, and quantitative geometrical analysis.

Modeling, Analysis, and Theory

The traditional operational workflow, in which the facilities provide the instrumentation and the users analyze the data, was effective when data rates were reasonably low (Figure 2). A preferred approach involves real-time feedback between visualization, modeling, and experiment parameters, but there is currently no way to realize such a workflow in data-intensive experiments.

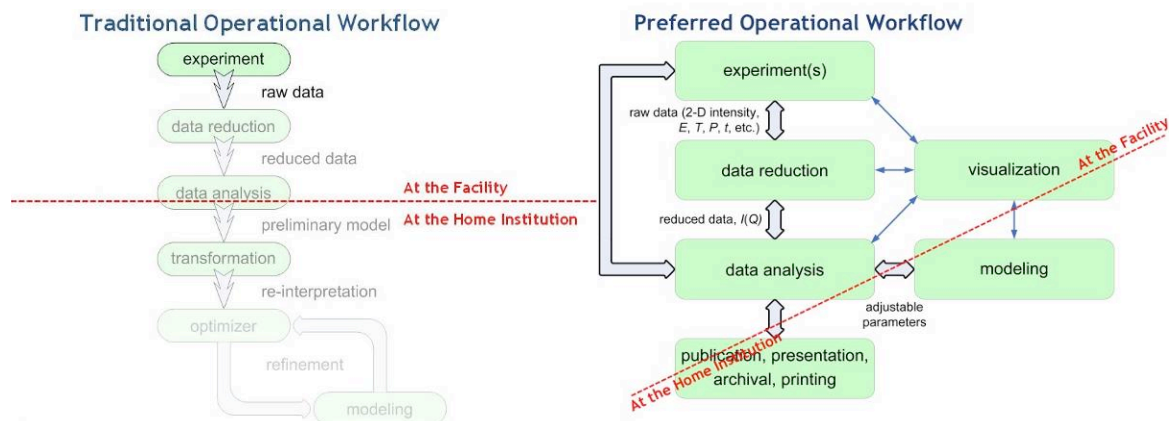


Figure 2. The traditional operational workflow at DOE large user facilities (left) includes data collection, data reduction, and the first steps of data analysis. Subsequently the data are handed to the users, who perform further analysis and modeling at their home institution. A preferred operational workflow (right) should include immediate data analysis performed at the facility to allow real-time feedback and optimization of the instrument parameters. Furthermore, the tasks that should be performed at the facility and at the home institution are better represented by the tilted dotted line.

What is needed is an advanced computational framework that allows the scalable integration of the data reduction procedures developed by the facility with analysis and modeling developed by the users, all made available to the user community so that there are no substantial barriers-to-entry. The technical challenges in producing the necessary infrastructure (high performance networks, accessible and searchable data repositories, libraries of advanced algorithms, software maintenance and distribution, adequate documentation) are sufficient to require a much greater investment in data handling and treatment than any individual facility has been willing to provide so far.

Simulation-Driven Science

Increasingly, the lines separating simulation, theory, and observations, the “three-legged stool,” which defined the role of simulation science since the 1970s, are becoming blurred. Scientific collaborations among the sub-disciplines are much closer today in nearly every field that employs simulation and modeling to a large extent. Often it is the simulation enterprise that provides the “big picture” integrated view of large, nonlinear problems, while measurements and theory are still needed to constrain the design and test the behavior of simulation codes. The rapid evolution of information, networking, and computer technology has transformed both modeling and the data sciences, thereby enabling highly interactive, integrated scientific collaborations to attack heretofore unsolvable problems. Uncertainty quantification (UQ) methods are now employed frequently in both data analysis and simulation, leading to more formal and rigorous comparisons of observations to simulation results, adding complexity that stresses the cyber-infrastructure supporting collaborations.

Although each application area will have differences, the next generation of scientific collaborations will require some common capabilities:

1. Centralized data catalogs and tightly federated data repositories containing increasingly diverse, dynamically generated, self-described data (observed and simulated) and well-defined abstractions for system services that will enable users to discover, access, and exploit the wealth of information and capture data provenance.
2. A new generation of flexible and extensible open-source software tools (models, analysis codes, and automated workflow drivers) and frameworks with well-defined APIs that will preserve the intellectual investment in their development and reduce the problem of rapid obsolescence as hardware systems (e.g., simulation engines, storage and network systems) evolve.

Without context, these requirements lack sufficient definition to motivate the research activities to design the next generation of collaboration technologies. Below are presented three examples that provide both context and motivation.

Combustion

Aggressive national goals aim to reduce greenhouse gas emissions by 80 percent by 2050 and petroleum consumption by 25 percent by 2020. Future engines will depend on a combustion modeling capability that combines high-fidelity direct numerical simulation, in situ analytics, and embedded uncertainty quantification. Models will be designed to differentiate alternative fuel effects on turbulence-chemistry interactions in high-pressure regimes of practical combustors. Their simulations will provide fundamental insight and validation data to guide development of predictive models used by industry to shorten the design cycle, promote economic competitiveness, reduce foreign oil dependence, and promote environmental stewardship.

Combustion models will rely on direct numerical turbulence simulation and therefore require adaptive mesh refinement to meet spatial resolution requirements. They will have simulation fidelity across a range of flow regimes from low Mach number to fully compressible formulations. Interpretation of the results requires in situ data analytics, because the models simulate data at rates that are too high for deferred analysis, and there is increasing emphasis on “steering,” or interactively altering the simulation design as a run progresses. UQ methods will ascertain the

impact of uncertainty of chemical parameters on predictive capability through quantitative comparisons with experimental data. This simulation paradigm (Figure 3) requires balancing simulation and analysis and is dependent on the definition of an optimal data structure.

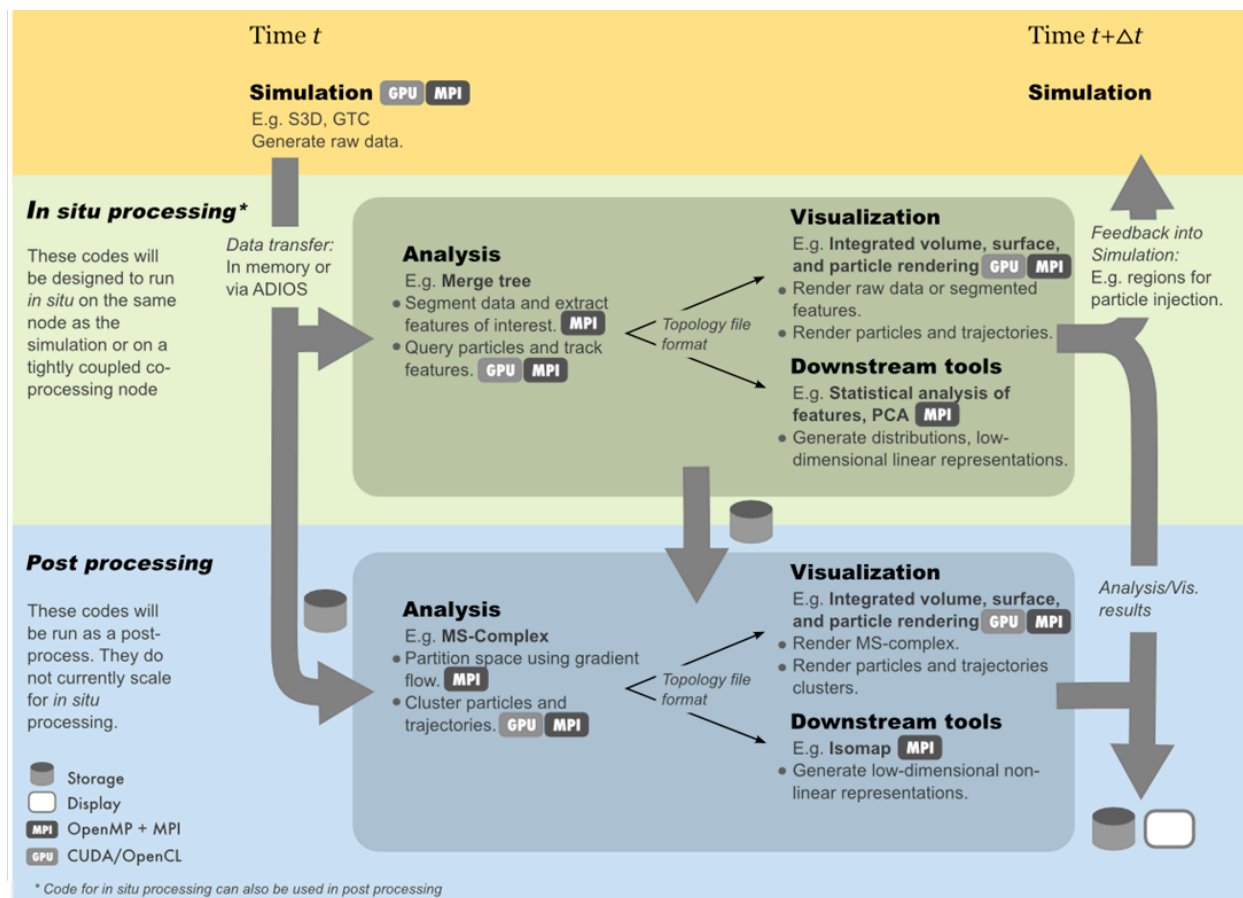


Figure 3. Schematic view of a combustion simulation with in situ analysis and visualization workflow. (Image: J. Chen, Combustion Research Facility, SNL)

Fusion Energy Science

One of the central plasma physics problems on the road to creating a working fusion power plant is understanding, predicting, and controlling instabilities caused by unavoidable plasma inhomogeneities. One consequence is the occurrence of turbulent fluctuations (microturbulence) which can significantly increase the transport rate of heat, particles, and momentum across the confining magnetic field — an overall effect that severely limits the energy confinement time for a given machine size and therefore the performance and economy of a tokamak device. Accelerated progress on this critical issue is especially important for ITER, because the size and cost of a fusion reactor is determined by the balance between such loss processes and the self-heating rates of the actual fusion reactions.

Fusion energy research has historically been a leading high performance computing (HPC) application domain, and has made excellent progress in developing advanced codes for which computer runtime and problem size scale well with the number of processors on massively parallel machines. The application of extreme-scale computations (spanning the range from multi-petaflops

to exaflops) of key plasma physics and materials science processes in fusion devices has been essential to advance the science. A good example is the effective usage of the full power of multi-teraflop to petaflop systems to produce three-dimensional, general geometry, nonlinear particle simulations which have accelerated progress in understanding the nature of plasma turbulence in fusion-grade high temperature plasmas (Figure 4). These calculations, which typically utilize multi-billions of particles for thousands of time-steps, would not have been possible without access to such powerful modern supercomputers together with advanced diagnostic and visualization capabilities to help interpret the results.




Figure 4. Petascale production-run simulations accounting for fully global 3D geometric complexity spanning micro and meso scales carried out on ORNL's Jaguar system with the XGC-1 PIC code, typically using 24M CPU hours, i.e., 100,000 cores \times 240 hours. (Image: C. S. Chang, SciDAC Center for Plasma Edge Simulations, PPPL)

Transformational changes in simulation, data analysis and collaboration technologies are necessary in order to move in a timely manner to producing HPC simulations with the highest possible physics fidelity. With the unprecedented resolution in a multi-dimensional phase-space enabled by access to HPC platforms at the petascale and beyond to the exascale, advanced kinetic simulation capabilities are expected to have direct relevance to existing experimental devices as well as to ITER. Validation of such advanced codes requires dealing with increasingly large databases, especially once ITER is operative. This will necessitate engagement of the network and the distributed systems communities to address the grand challenge issues inhibiting the development, deployment, and operation of scalable and secure extreme-scale scientific collaborative environments that involve multi-disciplinary, multi-facility, and/or distributed research teams in the next decade.

Climate Prediction

Climate simulation and prediction have entered a new and more demanding era with the recognition that global temperatures are increasing and will continue to increase at an accelerating rate from the accumulation of greenhouse gases in the atmosphere. Current global climate models reproduce historical global warming trends and provide insight into future climate change that will result from the continued accumulation of atmospheric greenhouse gases. Nevertheless, global models lack the fidelity to accurately predict, with quantified uncertainty, how climate will change on regional scales and the how that change will emerge in the next few decades. Developing the scientific and computational capacity to produce these predictions will usher in the era of seamless prediction across weather and climate temporal and spatial scales.

The DOE Climate Science for a Sustainable Energy Future (CSSEF) project endeavors to combine simulation, theory, and observations to build a climate simulation and prediction model two-generations more advanced than the current Community Earth System Model (CESM1). CSSEF will address a critical and relatively straightforward objective — to accelerate the incorporation of new knowledge, including process data and observations, into climate models and to develop new

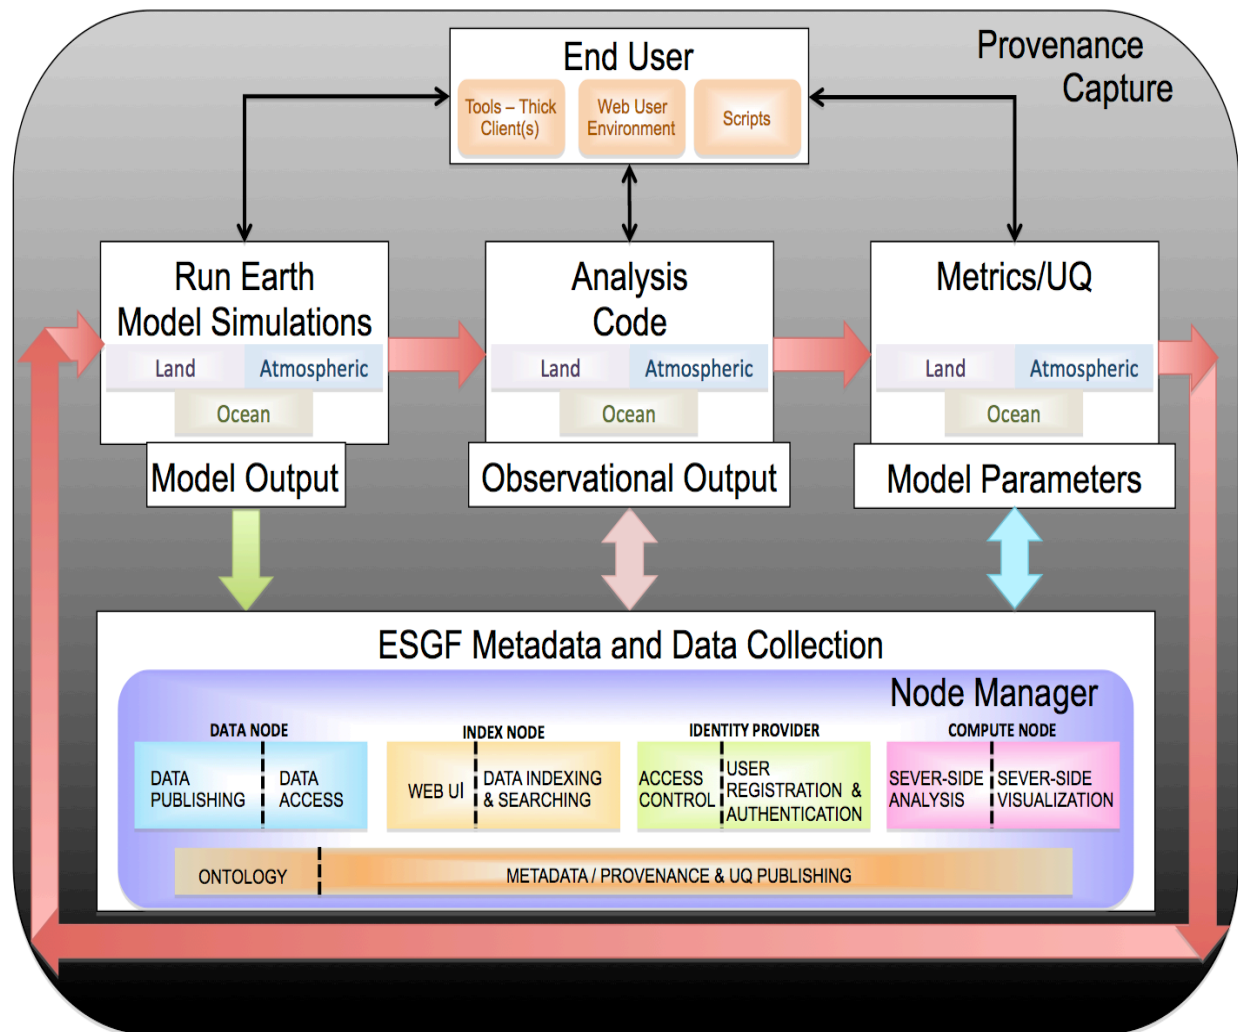


Figure 5: The CSSEF testbed architecture will incorporate an underlying ontology service to facilitate search and integration across formats. (Image: D. Williams, LLNL)

methods for rapid validation of improved models. A crosscutting objective of CSSEF is to develop novel approaches to exploit computing at the level of many tens of petaflops in climate models.

CSSEF will undertake several unique and potentially transformative research directions, including:

- the capability to thoroughly test and understand the uncertainties in the overall model and its components as they are being developed
- major scientific advances in the components that will achieve greater fidelity in modeling feedbacks in the climate system
- development of model evaluation procedures that allow the rapid ingest of observational data for model and component evaluation
- flexible dynamical cores that enable fine-scale simulations
- early adaptation of the model algorithms and code to the next generation of computers.

Figure 5 is a high-level conceptual view of the CSSEF testbed architecture and workflow. Working for all model components (atmosphere, land, and ocean), provenance capture is pervasive throughout the testbed, capturing and saving all user actions. The red arrow lines show the baseline ensemble loop, in which model simulations are conducted using a variety of input parameters generated by metrics and UQ ensemble drivers. At any stage, data can be collected and stored in the ESGF distributed archive. The black arrow lines show how access to the testbed is attained through desktop clients, web browsers, or scripts.

The Materials Genome

Materials discovery and innovation are crucial to a sustainable energy future. Applications areas such as photovoltaics, permanent magnets, cement, thermoelectrics, energy storage and generation, catalysis, etc. would all benefit from new materials with targeted optimized properties. The current materials innovation timeline from lab to commercialization is 15 years. Traditional trial-and-error or one-at-a-time experimental approaches are slow and lack systematic reliability.

Simulation of chemical and physical materials properties has sped the discovery of new materials. A revolution in how this process works is unfolding through the Materials Genome Initiative (materialsproject.org), which lends systematics and organization to the way simulation science is applied to materials. It is also inherently a collaborative approach.

The materials genome approach moves simulation from being driven by inspired guesses or expert knowledge about materials toward a fully systematic workflow that produces a database of all known materials and their properties. It is an approach to materials discovery which is inherently collaborative through sharing. The Materials Project applies density functional theory (DFT) to all the crystal structures in the International Crystal Structure Database (ICSD), a data-centric platform for rational materials design driven by the search for desired properties. The collaborative nature of this approach enters into the use, annotation, and curation of a shared, web-based database of materials. It replaces simulation by isolated researchers with an organized, shared resource for materials discovery.

Computational materials science has reached the point of accuracy that many properties can be reliably computed. This means that we can screen materials *in silico* for suitable applications, without having to make them in the lab. This saves an enormous amount of resources and time. The Materials Project has computed the ground state of ~20,000 compounds which are available for search and other applications at the materialsproject.org science gateway. Systematic collaboration

and search then replace individually inspired computations communicated by publication and a common understanding of the entire space of materials is made possible.

Data-Intensive Science

Collaborative science at large scales is sometimes characterized by very large datasets, collected or generated within a community, that are unusable by any individual scientist without collaborative tools for data discovery and access. An example of such data-intensive science can be found in collider-based High Energy Physics and Nuclear Physics (HEP/NP) research such as that being carried out at the Large Hadron Collider (LHC). The HEP/NP communities form large experiment-specific collaborations consisting of thousands of scientists and engineers to build detectors, operate the experiment, and process the data. The collaborations enable development of computing tools and infrastructures for their data-intensive research and have produced large distributed computing facilities which can illustrate future needs for collaborative science at extreme scales.

Collider-based HEP/NP experiments produce large amounts of data at facilities where compute resources are insufficient to meet the storage and processing requirements of the scientific mission. These projects routinely move data to remote centers for storage and analysis, introducing reliance on high performance networks and on experiment-specific data management tools to track data origins and ensure their authenticity. This characteristic data flow is occurring more frequently in science, from large telescopes in observational astronomy to climate simulations run at large data centers, and, as a growing trend, will expand the use cases for distributed workflows and data provenance mechanisms. The networks are the fabric upon which distributed workflows operate.¹ At extreme scales, processing demands will require network monitoring capabilities that enable new models for data placement, discovery, and processing that dynamically link science applications with compute and storage systems to maximize available resources for collaborative science.

Large collaborations, as in HEP/NP research, provide their members with tools to allow groups of scientists, often at different institutions, to work together on specific topics. Some tools target basic needs that arise in large associations, such as in membership and communication. Other tools attempt to address the complexities of processing large datasets and form the infrastructure in which the scientists work to mine data for information. In distributed environments, there remains a strong need for user-level tools that can create dynamic workspaces in which researchers can encapsulate all information in their work, greatly lowering the barrier to collaboration between scientists. A measure of the effectiveness of such collaboration tools is the extent to which member scientists work within the confines of the infrastructure rather than extract subsets of data for private use, a challenge that will grow at more extreme scales.

To confront this challenge, new infrastructures for data-intensive analysis need to support seamless access to data by all interested scientists regardless of their discipline and location, providing data at the granularity required by different analyses. Collaborative tools will interface with dataset definitions that are interpretable by scientists yet interoperable at the format and

¹ Here we assume (1) that the entire data path from end to end will be built to accommodate high data-rate flows, from the application/data source, through the network software stack on the host, through the NIC, the active devices in the LAN and WAN, and then up through the same environment at the receiving end; and (2) that network services match the application requirements for the network, such as minimum guaranteed bandwidth and engineered fail-over paths.

semantic level for cataloging and data access services. New processing models will need to invoke workload management systems that integrate information about data location, site capabilities, and network topologies to streamline all processing steps, constructing that workspace environment for large and distributed datasets that hides the at-scale complexity to satisfy the scientist's need for simple, direct access to the data.

Other unique data-intensive challenges in collaborative research arise with real-time data processing needed for immediate feedback to the remote instrument or scientist to guide adjustments to equipment, additional processing, or other actions. One example from the fusion community is the remote operation of tokamak devices for magnetic confinement, which requires data analyses and complex calculations be done in real time to maintain stable operations. Parallels can be found in other sciences that include remote operations of large detector systems with complex run-time requirements in which the time dimension makes these transactional analyses data-intensive.

HEP/NP experiments have been able to leverage the large collaboration sizes in developing complex data-intensive infrastructures. However, even this community recognized that partnerships between organizations were crucial for coordinating resources that span facilities, institutions, and even countries. Consortia such as the Open Science Grid, a participating organization in the Worldwide LHC Computing Grid collaboration, are an important layer for capturing requirements, providing common services, and disseminating information about best practices and technological solutions. Future research supporting scientific collaborations for extreme scale science should recognize the utility of this layer as a point of contact for emerging communities to engage in the development and utilization of collaborative toolsets, limiting one-off solutions without stifling innovation.

4. Mapping to Technology and Computer Science

Unifying Abstractions

Discussion during the workshop led to the identification of four simple but powerful unifying abstractions that, if supported by broadly adopted solutions, will accelerate progress towards the dream. These are:

- **Discovery:** All resources are easy to find and understand.
("I cannot use resources that I do not know exist!")
- **Centrality:** Standardized services reduce costs and encourage commonality.
("Don't make me install software or learn arcane details to collaborate!")
- **Portability:** Resources are widely usable in a transparent fashion.
("If I can't use your data or software, it isn't science!")
- **Connectivity:** Where information came from, and what other information it relates to, are easy to find.
("No information exists in isolation!")

Discovery: Researchers tend to be highly creative when it comes to using many resources (data, software, instruments, human expertise, computers) in their science. Yet as a result of the explosion in interesting resources the fraction known to individual investigators is surely declining. To combat this trend, we must ensure that (subject to security and privacy concerns) every resource is named, described, and discoverable, via appropriate descriptions and registries; that each name provides enough information to access the resource; and that each description provides sufficient information to permit a motivated and authorized individual to access and use that resource. We envision, for example, that every dataset, software package, technical article, and investigator within the DOE system is (subject, again, to security and privacy concerns) described and discoverable, to some degree at least. In addition to natural language search (the *sine qua non* of easy discovery), simple programmatic APIs should make all resources available to or from a science team readily available. By adding convenient mechanisms for associating user-generated annotations with resources a range of interesting new collaboration and research methodologies become feasible.

Centrality: Scientists need services, including access to facilities, that are as consistent and uniform as possible. Services should appear to be centrally managed and supported, even if the implementation is more distributed. In today's extreme-scale projects, the first steps are often agreeing on, installing, and configuring a Web server, Wiki server, email lists, source code management system, several other complex software systems, plus learning the arcana associated with the DOE facilities to be exploited. But this is 2012. Software as a service (SaaS) works. We should no longer need to install software or learn about parochial systems in order to collaborate. Appropriate centralized and standardized services can reduce dramatically the time and cost required to establish, participate in, and manage collaborative projects: ideally, establishing a "virtual private laboratory" to support a new collaborative (or indeed, private) project should require just a few mouse clicks. Such services can also encourage commonality, further reducing barriers to collaborative work. A major theme for future DOE collaboratory programs must be identifying the centralized services required for extreme-scale science — and then either acquiring or creating those services.

Portability: Science is increasingly dependent on the creation and use of data and software. Collaborative science is crucially dependent on the *sharing* of that data and software. Indeed, if a piece of data or software is not usable by others, then there is a strong case to be made that whatever process is being engaged in, it is not science because it is not reproducible [41]. Mechanisms that make knowledge encapsulated in data and software “portable” (i.e., reusable by others) will have a multiplicative effect on collaboration effectiveness. Portability for software can mean ensuring that it runs on many platforms and at many locations — or, in some cases, providing access to a single instance via service interfaces. Portability for data can encompass access protocols, representations, and semantics.

Connectivity: Next-generation scientific advances depend on analyses of data generated from simulations and experiments that span time, groups, and science domains. To enable these analyses, there needs to be a deep and pervasive ability to create scientifically meaningful connections between independent datasets. At extreme scales of data volume and diversity, this process requires automated data reductions and annotations. Safeguards against invalid or conflicting data must also be integrated without having to normalize large volumes of data to a standard format. Derived results depend on previous results, and so on back to the original measurement or computation; thus, validating, verifying, and reusing a given result depends on creating and maintaining this complex set of connections throughout the data lifecycle.

Overview of Technologies

The abstractions must guide the research, development, and implementations that deliver coherent and cost-effective technologies that enable or enhance collaboration. We decompose the technologies themselves into three areas:

1. **Technologies for Teams**, addressing the human aspects of collaboration
2. **Technologies for Data**, addressing the needs specific to sharing data
3. **Technologies for Processing**, addressing the sharing and co-scheduling of processing.

For completeness, we also identify underlying Synergistic Services that provide essential support for science.

Figure 6 shows the decomposition of technical topics used in this report. The figure makes a color-coded attempt to distinguish between technologies specific to collaboration and technologies that are essential for science, but less directly driven by the needs of collaboration

Technologies for Teams

Studies show that “research is increasingly done in teams across virtually all fields” and that “teams now also produce the exceptionally high impact research” [52]. Much of this work is digitally mediated, whether because team members are geographically dispersed, the resources used in their research are remote, or research products involve digital artifacts. Thus better technology can play a crucial role in making team science more effective. In this section we focus on the technologies needed to serve the interpersonal interactions that are central to collaborative science.

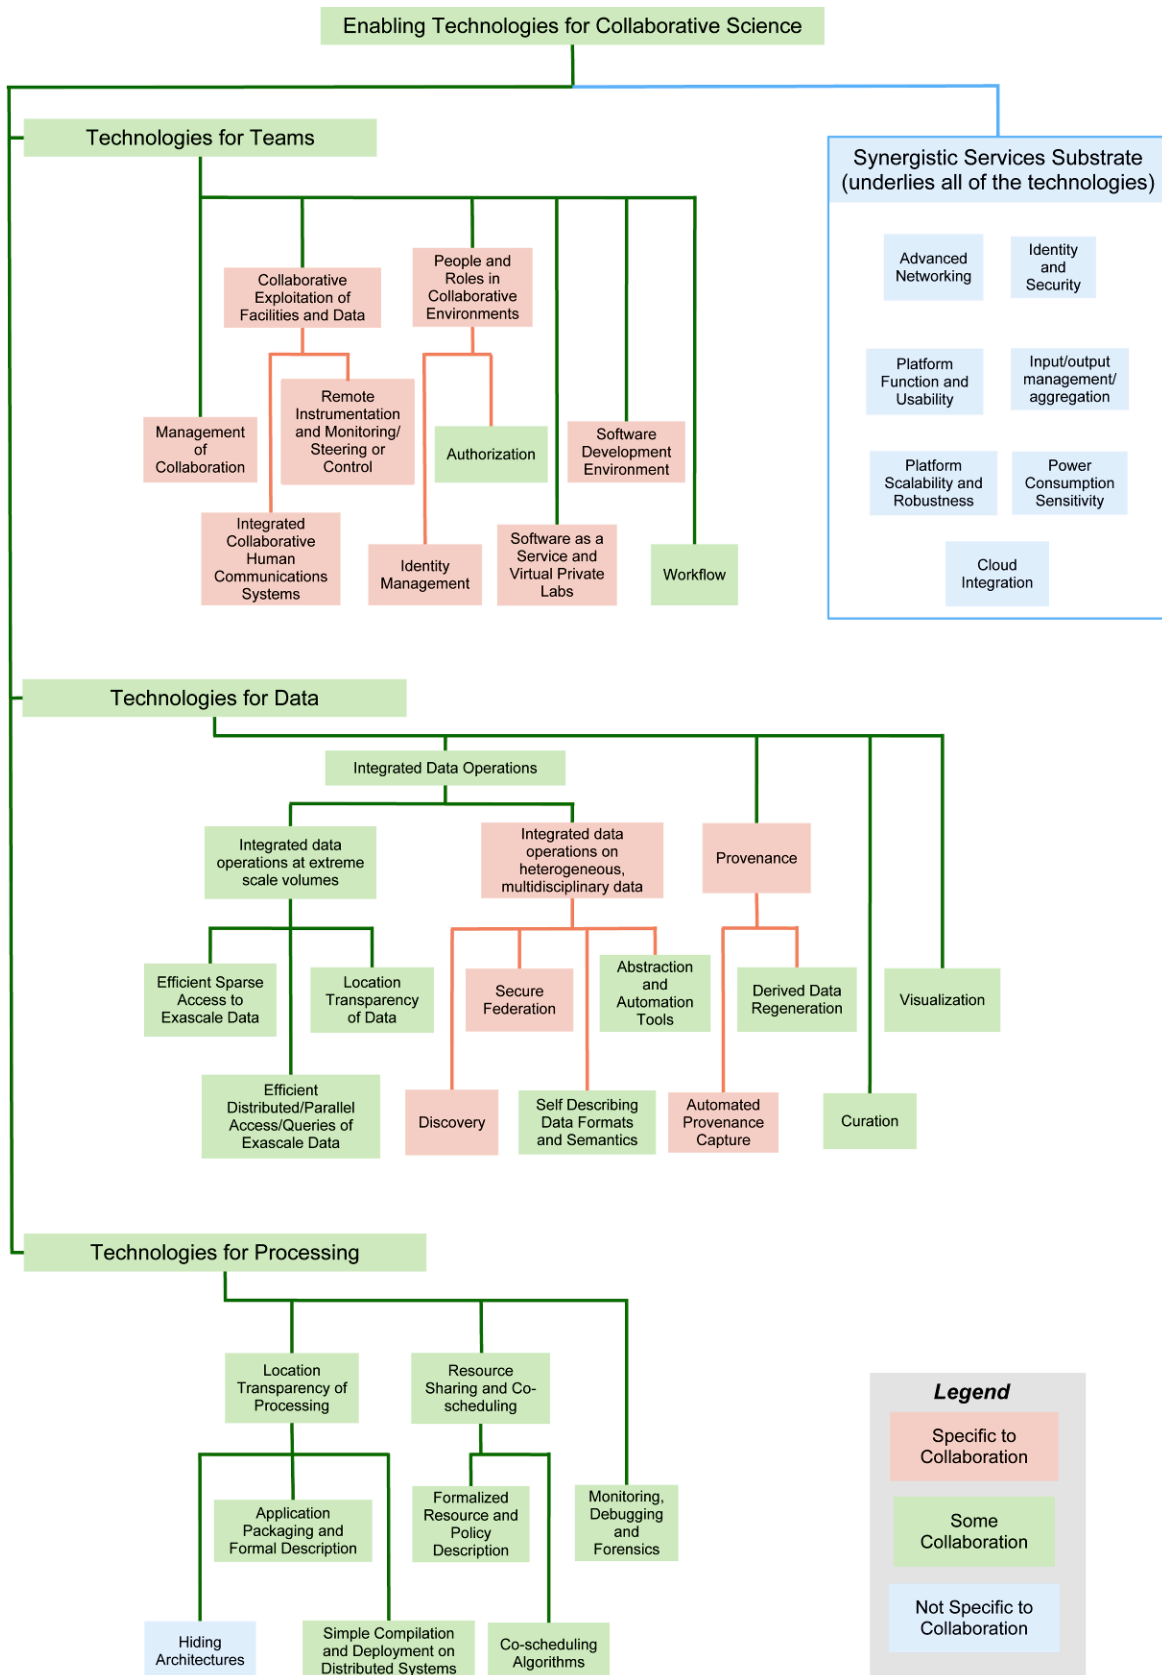


Figure 6. The technical capabilities and synergistic services examined in this report.

Management of Collaboration

The establishment, configuration, functioning, evolution, and management of collaborative science projects involves a wide variety of complex and time-consuming tasks, few of which are adequately supported by information technology. Locating required expertise and resources, communicating with participants, developing and testing new software, scheduling meetings, project planning, managing resource allocations, budgeting, progress tracking and reporting, logging and audit — these are just a few of the many tasks that frequently consume far too much time in collaborative projects. (One workshop participant described how his mid-sized scientific collaboration operated more than 15 different software services, each requiring a different user credential.)

We discussed in the subsection “Unifying Abstractions” the important role that centralized services can play in accelerating a wide range of collaborative tasks, by providing easy access to powerful common services without requiring that team participants install, configure, and operate software. Commodity services, like the Google Docs used to draft this workshop report, will surely play an important role in science. But a strong message from this workshop is that effective team science for the 21st century will require powerful new collaboration services that address the specific challenges of extreme-scale science. With such services, DOE can aspire to accelerate discovery and innovation in collaborative groups. Without them, increased funding on collaboration tools seems likely simply to add to an already large number of hard-to-install-and-use, non-interoperating software packages. This message applies to each of the subsections in this part of the report; it is expanded upon, in particular, in the subsection on Virtual Private Laboratories.

Collaborative Exploitation of Facilities

Individual users of DOE facilities face major new challenges due to greatly expanded data volumes and the growing sophistication of data analysis software. These challenges become yet greater when, as is increasingly often the case, the design, execution, and analysis of an experiment is a team effort. Large teams that are geographically dispersed face substantial challenges when attempting to conduct their science in a truly distributed fashion. For those teams whose focus is a large central experimental facility, the ability to bring remote collaborators into the scientific debate is particularly challenging. It is not unusual that an experimental session involves online hardware and software control adjustments. These adjustments are debated and discussed among the distributed team. This mode of operation places a premium on rapid data visualization and analysis to enable understanding in near real time by a geographically dispersed team.

Integrated Collaborative Human Communications Systems

Effective remote participation on the scale envisioned requires that it be easy to provision working environments in which personnel at different sites can engage in high-human-bandwidth dialogues enriched by up-to-date data. The diverse communication channels currently used in research — phone, audio, video, email, messaging, and data — need to be integrated into a common framework that supports both *ad hoc* and structured communication. Streaming of video and audio from remote control rooms needs to be integrated with *ad hoc* voice and video channels. Advanced interactive directory services (see subsection “Discovery,” p. 27) are needed so that people and data streams can be identified, located, scheduled, and accessed.

Since the ability to share complex visualizations and applications among remote participants is a necessary adjunct to scientific dialogues, interpersonal communications media needs to be enriched via remote sharing of displays and applications. Distributed shared displays will be an

important element supporting remote control rooms; these displays need to operate in the high latency environment inherent to long-distance communication. Important components of this infrastructure include an easy-to-use and easy-to-manage user authentication and authorization framework, and global directory and naming services.

Remote Instrumentation and Monitoring/Remote Steering or Control

It has not been necessary since the 1980s to camp out in the machine room to use a supercomputer. However, physical presence is still the norm for users of the majority of DOE's extreme-scale experimental facilities, and thousands of scientists travel every year for this purpose. Remote instrumentation and monitoring are increasingly used to make it unnecessary for most members of a collaborative team to travel to the experimental facility: one or a few team members can travel to a facility, while others monitor progress from home. However, remote instrumentation and monitoring mechanisms tend to be implemented in widely varying ways across facilities, hindering researcher mobility.

Remote steering and control, if implemented correctly, can avoid the need for travel to facilities entirely. In so doing, we can vastly reduce idea-to-data times for much of collaborative science. However, the realization of true remote steering and control places stringent demands on authentication, role-based authorization, and facility protection and collaboration systems.

People and Roles in Collaborative Environments

Few team science activities can take place without effective mechanisms for establishing user identity and for expressing and managing policies concerning access to resources. These mechanisms must be able to co-exist with diverse site security policies and integrate cleanly with popular commodity collaboration technologies. (I probably won't be able to use my Facebook identity to access a leadership class supercomputer—but I may well want to use Google Docs to develop collaborative documents.) As team science grows in scale and importance, and attackers simultaneously grow more sophisticated, increasing emphasis must also be placed on protecting not just individual resources but the distributed resource collections and services used by scientific collaborations.

Identity Management

The secure establishment of identity for different purposes and the management of the multiple identities that any individual inevitably possesses are fundamentally difficult problems. While these problems are by no means unique to extreme-scale science, they are particularly challenging in collaborative scientific environments due to the many different individuals and institutions that may be involved in a collaboration, the wide variety of resources used, and the often high degree of trust required to achieve remote access. A single project may involve participants from multiple institutions and countries, who then need to share resources ranging from collaborative documents and email to supercomputers and petabyte storage systems. A typical bad dream (if not nightmare) is setting up a workshop agenda to which few participants can upload their talks since the majority of their (many) identities are unknown to the agenda management software. And each individual may be involved in multiple such projects and meetings at any one time, and in many such projects over an extended period. Given the complexity of DOE science, progress will require both research and effort applied to development and deployment. Solving this problem may be at least as much an issue of leadership and management as it is of technology.

Authorization

Once a person's identity has been determined, various parties need to determine what the person should be allowed to do. This decision can depend on many factors, including the nature of the credentials provided by the user, the person's role in the activity or project, and the details of the intended task. In a single line-managed organization, managing and auditing such authorizations are continuing challenges. In collaborations for extreme-scale science, it is a major impediment to progress. As with identity management, progress requires a combination of new ideas, better development and deployment of known methods, and leadership and organization.

Virtual Private Laboratories

Despite considerable advances in networking, security, and data movement (three areas in which DOE innovations have been influential), establishing and securely and efficiently operating resource sharing and teaming relationships remains inordinately difficult. It is not uncommon for the apparently simple act of sharing a single dataset to involve many email exchanges and manual manipulation of accounts, permissions, and privileges — with results that are frequently inefficient, nonintuitive, and even insecure.

A major reason for such difficulties is the continuing bespoke nature of much of the technology used to enable collaboration. Requiring researchers to install software, manage research data, configure collaboration tools, or manage permissions is like requiring them to install and run their own telephone exchange to make a phone call — a telephone exchange that will not interoperate with those installed by other teams.

Workshop participants concluded that an answer to this problem may be to get individual researchers, research teams, and laboratories out of the business of installing and operating the software and other information technology (IT) used in their research. Instead, responsibility for much of that IT should be transferred to software as a service (SaaS) providers. In industry, this approach has been tremendously successful: many small and medium businesses today obtain essentially all of their IT functions (payroll, accounting, Web presence, email, customer relationship management, etc.) from business SaaS providers. Equivalent services for research would allow an individual or team (a small or medium research laboratory) to outsource the tasks of establishing and operating the collaboration, data management, security, and other tools required for an effective resource sharing or teaming relationship. Research SaaS providers would provide on-demand access to powerful virtual laboratories that individually or collectively deliver all of the IT required for effective research.

The successful realization of this vision of outsourced IT will require answers to a range of challenging research questions. What are the critical processes that underpin modern research — the equivalents for small and medium research teams of payroll, accounting, and customer relationship management for small and medium businesses? What are the foundational elements on which can be built robust, secure, and scalable research data management and collaboration solutions? How can these elements be integrated with campus and national cyberinfrastructure systems, such as supercomputer centers? How do we scale solutions to massive data, large teams, and high-throughput processes? How do we integrate the audit and provenance mechanisms required for reproducible research, without creating onerous requirements on investigators? What user experience (UX) elements are important in research? (Companies such as Netflix, Google, Apple, and Amazon have pioneered approaches to consumer UX that have proved transformative in their usability. Will similar methods work for science?) How does outsourced IT change the security equation? (SaaS can enhance security by professionalizing its operations. But presumably it also

introduce new risks. How can those risks be managed?) It will also be important to study the economics and sociology of DOE science. What factors may hinder or encourage adoption? How will we sustain such outsourced services, and who will provide them? What may be the unexpected consequences?

Software Development Environment

Algorithm and software development is an integral part of the intellectual process of collaborative science. Collaborations have been using collaborative tools like the CVS version control system [20] for decades. Although existing tools could certainly be better, one of the greatest challenges for the future is the automated tracking and recognition of the authorship of algorithms and software ideas that will accelerate extreme-scale science.

Workflow

Workflow is concerned with the specification, coordination, and documentation of actions performed on different resources, by the same or different people, over a variety of time scales. While the need for workflow technologies is not peculiar to DOE science, extreme-scale science introduces unique requirements associated with, for example, large data volumes (e.g., 100+ MByte/sec flows from light source beamlines), specialized audit and approval processes (e.g., in ITER, collective control of an experimental shift), and real-time monitoring and control. Workshop participants identified the end-to-end problem (i.e., workflows that capture entire science processes, efficiently and correctly) as particularly important for DOE to address. The ability to discover (Discovery), outsource (Centrality), and reuse (Portability) workflows are vitally important.

Technologies for Data

Data are central to scientific computing. The challenges unique to extreme scale science are the scale, heterogeneity, and distributed organization of the data used in an analysis. Consequently, data technologies must be scalable in several different dimensions: the *complexity* of individual datasets, the *volume* of datasets, and the *diversity* of data types and domains. Data technologies that enable the ultimate goal of collaborative analysis and interpretation of the data can be divided into several categories: integrated data operations, provenance, curation, and visualization. Of these technologies, the most specific to extreme-scale collaboration are integrated data operations and provenance.

Integrated Data Operations

Integrated data operations combines the unifying concepts of data *discovery*, *portability*, and *interconnectivity*: the datasets must be found, understood, and related to each other. Particularly at extreme volumes, technologies for data must take a broad view of “the data” that includes derived in situ and other reduced views of the original raw dataset. Collaborative technologies must address this complexity through automation and/or the use of underlying tools that can handle semi-structured data, such as schemaless data stores.

Integrated Data Operations at Extreme Volumes

At extreme volumes, access to even a single dataset will require new technologies that provide the following abilities:

Efficient sparse access. Sparse access allows work on just the relevant subset of the data. The user must be able to select a subset by combinations of attributes (which requires detailed indexing). This is complicated by the fact that extreme-volume datasets are distributed across file systems and physical locations. Where broadly used summaries are expensive to compute (e.g., seasonal averages for the CMIP5 climate model dataset), data access mechanisms need to consider the tradeoffs between storage and re-computation.

Distributed and parallel access. These technologies enable distributed query processing and result-set retrieval from local clusters as well as fully distributed filesystems and databases within a single domain. This functionality may require technologies for parallelization as well as technologies to handle the vagaries of MANs and WANs. For the latter, understanding and adaptation to network performance is crucial.

Location transparency. Location transparency hides details of distributed access from either the user or other layers of automation. This is a cross-cutting technology that may exist at several layers and places in the data pipeline and associated storage stack.

Integrated Data Operations on Heterogeneous Multidisciplinary Data

Another challenge is the extreme heterogeneity and complexity of the various data needed at a given stage in the data pipeline, for example during data generation, storage, (re)distribution and analysis. The major technologies needed to cope with this heterogeneity are:

Discovery. The dream of extreme-scale discovery requires normalization of semantics and vocabularies. In more complex scenarios, ranges of values will have to be understood in multiple unit systems or in terms of the discrete values they span (e.g., atoms with two or more valence electrons). A number of scientific disciplines have or are developing standardized semantic descriptions of their common data formats that the data technologies must be able to exploit. These semantic descriptions support integration of data both within and between different science communities. Global directory and naming services are a technical prerequisite for full exploitation of normalized semantics.

Secure federation. Secure federation involves virtual gathering, replication, and redistribution of datasets. Extreme-scale collaboration inherently requires new technologies for distributed data access across both local- and wide-area networks. Most data-intensive collaborations host data in distributed, heterogeneous facilities requiring new methods for connecting (or federating) these resources. Computing resources located at one site require data resources at another and immediately confront a range of complex issues which are highly context-sensitive: staging-in files or streams over the network (latency); caching full datasets, files, or even sub-file fragments (caching capabilities at the computing site); employing a uniform name space across the collaboratory and a single point of entry (accessibility); creating and optimizing cache and read-ahead buffers in the client when needed (CPU and IO performance); and a host of data management issues for temporal or custodial datasets. These technologies must integrate the trust models and technologies (e.g., PKI, Shibboleth) used within a collaboration with the policies at the data sites. To

move data efficiently, these technologies should also be integrated with advanced capabilities of high performance data transfer tools.

Self-describing data formats and semantics. These are important tools for the automation of data integration in the data pipeline. The standardization of domain semantics and formats, while powerful and useful, is both difficult and inevitably somewhat out of date. Self-describing formats, which use general-purpose technologies to handle many domain-specific formats and semantics, fill this large and ever-present gap for scientific collaboration. Examples of self-describing formats are XML and HDF5. The issues around data semantics are trickier to automate, but collaboration technologies should leverage the progress made in biological and social sciences domains with semantic technologies such as OWL and controlled vocabularies.

Abstraction and automation tools. These can use data semantics to transform datasets from one format to another, or perform on-the-fly transformations. In turn, scientific collaboration technologies can provide programmatic interfaces to enable new analysis tools to access data using standardized semantic descriptions, and thus greatly simplify the roll-out of new tools that can be made available to all members of the scientific community.

Provenance

We define provenance as information recording how the data were acquired and processed. Provenance is used to assure and document the quality of the data, perform forensics (e.g., to detect which parts of data were affected by the instrument, a faulty algorithm, or a bad processing node), and as a recipe how to reproduce intermediate or output data from input data [11, 14]. This includes the complete who, what, where, when, and why of the data capture and processing. Provenance for extreme-scale collaboration requires the following technologies:

Automating provenance capture. Provenance must be available not only to file-based data, but also to track updates, etc., to record-based data in databases. Also, provenance must handle extreme semantic heterogeneity and allow adding new categories of information easily, as the data are used, fused with other datasets, and examined in new ways. For extreme volumes of data, the organization of provenance information must be efficient and flexible. Provenance may not be captured per fine-grained data object; instead, objects with the same provenance will be grouped. These groupings may need to be modified as objects are processed and re-processed by independent codes.

Derived data regeneration. Simply put, can the data and provenance be combined to verify a published result? Technologies addressing this problem could extend existing work such as the Karma system [49] and the Open Provenance Model (OPM) standard [38]. A crucial challenge is the large number of dependencies (software, instruments, inputs) that lead to a given result. Approaches are also needed that scale to potentially massive amounts of provenance data. Techniques such as graph-based datastores could help to organize the complex datasets. The rigorous ability to regenerate derived data also enables a world of virtual data, where derived datasets can be instantiated or remain virtual to optimize the total cost effectiveness of storage and processing resources.

Curation

Another technology that is needed for scientific collaboration, though not entirely specific to it, is curation of data across changes of technology and representation. Processing and storage

technology changes to a new generation every three to five years, while the data may remain research-relevant over decades, thus requiring curation technologies that transparently and efficiently migrate data to cheap long-term storage (e.g., tape) and the associated processing software to migrate to new hardware architectures and build environments. As new data representations emerge, the tools and support available to access and use the data in their original representation become obsolete, necessitating a migration to a newer representation. To enable this, it is necessary to bundle high-quality metadata and documentation with the data.

Visualization

Visualization is a powerful tool for analysis and can turn large volumes of data into knowledge by identifying structure and patterns. In extreme-scale collaborative science, one challenge is to connect data with visualization tools for which it may not have been designed. Visualizations of the corresponding metadata may also help users to find and navigate through datasets. Extreme-scale datasets are inherently difficult to visualize, as the data are typically too large to portray meaningfully in lists and tables, and even plots and graphics require multi-level resolution to allow visualization at many different scales (e.g., Sky in Google Earth). Powerful visualization tools typically require the data to be represented and organized in very tool-specific ways, a barrier to future use when the representation or technology base changes. Thus, the visualizations (of science or metadata) must also be treated as derived data, which may be represented with self-describing formats, have provenance, be curated, etc. Sharable visualization environments empowering geographically distributed teams are vital tools for extreme-scale collaboration. Although less specific to collaboration, extreme-scale science also demands progress in user interfaces presenting multiple visualization technologies, the treatment of multi-dimensional data, and semi-automated segmentation and feature extraction.

Technologies for Processing

Collaborations face high barriers to entry when they attempt to use extreme-scale processing resources. These include issues related to location transparency of processing, resource sharing and co-scheduling, and monitoring and forensics.

Location Transparency of Processing

Location transparency is a requirement for both processing and data. Processing presents a particular challenge in that facilities are necessarily disparate in their properties, but also unnecessarily different in the face they present to scientists. Movement of processing and data is not simple or transparent, particularly when processing at one center is dependent on processing at another. X-Stack is leading efforts to improve portability of applications to make it easier for a single code to run on multiple platforms; but once built, these applications need to be able to run transparently across the platforms, including their configurations and surrounding environments/ecosystems. This requires a degree of encapsulation that goes beyond just executable portability. Collaborations have increased difficulties in this area because they need to compose multiple, diverse codes and use diverse resources. The key areas where extreme-scale science needs research and development include:

Hiding architectures. Today, the most successful computational scientists often code explicitly for the differing multi-core and GPU architectures. This represents an increasing intellectual effort that

could bring wide benefits to extreme-scale science if it were complemented by research into how the techniques could be abstracted and made more widely available.

Application packaging and formal description. Location transparency does not extend to running applications on facilities for which they are totally unsuited. Formally describing applications and making them available in standardized packaging, likely with their execution environments, will be vital. Techniques for formal description should be developed in a research program that is closely tied to the practicalities of delivering location transparency.

Simple compilation and deployment on distributed systems. An integral part of the research and development program aiming at location transparency should be the early delivery of a simple (from the scientist's perspective) compilation and execution environment for distributed systems. Necessarily, the more difficult use cases will be initially set aside.

Resource Sharing and Co-scheduling

Today, not only is location transparency a dream, it is even difficult to discover available and appropriate resources across the DOE complex (and other agencies). Even when a collaboration receives allocations of those resources, it is difficult to optimize their usage. Each processing center has its own idiosyncratic policies which may not be understandable to humans, let alone scheduling software. In the future, this will be exacerbated by the need to synchronize experimental data gathering or preparatory processing at smaller sites with the main processing at expensive exascale machines, even in real time.

As an example, an interaction file generated at NERSC might be combined with observational data from SLAC and then fed into a simulation at OLCF with post-processing leading to visualization for the PI at PNNL. Orchestration of these tasks requires scheduling of compute, storage, and network resources (and beam time) across multiple facilities and team members. Today orchestration is done manually and is error-prone, impeding the scientific end result.

Formalized resource and policy description. Closely paralleling the need for formal descriptions of applications is the need to apply similar rigor to computational resources and the policies that govern their use.

Co-scheduling algorithms. Once resources, policies, and application requirements are formally described, the next frontier will be to implement and iteratively improve co-scheduling algorithms. Decades of experience with the evolving needs for “simple” batch scheduling algorithms make it certain that this much more complex environment will present a substantial intellectual challenge.

Monitoring Debugging and Forensics

Observation of a computation is just like observation of a physics experiment; it requires the same level of tools for data reduction and analysis, just in a different domain. These tools must be collaborative and able to deal with both highly distributed computations and highly distributed collaboration members. Just as physics experiments are producing ever-larger volumes of often heterogeneous data, computations are producing ever-larger volumes of heterogeneous monitoring information. Failure interpretation and debugging is more complex when multiple collaborating platforms are involved.

Software systems of any level of complexity typically require more code dealing with monitoring, error handling, and debugging than is required to provide the core algorithmic function. In the distributed extreme-scale environment, this imbalance must grow, and the success of the scientific endeavor will depend crucially on research into the automation of monitoring and debugging.

Extreme-scale science involves collaborative exploitation of resources with extreme value. A simple assumption that they will not be misused is unacceptable. There is a vital need to be able to detect misuse and to demonstrate that this detection can be trusted both for current and past activities.

Synergistic Services

The computational tools and domain solutions for collaborative science depend crucially on a robust, sustained set of underlying synergistic services across the DOE computational facilities, local community clusters, and individual investigator laptops and mobile devices (Figure 7).

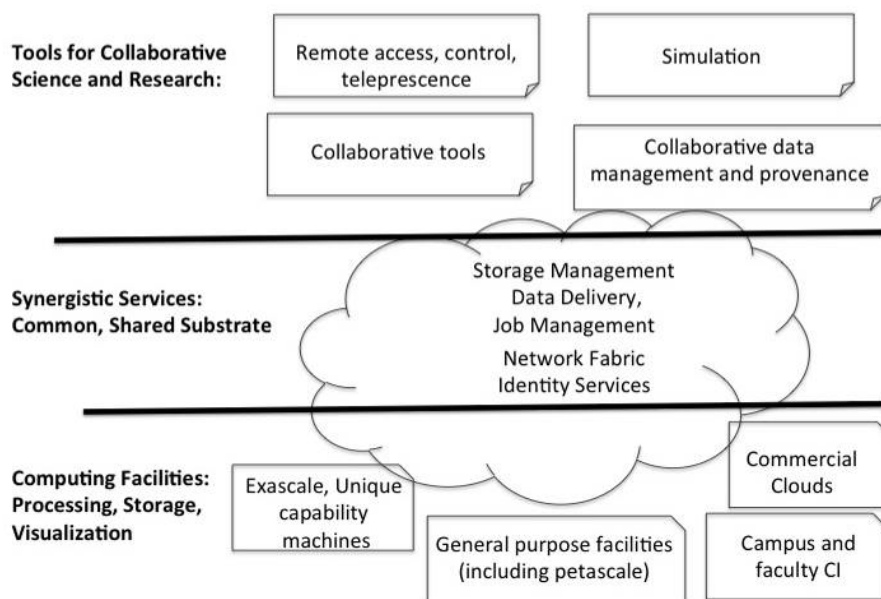


Figure 7: Layering of Services for Collaborative Science

It has been the goal of the SciDAC collaboratories program over the past decade to research, develop, acquire, and deploy the necessary services as a common underlying substrate across the DOE complex — communities and facilities — to ensure effective and reusable solutions for the existing DOE science communities and as well as other communities. The objectives have been to provide an integrative layer of shareable common synergistic services across the diverse and evolving facility capabilities and upon which all DOE collaborative science can rely.

As the scale and complexity of the science extends to the extreme over the next few years, it is essential not only that these services are sustained but that the providers are engaged in architectural and design decisions and also evolve their technologies to support the new capabilities needed to ensure the most (cost) effective and usable solutions possible.

The synergistic services provide an excellent layer to deliver unifying services that meet the abstractions identified in this report:

- to discover and use the diversity of DOE facility resources (e.g., information/data catalogs and publishing)
- to include central services used by all communities (e.g., federated identity management, data delivery)
- to ensure the connectivity needed between the science applications and the facilities (e.g., network management and monitoring)
- and to facilitate and test the portability of applications and data across the enterprise (e.g., provide common interfaces to diverse facilities and implementations).

Over the next decade, as the facilities continue their order-of-magnitude scaling in processor power (e.g., ubiquitous deployment of architectures such as exascale, MIC, GPU) and storage (e.g., exabyte data caches, zetabyte tape archives), well defined interfaces and capabilities must be provided for the synergistic services so that all communities can depend on them and integrate their use into their systems. The providers and supporters of these services must contribute to and abide by the common architectural agreements, principles, and boundaries. Gaps in the capabilities and robustness of today's implementations of synergistic services were identified in the workshop. These need to be addressed in concert with all the science communities and facility providers.

5. Characterizing Technical Capabilities

This section captures the perspective of the workshop participants on the issues that must be taken into account when defining a program of research and development aimed at scientific collaboration. Each of the needed technical capabilities was examined in terms of its:

1. **Benefit to science:** The estimated impact on scientific productivity.
2. **Specificity to collaboration:** Does this capability benefit collaborative science much more than science in general.
3. **Exclusivity to DOE science:** Does this capability benefit the science pursued by DOE much more than science and commerce in general.²
4. **Engineering versus computer science:** Is the provision of this capability principally an engineering task, or does it require advances in computer science.
5. **Cost:** The order of magnitude of the effort required.

The relationship of the technical capabilities that were examined is illustrated in Figure 6 on page 22.

During the report writing phase, workshop participants³ contributed their personal judgment on each of these issues. A summary appears in the table below. The numbers in the columns range from 1 to 5 as follows:

- **Benefit:** least to most benefit to science
- **Spec:** least to most specific to collaboration
- **Excl:** least to most exclusive to DOE
- **En/CS:** mainly engineering to mainly computer science
- **Cost:** lower than 1 FTE years to more than 30 FTE years

² Answers to this question were based only on the nature of DOE science in comparison with other science and technology. Issues such as “Is this already or is it likely to become an assigned DOE responsibility after discussions with other Federal agencies” were ignored.

³ All workshop participants, apart from DOE program managers, were polled. Responses were received from 26 participants, about half of the total.

**Table 1: Characterization of technical capabilities
based on the personal judgment of workshop participants**

Technical Capability	Benefit	Spec.	Excl.	En/CS	Cost
Management of Collaboration	3.2	4.4	2.6	1.9	2.4
Integrated Collaborative Human Communications Systems	3.0	4.2	2.2	2.6	2.7
Remote Instrumentation and Monitoring/Steering and/or Control	3.3	3.2	2.9	2.7	2.7
Identity Management	3.0	3.3	2.0	2.2	2.0
Authorization	3.0	2.8	2.3	2.3	2.1
Software as a Service and Virtual Private Laboratories	3.3	3.5	3.1	3.3	3.3
Software Development Environment	3.5	3.3	2.2	2.5	2.9
Workflow	3.8	2.8	2.1	2.7	3.3
Efficient Sparse Access to Exascale Data	4.0	2.8	3.3	3.5	3.5
Efficient Distributed/Parallel Access/Queries of Exascale Data	3.9	2.8	3.5	3.6	3.1
Location Transparency of Data	3.2	3.2	2.4	3.1	2.8
Discovery	4.1	3.8	2.7	3.5	3.0
Secure Federation	2.7	3.4	2.7	2.9	2.5
Self Describing Data Formats and Semantics	3.5	3.0	2.3	3.7	2.7
Abstraction and Automation Tools	2.8	2.6	2.1	3.7	2.5
Automated Provenance Capture	3.9	3.5	2.7	3.2	2.9
Derived Data Regeneration	3.4	2.7	2.3	2.6	2.8
Curation	3.5	2.8	2.2	2.6	2.8
Visualization	3.7	2.6	2.2	3.2	3.1
Hiding Architectures	2.8	2.3	2.4	3.2	3.3
Application Packaging and Formal Description	2.7	2.5	2.5	2.3	2.6
Simple Compilation/Deployment on Distributed Systems	2.8	2.7	2.2	2.1	2.2
Formalized Resource and Policy Description	2.9	2.8	2.5	2.7	2.1
Co-scheduling Algorithms	2.8	2.5	2.5	2.9	2.5
Monitoring, Debugging and Forensics	3.2	2.6	2.4	2.6	2.5

6. Non-Technical Issues

Sociological Barriers to Collaboration

One of the biggest impediments to collaborative research is effective communication. Every field has its own terminology, where often the same term means different things to researchers in different fields. Additionally, certain basic information is internalized by every group as known, and is therefore not articulated in discussions. An added impediment to collaborations arises from the differences in priorities among different groups. Finally, typical single-discipline research communities tend to be insular. People prefer talking to like-minded people and often view contrary ideas with suspicion. All of these factors hinder the act of collaboration when teams are colocated; but add separation by a distance that prevents seamless face-to-face meetings, and the problem is compounded.

Research in the sociology of science has determined that technologies such as email, instant messaging, or conferencing technology (both video and phone) did not provide sufficient advantages to distributed teams [19]. That research identified some requirements of technology that might help change this outcome. These requirements included: (1) tools to manage and track the trajectory of tasks over time, (2) tools to reduce information overload, (3) infrastructure to facilitate ongoing conversations, (4) infrastructure to encourage awareness with reasonable interruption for spontaneous talk, (5) tools that support simultaneous group decision making, and (6) tools and infrastructure capable of supporting presentations and meetings across a distance. Many of these are now available in a variety of forms but have failed to become easy to use and to integrate into typical scientific teams and workflows.

Finally, there are social barriers surrounding data and its part in collaboration. Of course, sharing data is part of the basic process of science; but the distinction here is that data at all stages of the processing pipeline are being shared, so that one person may use “raw” data that were never included in a peer-reviewed analysis. What is fundamental here is trust, both of the data source — not to have their hard-won data used without credit — and of the data recipient that the data are not junk. There is a strong social component to this trust, e.g., from personal contact or position within respected organizations.

The establishment of trust speeds up the science by focusing researchers on the data with the highest quality or relevance, in much the same way that a professor accelerates learning by guiding students. At the same time, humans are fallible, and this can result in trust being misplaced. One good example of this is the scandal at Duke around “personalized medicine” [7], which raised questions about the entire sub-field of research. This example occurred with peer-reviewed journal publications, which are presumably more controlled than sharing of raw datasets. How can this phenomenon be avoided in the context of collaborative science? Easily used tools for cross-validation of results across projects and groups could help alert scientists to discrepancies. The results of these cross-validations need to be made accessible along with the original data, so that external viewers can be made aware of criticisms and anomalies; one of the most worrying aspects of the study cited above was the ability, for a time, of a single institution (Duke) to bend the peer review process to squelch critical analyses from external sources.

At the end of the day, collaboration is a complex effort that is studied at many different levels; the establishment of trust and clear communication paths is key to a successful outcome. Any

environment or tools constructed need to promote and simplify the establishment of trust and communication.

Removing Barriers to Collaboration

One of the barriers to successful and beneficial collaborative research remains the ability of the groups to achieve the effectiveness of a common organization. Lessons learned by the Earth Systems Grid, high energy physics, astrophysics experiments and others, point to the need for attention to the community approach. This attention must be continued through the lifetime of the community and in the face of conflicting group and individual priorities, schedules, and goals. Challenges remain to ensuring that the diversity of the community is reflected by the diversity of the leadership and management structures — moving the culture from “they” to “we.” Community-wide ownership of the charter, objectives, and governance of the collaborative enterprise, with clear decision making responsibilities and decisions made by consensus, can contribute to successful outcomes. In the research environment the ability to be flexible, agile, and responsive to the changing conditions is a must. In a diverse, multi-cultural environment, the autonomy of the contributors and collaborators, together with mechanisms to avoid compromising the overall community goals, help ensure the appropriateness of the collaborative response.

Teams face the additional barrier of buy-in to and credit for the “less desirable work.” In computational projects this includes testing, documentation, and troubleshooting. Facing the challenge includes ensuring attention and credit is given to the needs of the individual, including allowing opportunities on the ground for the technical and professional growth of the individual.

Gaps in the Office of Science Strategy for the Support of Collaboration

Several challenges exist in the R&D into, and long-term support of, software and services that enable scientists to collaborate. Here we highlight three significant challenges of this type.

No one party, DOE/SC or otherwise, entirely owns this space. As such there is potential to assume responsibility lies with other parties, as when DOE says, “We don’t do x because NSF [or industry] does it.” But these problems are hard, and often need all available intellect applied to them. In addition, extreme-scale DOE science often has unique requirements. There is no substitute for a careful examination of DOE needs in relation to the wider scientific and commercial environment.

Transitions from research to development to implementation are notoriously hard. If facilities and science disciplines do not participate in R&D by at least communicating their needs and desired outcomes, there exists significant risk that useful R&D prototypes cannot make the transition to delivering positive impact in advancing science. To the extent that joint plans can be made early in the process which target tangible science outcomes well beyond the prototype, this gap can be narrowed. In addition, mechanisms are urgently needed that will permit DOE/SC science projects that use technologies to contribute to their development and support.

Computer science metrics for R&D success sometimes do not include the end-to-end impact on science of new software and methods. The value of discovering and delivering performance gains, novel features, and overall utility to scientists can be measured in differing ways. It is important not to conclude the research process too early, as a mere proof of principle will often miss many

interesting questions that arise in the end-to-end scientific process. DOE's history of mission science focuses on large long-term research campaigns well beyond what universities and other agencies are capable of sustaining. As such it requires focus on long-term support of the means through which collaboration can be sustained and the proper metrics are science impact rather than simple prototypes.

Focusing the SC Strategy for the Support of Collaboration

Translating the potential of scientific collaboration into capabilities that empower science at the extreme scale will require a coordinated and sustained research and development program. A broad community of scientists, community leaders, and software developers will have to pull together their intellectual and physical resources in a coordinated effort to address the challenges faced by ever larger and more diverse collaborations. The DOE Office of Science, which had the vision to launch a National Laboratories program in 2001 and is today responsible for the world's largest collection of scientific user facilities, is uniquely positioned to create and lead such a program. The lessons learned from building and operating the existing laboratories program (described in Appendix B) as well as the requirements and expectations identified by the broader DOE science community will provide the technical foundation and guidance to the researchers and developers who will be engaged in developing novel collaboration frameworks and in building powerful software tools. These tools will enable scientific collaborations that operate at the extreme scale to *discover*, *centralize*, and *port* resources across organizational and community boundaries.

Many aspects of this innovative program call for a community-based entity that will provide an intellectual and organizational anchor. This entity will enable the scientific collaboration community to sustain the needed communication and engagement to meet the vision, to remove current barriers to collaboration and ensure they remain lowered, and to provide a forum for developing and maintaining the unifying set of abstractions.

Such a forum (consortium, institute, foundation) for scientific collaboration in a computational environment will play a critical role in transitioning the field to a new level in addressing the current and future challenges of collaboration at the extreme scale. The forum needs to be inclusive across projects that focus on computer science research, those that develop and deploy tools into production, and the scientific collaborations themselves. The forum will bring together existing and new projects to ensure ongoing dissemination of knowledge, widespread adoption, and training in current best practices, methodologies, and techniques. It will serve as a vehicle to formulate needs, recognize capabilities, and identify commonalities across the DOE science community. It will also help in training and advising about the ins and outs of scientific collaboration.

It should be noted that this is an extension from and takes an expanded approach from the scope of the SciDAC Outreach Center [48]. This forum (consortium, institute, foundation) has the mission to provide a proactive, centrally organized core of effort to work on the technical and architectural frameworks, not just outreach, and to be an effort across the types of community (engineering, science, CS) involved. It is crucial that the engagement and intellect of the many are engaged in making the vision (charter, governance, strategies) and further in bringing the vision to fruition.

Appendix A: Science Narratives

The DOE Office of Advanced Scientific Computing Research (ASCR), in partnership with the other Office of Science programs, sponsored a Scientific Grand Challenges Workshop Series over the past few years. Each workshop focused on the grand challenges of a specific scientific domain and the role of scientific computing in addressing those challenges. The primary goal of these meetings was to engage the relevant scientific communities in a dialogue to identify the opportunities and challenges ahead and the role for scientific computing and multidisciplinary partnerships.

At each workshop, approximately 50 invited technical leaders in the field of extreme-scale science and computing discussed ways to overcome the most challenging technical issues and produced a report that includes a list of actionable recommendations. These reports are available to the public at <http://science.energy.gov/ascr/news-and-resources/workshops-and-conferences/grand-challenges/>.

Using those reports as a resource, this appendix highlights the extreme-scale collaboration needs of specific research communities.

Accelerator Science

Particle accelerators are among the most versatile and important of the DOE's extraordinary tools for extraordinary science. The nation's accelerators are responsible for a wealth of advances in materials science, chemistry, the biosciences, high-energy physics, and nuclear physics. They also have important applications to the environment, energy, and national security, such as studying bacteria for bioremediation, exploring materials for solar cells, developing accelerator-driven fission- and fusion-energy systems, and developing accelerator-based systems to inspect cargo for nuclear contraband. The technologies developed using particle accelerators also have a huge impact to the nation's economy by helping to maintain U.S. leadership in science and technology. Accelerators also have a direct impact on the quality of people's lives through applications in medicine, such as the production of medical radioisotopes, pharmaceutical drug design and discovery, and through the thousands of accelerator-based irradiation therapy procedures that occur daily at US. hospitals.

The next generation of accelerators presents major opportunities as well as major challenges. Designing future accelerators will involve the best minds nationwide and worldwide collaborating to develop designs that operate in new regimes of beam physics, meeting increasingly challenging design goals in terms of beam quality, beam intensity, beam energy, and other properties. Future accelerators will also involve optimizing three-dimensional, extraordinarily complex electromagnetic structures that are used to guide, manipulate, and control the beams. Collaborative design involving extreme amounts of data will require advanced, high-speed networks as well as advanced computational tools to explore and analyze the data.

At present the most challenging accelerator modeling problems are performed on more than 100,000 processors, and in the future that is expected to rise to more than 1 million processors. The simulations now use billions of simulation particles and billions of mesh points. At each step of a time-dependent simulation, the amount of data can easily exceed a terabyte, leading to full datasets in excess of a petabyte. It is not feasible to move entire datasets to all the desired remote locations for analysis. Instead, researchers will depend on collaborative, parallel, remote visualization and

data analysis tools that at present do not exist. These tools, and the infrastructure supporting them, will be essential for accelerator design optimization, risk reduction, and cost reduction.

Extreme-scale collaboration is also essential for exploring advanced accelerator concepts — involving extreme regimes of beams, plasmas, and radiation — where physical measurements and experiments may be too difficult or too expensive to perform. In such situations large-scale simulation, combined with collaborative data analysis, is a tool to explore physics that is otherwise inaccessible, and such simulations are a window to new discoveries in accelerator and beam physics.

In addition to the extreme-scale collaborative technologies for accelerator design and accelerator science discovery, such technologies are also needed for analysis of the data produced in experiments at accelerator facilities. These requirements fall under the heading of collaborative tools for large-scale experiments that are described elsewhere in this document.

In summary, DOE researchers are involved in the design and development of many new accelerators that will have a major impact to science, to people's lives, and to the nation. These facilities — including fourth-generation light sources, high-intensity proton drivers, rare isotope facilities, electron-ion colliders, electron/positron and muon colliders, facilities for high energy density physics research, and facilities for accelerator-based energy production, as well as advanced concepts including laser-plasma and beam-plasma accelerators — will all involve extreme-scale collaboration for design, optimization, exploration, and discovery.

Astronomy — LSST

The Large Synoptic Survey Telescope (LSST) [31] is a large-aperture, wide-field, ground-based telescope that will survey the visible sky every few nights in six photometric bands. As a public facility, the LSST will make available its images, alerts, and resulting catalogs to the United States and Chilean communities with no proprietary period. The planned 10-year survey will produce a database suitable for answering a wide range of pressing questions in astrophysics, cosmology, and fundamental physics. The same dataset can be used to characterize the properties of dark energy and dark matter; produce nearly instant alerts of detected optical transients like exploding stars in distant galaxies; discover and provide orbits for potentially hazardous near-Earth objects; and catalog billions of objects with astrometric precision and to unprecedented photometric depths. The LSST project has been endorsed by the 2010 decadal survey and several other national studies of priorities for astronomy and physics.

The massive data produced by the LSST must be managed efficiently and analyzed in real time, thereby providing an important testbed for new approaches to data management. The LSST data management system must employ innovative, large-scale database techniques including parallelization of queries, memory-based indices, and data partitioning and clustering. Also, LSST will use supercomputing technologies and create a general-purpose data and algorithm-parallel framework that will be available as open source software and reusable on any high-performance, parallel scientific application. As a result, the astronomy and physics communities will have an open source example to leverage for future projects.

Finally, based on the science experiments described in the LSST Science Book and other as-yet undefined science requirements, there are large computational requirements associated with analyzing the LSST data. It is impractical to move the large volume of data to every researcher, so

the data management system must provide supercomputing-scale data centers that host the data and allow for processing by externally developed analysis codes, often integrated with LSST open source codes. Even those centers will not be sufficient for the most demanding analyses, so software portability to grid and national supercomputing centers is a requirement. Collaboration infrastructure supporting these researchers, while mostly beyond the scope of the LSST project itself, will be critical for the successful exploitation of the data and software.

Astrophysics Simulations

Astrophysics simulations are typically multiphysics and multiscale. The deaths of stars in explosive events provide some of the most striking examples of these characteristics, involving simulations that evolve dozens of degrees of freedom representing a raft of physical processes at each spatial grid point. These simulations also typically require prodigious use of adaptive mesh refinement to attain the requisite spatial resolution given the vast range of scales involved in the problems. Particular examples include the study of the deflagration and detonation phases in the Gravitationally Confined Detonation (GCD) model of Type 1a (Snl) supernovae [28, 42] using FLASH [21, 25] (Figure 8), and core-collapse supernovae using Chimera [36, 37]. These simulations produce vast quantities of data, and the process of scientific discovery involves analyzing these data repeatedly and via a variety of techniques.




Figure 8. Snapshot of gravitationally confined detonation simulation of the Snl just before the detonation.

In the Snl studies, for example, at least three more stages of computing are necessary. Each stage uses a different code developed by individuals and/or groups with different expertise. The Lagrangian tracer particle time-histories produced by the FLASH simulations are used to drive

post-processed simulations of explosive nucleosynthesis in the event. The output of this analysis is mapped to a mesh suitable for, and understood by, the radiative transfer codes which compute the light curves and spectra that can be compared to the observed data. Core-collapse simulations are similar in their analysis workflows; e.g., tracer trajectories are routinely processed for nucleosynthetic products. In addition, gravitational wave signatures and neutrino signals in terrestrial detectors are calculated for these events using codebases produced by experimental groups that include estimates of detector response and signal-to-noise information.

The FLASH Center has achieved significant success in automating the first phase of simulating the GCD model of deflagration followed by denotation using Smaash, an in-house simulation management and analysis system [27]. Smaash seeks to solve the challenges of monitoring simulations, cataloging the metadata of the run, and capturing and cataloging the output analysis data by offering what approaches a single point of control and analysis, and a metadata-base. In a similar vein, the Oak Ridge group has developed a multi-faceted software support system for scientific simulation with CHIMERA called Bellerophon [30]. Bellerophon provides a “one-stop shop” for code development and scientific workflow management, fully automating code verification tasks by compiling and executing the latest revision of CHIMERA daily. The test results can then be analyzed directly via Bellerophon’s software tools. One of Bellerophon’s most powerful capabilities is the automated generation of concurrent visualization for ongoing CHIMERA simulations. Once a simulation is bound to the Bellerophon system, computational scientists can access and customize the resulting visualization, in near real time. Another software feature in Bellerophon delivers a plethora of dynamically generated code development statistics (e.g., code churn, LOC, commits).

However, much of the research pipeline continues to be operated manually. The challenges include: (1) the heterogeneity of the input/output formats of various codes; (2) involvement of geographically disparate groups with different subcultures and preferences; (3) currently available tools are not equal to the task; and (4) complexity of the task combined with lack of adequate resources makes it impossible to develop adequate tools in-house. One difficulty that has been brought to light by our experiences is that it requires a great deal of effort to maintain a logical storage and semantic association between the various datasets, analyses, and results that comprise different phases of a unitary scientific project. All too easily, various intermediate and final data products, analysis scripts and binaries, visualizations, and analysis products can wind up strewn across many storage systems, workstations, and even laptops. This fragmentation becomes more likely the larger and more articulated the collaboration. The result can be the permanent loss of valuable information, as decentralized knowledge of the places of residence and detailed nature of various analyses and data products is ephemeral. For example, it is not an infrequent occurrence to lose access to such scientific results “in the can” in consequence of a personnel change, simply because nobody remaining in the collaboration knows how to interpret the files embodying those results (if they can even be identified in the file noise of a large storage system).

A desideratum of such collaborations, then, is some kind of tool that acts as a project “lab notebook,” keeping together files that are part of a unitary scientific project, imposing some clarifying structure on such files by means of an emphasis on operator/scientist notes, file manifests, and other such documentation, keeping track of such disparate elements as primary simulation data, analysis pipeline scripts and codes, intermediate data, visualizations and plots, and final “press-ready” scientific output.

Biosciences

The last decade has seen a dramatic increase in the amount of data originating from ever more automated high-throughput “omics” approaches, resulting in a overwhelming number of genomes, metagenomes, transcriptomes, proteomes, metaboloms that are stored in ever-increasing databases. So while we have an exquisite inventory of genes, proteins, and metabolites, yet we are only beginning to understand how living systems work. What is missing for a complete systems biology understanding is how all these components are organized spatially and temporary in their specific cell and tissue/microbial community context. Deviations from the compositional and architectural organization are the base for disease development and/or breakdown of normal environmental function (e.g., plants and microbes) with devastating consequences. Such systems biology understanding is key to the success of global challenges such as bioremediation, the biological processes in the carbon cycle, as well as plant biomass-derived transportation biofuels.




Figure 9. New imaging tools are producing tens of terabytes of biological data.

While interpretation of imaging data is hardly a novel idea, and consciously or subconsciously is carried out on a regular basis by every scientist, it is the overwhelming large volume of datasets that can be generated which makes it a challenge. With the recent development of advanced electron microscopy imaging approaches, it is now possible to acquire 3D architectural datasets of cells and tissues/communities comprising 32,000 by 32,000 by 20,000 voxels, hence resulting in tens of terabytes of information, which can no longer be dealt with by human visual inspection alone and by manual image analysis and feature extraction (Figure 9). Furthermore, such datasets

are likely to be of such complexity that biologists of different disciplines will work in isolation or collaboratively on the different aspects of such datasets, ideally along with the imaging and image analysis scientists. If we do not want such datasets to simply vanish into the black hole of a data repository, we will need to develop tools for sharing such datasets between multiple sites around the world, and provide tools for data annotation, feature extraction, model building, animation, and quantitative analysis. It is conceivable that entire organisms such as nematodes, fruit fly larvae, and zebrafish larvae can be imaged at macromolecular resolution, and hence we can envision taking anatomical dissection to the level of the cells and macromolecular complexes.

In order to realize the potential of true collaborative analysis on multidimensional datasets, the 3D data have to be stored in a common repository and made available to the scientific community at large. Collaborative technologies need to be developed that enable the remote analysis and interactive visualization of data as well as the effective interaction between groups of scientists that may be geographically distant.

Experimental Magnetic Fusion Energy

The long-term goal of Fusion Energy Science (FES) research is the development of a reliable energy system that is environmentally and economically sustainable. Nuclear fusion, the power source of the stars, has been the subject of international research for over five decades. To achieve this goal, it has been necessary to develop the science of plasma physics, a field with close links to fluid mechanics, electromagnetism, and nonequilibrium statistical mechanics. For the experimental component of FES research, progress has been paced by facilities of increasing size, complexity, and cost, resulting in a decline in their numbers. This path has led to a concurrent growth in the importance of collaborations among large groups at the experimental sites and smaller groups located throughout the world. For U.S. scientists, the future focus of experimental research is shifting towards Europe and Asia. As a result of the highly collaborative nature of FES research, the community is facing new and unique challenges.

Experimental magnetic fusion research in the U.S. is centered at three large facilities (Alcator C-Mod, DIII-D, NSTX) with a present day replacement value of over \$1 billion (Figure 10). Teaming with this experimental community is a theoretical and simulation community that concentrates on the creation of realistic nonlinear 3D plasma models. Working together to advance scientific understanding and innovation, these two groups represent over one thousand scientists from over forty institutions. As the



Figure 10. A view of the inside of the DIII-D Tokamak.

capabilities of wide-area networks have increased over the past decade, more researchers have begun to collaborate with the experimental institutions from their home laboratory rather than travelling for experiments or permanently relocating. While the community has made significant

progress in accommodating this new pattern of use, much remains to be done if we wish to take full advantage of newly emerging technologies.

Magnetic fusion experiments operate in a pulsed mode where in any given day, 25–35 plasma pulses are taken with approximately 10 to 20 minutes in between each ~ 10 s pulse. Throughout the experimental session, hardware/software plasma control adjustments are made as required by the science. These adjustments are debated and discussed among the experimental team within the roughly 20-minute between-pulse interval. This mode of operation places a large premium on rapid data visualization and analysis that can be assimilated in near real time by a geographically dispersed research team.

The magnetic fusion community has considerable experience in placing remote collaboration tools into the hands of real users, and the community's desire for more efficient collaboration is well known, as was identified in a review by the National Research Council [39] in 2000. Efforts to improve collaboration within the community have included sharing of resources and co-development of tools, but these have been done in mostly an ad hoc manner. The ability to remotely view operations and to control selected instrumentation and analysis tasks was demonstrated as early as 1992 [22]. Full remote operation of an entire tokamak experiment was tested in 1996 [26, 34]. Today's experiments invariably involve a mix of local and remote researchers. Support for remote collaborations was facilitated by the adoption of a common data acquisition and management system, MDSplus [35]. Based on a client/server model, it is currently in use on 30 experiments worldwide, storing digitized, analyzed, and simulation code data, becoming a de facto standard. Most recently, the work of the National Fusion Collaboratory Project (2001–2006) has created FusionGrid for secure remote computations and has placed collaboration tools such as Access Grid, VRVS, shared display walls, and application sharing into operating control rooms [29, 47].

Looking toward the future, the U.S. DOE Facilities for the Future Report listed ITER as the highest priority for the Office of Science. ITER is a burning fusion plasma magnetic confinement experiment under construction in France and is the next major step in an international program aimed at proving the scientific viability of controlled fusion as an energy source. It will be run as an international collaboration, with researchers from China, Europe, India, Japan, Korea, Russia, and the U.S. sharing operational and scientific responsibilities. Although operation is approximately a decade away, the U.S. fusion program is increasingly oriented toward ITER; even now a significant portion of the national program is organized around coordinated efforts to develop promising operational scenarios. U.S. scientific involvement on testing these and other scenarios for ITER operation are occurring on the three large U.S. machines but also increasingly on the EAST tokamak in China and the KSTAR tokamak in Korea. These two machines have capabilities in certain areas that are beyond those of the U.S. machines (e.g., long-pulse), and U.S. participation during planning and operations continues to expand.

Effective remote participation on the scale envisioned requires the provision of a working environment in which personnel at different sites can engage in dialogues enriched by data flows and where the “friction” and barriers to communication are significantly reduced. Experimental programs will greatly benefit through the provision of a flexible, standards-based collaboration space, which includes advanced tools for ad hoc and structured communications, shared applications and displays, enhanced interactivity for remote data access applications, high performance computational services, and an improved security environment. Since the ability to share complex visualizations and applications among remote participants is a necessary adjunct for scientific dialogues, interpersonal communications media needs to be enriched via remote sharing

of displays and applications among researchers. Distributed shared displays will be an important element supporting remote control rooms and will need to operate in the high latency environment inherent to long distance communication. While FES has a significant track record for developing and exploiting remote collaborations, with such large investments at stake, there is a clear need to improve the integration and reach of available tools.

Environmental Molecular Sciences Laboratory: A Long-Term View of Data Sharing, Analysis and Collaboration

The Environmental Molecular Sciences Laboratory (EMSL) is funded by DOE's Office of Biological and Environmental Research, which supports world-class research in the biological, chemical, and environmental sciences to provide innovative solutions to the nation's environmental challenges as well as those related to energy production. EMSL's distinctive focus on integrating computational and experimental capabilities as well as collaborating among disciplines yields a strong, synergistic scientific environment. In fact, EMSL is unique as a DOE user facility in bringing computing as a capability on par with experiment. Bringing together experts and state-of-the-art instruments critical to their research under one roof, EMSL has helped thousands of researchers use a multidisciplinary, collaborative approach to solve some of the most important national challenges in energy, environmental sciences, and human health. These challenges cover a wide range of research, including synthesis, characterization, theory and modeling, dynamical properties, and environmental testing.

EMSL science is primarily based on three themes that help to define and direct the development of key capabilities and collections of user projects that will enhance scientific progress in the areas of environmental molecular science most critical to DOE and the nation. Although each science theme focuses on drivers important to that field of science, the themes share significant overlap and linked areas of common interest. Thus, the scope of a research project in EMSL may impact all three science themes. The three science themes are as follows:

1. **Biological interaction and dynamics**, whose goal is to develop a quantitative, systems-level understanding of the dynamic network of proteins and molecules that drive cell responses and how groups of different cells interact to give rise to functional cell communities.
2. **Geochemistry/biogeochemistry and subsurface science**, whose aim is to study reaction mechanisms at the mineral-water, microbe-mineral, and fluid-fluid interfaces at the molecular scale and understanding the effect of these mechanisms on the fate and transport of contaminants.
3. **Science of interfacial phenomena**, where the main purpose is to develop an understanding and gain control of atomic- and molecular-level structure-function relationships at interfaces that enable the optimization of interfacial properties, such as the control of catalytic activity and selectivity.

An ongoing effort at EMSL is the collecting and storage of all data generated at the facility (over 150 instruments) and providing the necessary tools, primarily open source, to allow researchers to interact with the data and potentially data from other data sources via collaborative tools. We recognize that over the next decade, the increasing volumes, rates, and heterogeneity of data produced (multiple instruments same sample to multiple types of data — proteomics, transcriptomics, etc.) will require a shift in our approach to data integration and analysis. Due to a combination of factors, including the sheer volume of data produced and the increasing desire for

real-time data analysis, organizational policies that limit data sharing, and/or data ownership issues will require a reevaluation of the way analysis will be performed in a potentially highly distributed environment. To enable researchers to gain access and to collaborate with other researchers regardless of location in an effective way will require with data the development of a number of innovative methods for knowledge creation and discovery through distributed adaptive data analysis by developing computational environments that integrate and analyze data where it resides (rather than requiring its transfer to collection in a single location for analysis) and that provide collaborative tools for real-time interaction.

Lattice Field Theory for High Energy and Nuclear Physics

Quantum chromodynamics (QCD), the fundamental theory for nuclear interactions, is the strong nonlinear quantum field theory sector of the Standard Model of elementary particle physics, which describes how quarks and gluons are bound together to form the neutron and proton that build and interact to form atomic nuclei. The Standard Model has been enormously successful; however, our knowledge of it is incomplete because it has proven extremely difficult to extract many of the most important predictions of QCD, those that depend on the strong coupling regime. To do so from first principles and with controlled systematic errors requires large-scale numerical simulations within the framework of lattice gauge theory. These simulations provide the underlying strongly interacting “gluon plasma” that binds quarks into observable particles (Figure 11)

In addition, there is a new strongly interacting quantum field theory that is a serious contender to explain the new physics beyond the standard model (BSM) under investigation at the Large Hadron Collider (LHC) experiment at the forefront of the HEP energy experimental frontier. Over the last decade through the DOE SciDAC program, a unified software stack (Figure 12) has been constructed that serves all lattice field theorists under the direction of the USQCD Executive Committee (see usqcd.org). The results have been to accelerate the integration of new methods into application codes and the rapid adoption of each generation of high performance computer architecture.

However, lattice field theory is now undergoing a new period of rapid expansion into new physics domains just as computer architectures are becoming increasingly heterogeneous and hierarchical in order to reach to the exascale. Both trends are very promising but also disruptive, and they imply an increasingly complex algorithmic and software environment needed for lattice field theory. Keeping pace with these development calls for increased resources and interdisciplinary collaborations between physicists and applied mathematicians. For example, as the capabilities of




Figure 11. Snapshot of the “gluon plasma” that bind quarks in nuclear states. Large simulations using symplectic integrators generate these quantum space-time distributions on 4D hyper-cubic lattices. A linear algebra elliptic solver for the Dirac equation is a critical kernel that tracks the motion of quarks that bind into the observed protons and neutrons.

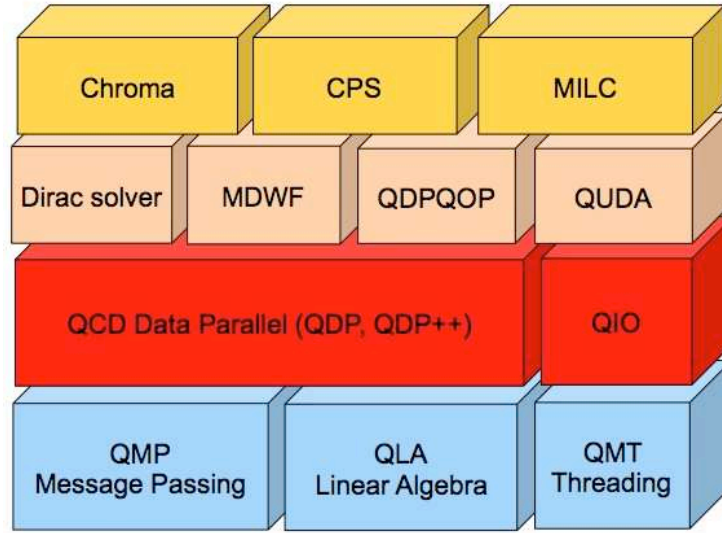


Figure 12. Software stack for USQCD develop by the SciDAC project on Software Infrastructure for Lattice Field Theory. Specialized architecture-aware libraries for communication (QMP/QMT) and arithmetic kernels (QLA) optimized support for a data parallel DSL for highly optimized linear algebra routines and the full array of application codes for nuclear and high energy physics (Chroma/CPS/MILC).

computers are enabling us to reach the physical pion mass, the state-of-the-art Krylov single scale conjugate gradient inverters are slowing down due to the increasingly singular nature of the Dirac operator. In response, new multi-scale algorithms are just beginning to be designed. This project goes back 20 years [13], but only in the last four years has a collaboration between applied mathematicians (TOPS) and lattice field theories (USQCD) invented an appropriate adaptive multigrid algorithm for the Wilson-clover discretization of Dirac operators [5]. Already the impact on light quark lattice demonstrates a 10–20 fold speedup on the BlueGene/P.

Lattice field theorists have a long history of interaction with hardware vendors, for example, developing the early QCDOC architecture leading to IBM's BlueGene series, and the recent collaboration with NVIDIA to optimize the high performance GPU library for QCD [9, 18, 6, 44]. There is clearly a need to establish a more permanent and continuing framework to support for this interdisciplinary work. Past experience demonstrates that success depends on a sustained interaction between application physicists, applied mathematicians, and computer hardware and software architects. New and optimal methods require that at every stage of development a continuous design-testing-implementation cycle be guided by the application itself with testing in actual production to prove its efficacy. Critical to this are unique individuals who are trained, gain experience, and find employment that rewards expertise in dual disciplines.

Appendix B: Previous ASCR Collaboratories Projects

Open Science Grid

As one of the existing collaboratories funded by the ASCR SciDAC program as well as the NSF, the Open Science Grid (OSG) [43] has evolved into an internationally recognized key national element enabling scientific discovery depending on distributed high throughput computing across a broad range of disciplines. This has been accomplished by a committed partnership that cuts across science disciplines, areas of technical expertise, and U.S. research institutions. OSG broadly promotes the adoption and advances the state of the art of distributed high-throughput computing (DHTC) — shared utilization of autonomous resources where all the elements are optimized for maximizing computational throughput. The U.S. LHC scientific program, the Fermilab Tevatron and High Intensity Frontier programs are committed to contributions to and use of the OSG fabric of services. The services span three groups: software services (the Virtual Data Toolkit); support services like education, training, consulting and; and an infrastructure of DHTC services (referred to as production services) for those who would like to join the OSG DHTC environment (Figure 13). These services support the broader community that builds and operates their own DHTC environments (e.g., LIGO, nuclear physics experiments, campuses) as well as supporting the DHTC environment of OSG.

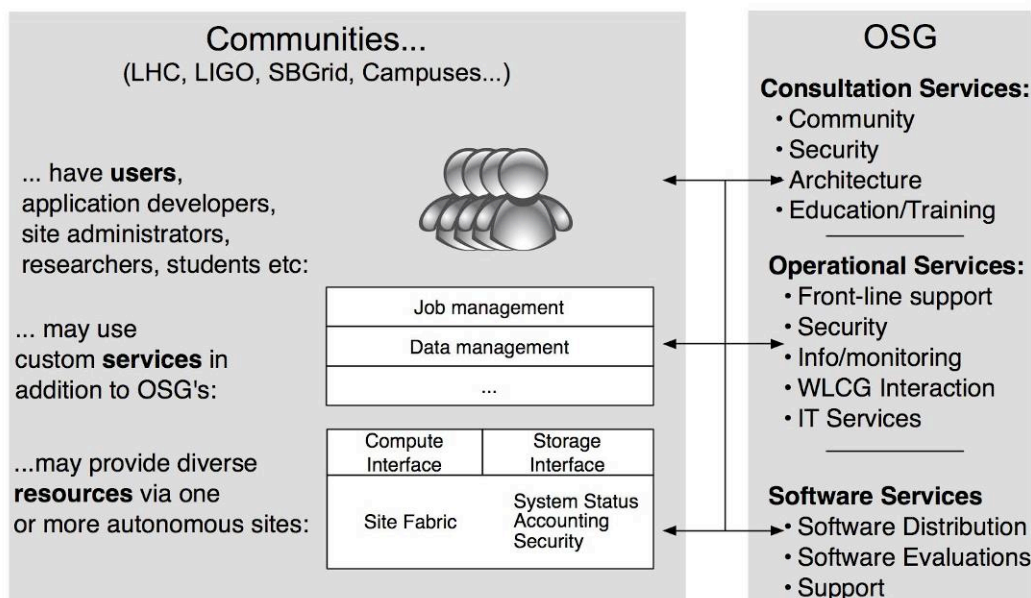


Figure 13. OSG's community-focused architecture. There is a sharing of software, operational services, and knowledge between the communities and OSG in each of these areas.

The National Fusion Collaboratory Project

Developing a reliable energy system that is economically and environmentally sustainable is the long-term goal of Fusion Energy Science (FES) research. As magnetic fusion experiments have increased in size and complexity, there has been a concurrent growth in the number and

importance of collaborations among large groups at the experimental sites and smaller groups located nationwide. Looking toward the future, the large-scale experiments needed for FES research are staffed by correspondingly large, globally dispersed teams. As a result of the highly collaborative nature of FES research, the community is facing new and unique challenges. While FES has a significant track record for developing and exploiting remote collaborations, with such large investments at stake, there is a clear need to improve the integration and reach of available tools.

The National Fusion Collaboratory (NFC) Project [29, 47] was funded by DOE under the SciDAC program to address these challenges by creating and deploying collaborative software tools. A five-year project that was initiated in 2001, it built on the past collaborative work performed within the U.S. fusion community and added the component of computer science research done within ASCR. The project was a collaboration itself, uniting fusion scientists from General Atomics, MIT, and Princeton Plasma Physics Laboratory and computer scientists from Argonne and Lawrence Berkeley national laboratories, Princeton University, and the University of Utah to form a coordinated team. The group leveraged existing computer science technology where possible and extended or created new capabilities where required.

The main objective of the NFC project was the development and deployment of a production national FES grid (FusionGrid) that would be a system for secure sharing of computation, visualization, and data resources over the Internet. The goal of FusionGrid was to allow scientists at remote sites to participate as fully in experiments and computational activities as if they were working on site, thereby creating a unified virtual organization of the geographically dispersed U.S. fusion community. The vision for FusionGrid was that experimental and simulation data, computer codes, analysis routines, visualization tools, and remote collaboration tools are to be thought of as network services. In this model, an application service provider (ASP) provides and maintains software resources as well as the necessary hardware resources. This grid's resources would be protected by a shared security infrastructure, including strong authentication to identify users and authorization to allow stakeholders to control their own resources. In this environment, access to services is stressed rather than data or software portability.

Deployment of the production FusionGrid infrastructure was made possible through work supported by the NFC project, by leveraging other SciDAC programs, and with the base funding for FES research. Access to fusion data was made secure via the integration of MDSplus and the Globus toolkit. MDSplus is a data acquisition/management system developed within the fusion community and in use at the three large U.S. experiments as well as other sites worldwide. The TRANSP code was released as a FusionGrid computational service along with supporting infrastructure development (data storage, monitoring, user GUI). This FusionGrid service was so successful that it became the production system for TRANSP usage in the United States and internationally. The development of the desktop-based Access Grid (AG) node combined with the concept of application sharing in the AG venue allowed usage of this capability in the tokamak control room. Shared display walls combined with new application sharing software were demonstrated in tokamak control rooms, resulting in fusion programs permanently installing new hardware to fully deploy this new functionality (Figure 14). The SCIRun visualization software was integrated with MDSplus storage of fusion simulation datasets to facilitate experiment/simulation comparison. Finally, interacting with FusionGrid's security system was simplified by the deployment of MyProxy [40] to manage scientists' credentials as well as through a new web-based centralized authorization system.



Figure 14. The control rooms of NSTX (a), DIII-D (b), and C-Mod (c) with shared display walls being used to enhance collaboration. On the DIII-D display is also video from remote collaborators in Europe who were participating in that day's experiments.

In the intervening years since the conclusion of the NFC project, the demand for better remote collaboration capability has only intensified. For example, the full operation of EAST in China and KSTAR in Korea has resulted in increased collaboration with these partners. These remote activities not only include experimental planning but also support for experimental operations that includes issues associated with plasma control and pseudo-real-time data analysis. With construction of ITER well under way, prototyping has begun for software solutions in a variety of areas, including secure remote control, networks, and data archiving. With such large investments at stake, there continues to be a clear need to improve the capability, integration, and reach of supporting tools.

The Earth System Grid

Earth System Grid (ESG), which evolved into the international Earth System Grid Federation (ESGF), was established to provide the worldwide climate research community with access to the data, information, model codes, analysis tools, and intercomparison capabilities required to make sense of enormous climate datasets. Its specific goals were to (1) provide an easy-to-use and secure web-based data access environment for datasets; (2) add value to individual datasets by presenting them in the context of other datasets and tools for comparative analysis; (3) address the specific requirements of participating organizations with respect to bandwidth, access restrictions, and replication; (4) ensure that the data are readily accessible through the analysis and visualization

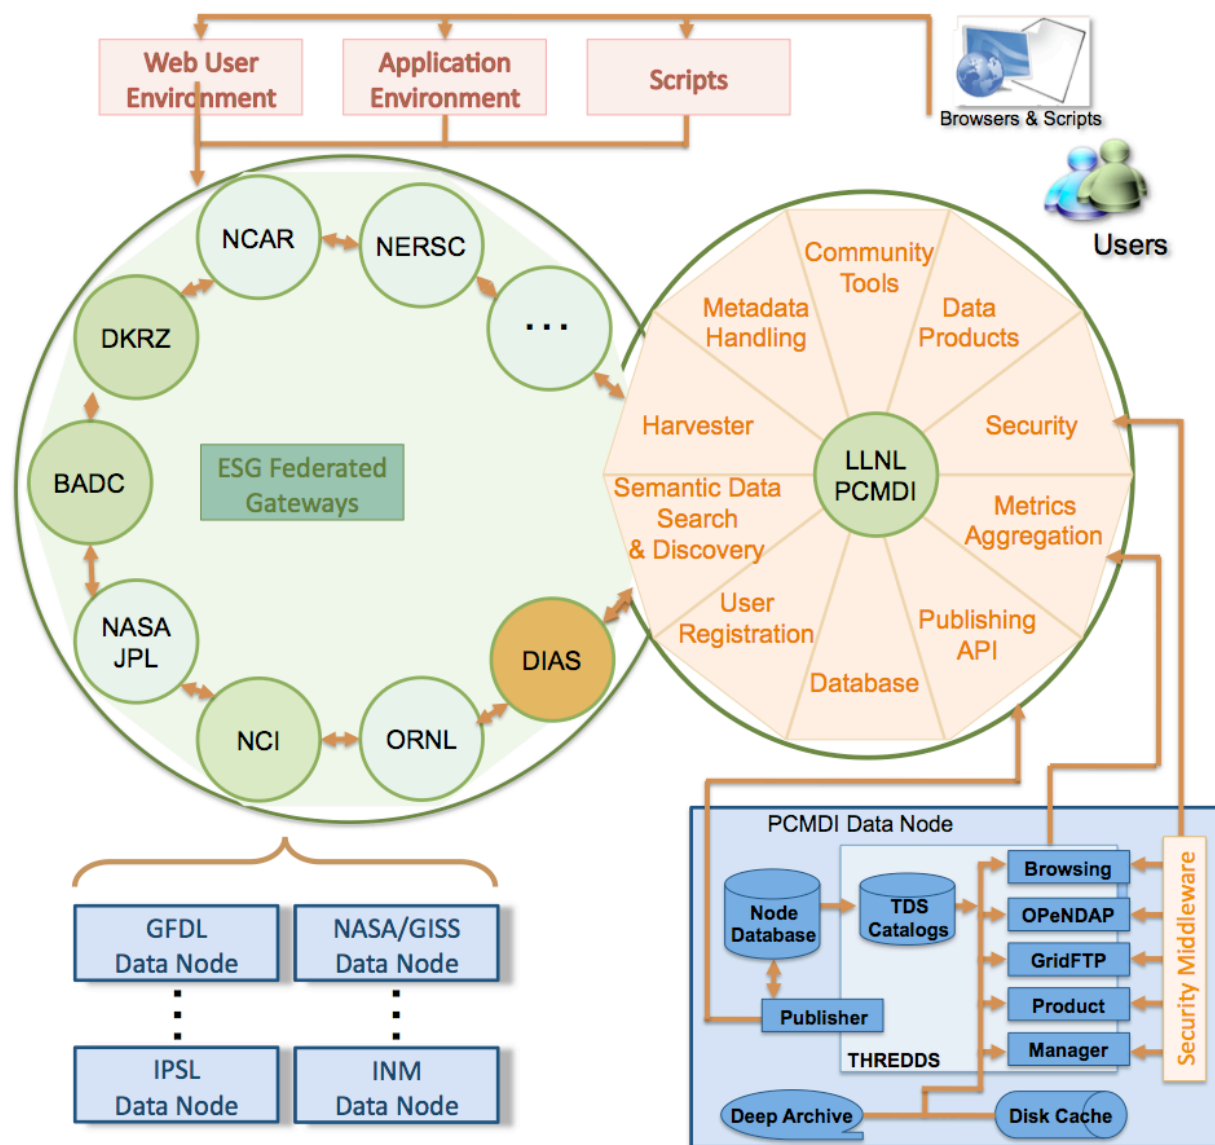


Figure 15. Earth System Grid Federation production environment.

tools used by the climate research community; and (5) transfer infrastructure advances to other domain areas.

ESG R&D has been supported for more than a decade by ASCR and BES, via the DOE Collaboratories and SciDAC programs. This support has allowed the ESG team to lead international development and delivery of a production ESGF environment for managing and accessing ultra-scale climate data (Figure 15). This production environment includes multiple national and international climate projects (such as the Community Earth System Model and the Coupled Model Intercomparison Project), ocean model data (such as the Parallel Ocean Program), observation data (Atmospheric Radiation Measurement Best Estimate, Carbon Dioxide Information and Analysis Center, Atmospheric Infrared Sounder, etc.), and analysis and visualization tools, all serving a diverse user community. These data holdings and services are distributed across multiple ESG-CET sites (such as ANL, LANL, LBNL/NERSC, LLNL/PCMDI, NCAR, and ORNL) and at unfunded partner sites, such

as the Australian National University National Computational Infrastructure, the British Atmospheric Data Centre, the National Oceanic and Atmospheric Administration Geophysical Fluid Dynamics Laboratory, the Max Planck Institute for Meteorology, the German Climate Computing Centre, the National Aeronautics and Space Administration Jet Propulsion Laboratory, and the National Oceanic and Atmospheric Administration.

ESGF software is distinguished from other collaborative knowledge systems in the climate community by its widespread adoption, federation capabilities, and broad developer base. It is the leading source for present climate data holdings, including the most important and largest datasets in the global climate community. Recently, ESG extended its services beyond data file access and delivery to include more detailed information products (scientific graphics, animations, etc.), secure binary data-access services (based upon the OPeNDAP protocol), and server-side analysis. The latter capabilities allow users to request data subsets transformed through commonly used analysis and intercomparison procedures.

ESG has been a tremendous success for the ASCR Collaboratories program. It is also illustrative of a broader challenge for DOE collaboratories. Because ESGF is providing production services, ASCR has ceased funding for the activity. However, production operation in fact involves not only substantial operations costs but also new, yet more demanding R&D. In particular, the ESG team faces substantial technical challenges due to the rapidly increasing scale of climate simulation and observational data, which will grow, for example, from less than 50 terabytes for the last Intergovernmental Panel on Climate Change (IPCC) assessment to multiple petabytes for the next IPCC assessment. The convenient access provided by ESGF also leads to increased demand, as many more users seek to understand, process, extract value from, visualize, and/or communicate climate data to others.

GridFTP and Globus Online

Modern science and engineering must increasingly deal with enormous amounts of data: terabytes (a thousand gigabytes or a million megabytes) are commonplace; petabytes (a million gigabytes) are on the horizon; and exabytes (a billion gigabytes) are projected for five to ten years from now. As data are produced, shared, and analyzed, they must frequently be moved from place to place. Thus, the ability to move data rapidly across networks has become fundamental to many collaborative projects in both research and industry. Unfortunately, despite rapid growth in network performance, the transfer speeds actually achieved by applications were for a long time embarrassingly low, often 1% or less of available network bandwidth. The effects on users are devastating. Moving one terabyte over a 1 gigabit/s (10^9 bits per second) network could take 10 days rather than a few hours, essentially preventing the use of remote systems and distributed collaboration.

The rapid, reliable, and secure movement of data from one file system to another is a challenging end-to-end problem. A single transfer may involve many disks, switches, firewalls, routers, and networks, as well as file systems, network protocols, and authentication protocols. High performance may often require the exploitation of parallelism at the storage system and network level. Application-level software may want to move a single large file, stream data, or transfer many small files. Individual hardware and software components may vary greatly in their performance and reliability. Failures may occur at the hardware and software levels.

Thus, we face three related challenges. First, we must develop methods that allow data transfer services to orchestrate rapid, reliable, and secure data movement in such diverse, heterogeneous environments with minimal user intervention. Second, we must define a protocol (a set of conventions for the messages exchanged by services) that is simple enough to be implemented efficiently but also rich enough to provide services and the information required to optimize data transfers. Third, we must persuade people to adopt the approach, so that we can be confident that when we want to move data, both sender and receiver are able to participate in a high-performance transfer.

DOE computer scientists have overcome these challenges via a combination of research, software development, and standards development. Working within the Open Grid Forum, an international standards organization, they led an effort to define extensions to the popular file transfer protocol (FTP) that enable secure, reliable, and high-performance transfers. (The resulting GridFTP extensions [2] range from messages for negotiating striping parameters to restart markers that permit an interrupted transfer to be resumed.) They also developed high-performance implementations of this protocol, within the Globus [1] and dCache projects, that use specialized methods to achieve extremely high performance in a variety of settings. For example, striping across multiple servers permits a large file (or many small files) stored in a parallel file system to be moved across a multi-gigabit/sec network (at up to 27 gigabits/sec to date), while pipelining [12] allows many small files to be moved efficiently. Finally, they succeeded in generating worldwide acceptance of the standard and software, such that experimental and simulation scientists worldwide use GridFTP for data movement.

Optional usage reporting in the Globus GridFTP implementation shows that the number of transfers reported (certainly a subset of the total performed) frequently exceed 12 million per day (Figure 16) and move more than half a petabyte. Experiments at the Large Hadron Collider make extensive use of GridFTP, transferring tens of terabytes daily. DOE supercomputer centers use GridFTP to move data in and out of their mass storage systems, achieving wide area performance 20 or more times what was achieved previously. The Advanced Photon Source uses GridFTP to allow U.S. and international users to access experimental data, again at vastly greater speeds than before. In addition to raw performance, GridFTP users like its built-in security (integrating with site security systems) and interfaces to different storage systems.




Figure 16. Monthly totals of GridFTP transfers, for those servers reporting

In more recent work, the Globus team has addressed usability challenges associated with GridFTP by developing Globus Online [3, 23], a software-as-a-service (SaaS) system for managing GridFTP transfers. Globus Online allows users to hand off simple and complex transfers to a service that handles performance optimization, error detection and recovery, and other otherwise time-consuming tasks, while requiring no local software installation (Figure 17).

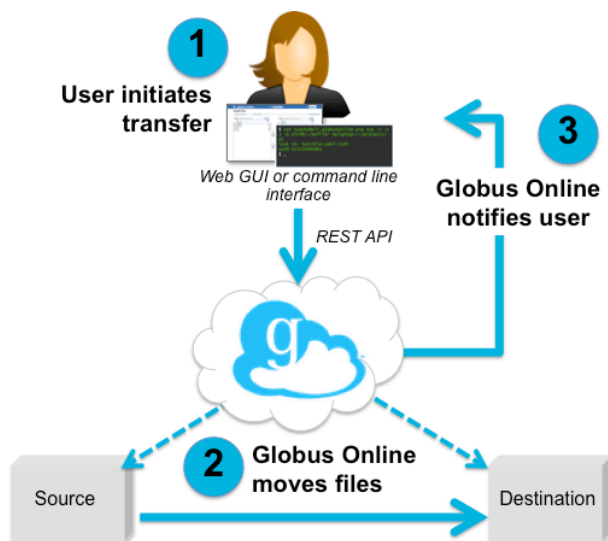


Figure 17: Globus Online data transfer service.

GridFTP and Globus Online development has been supported by ASCR for several years, most recently by the Center for Enabling Distributed Petascale Science.

Grids and Virtual Organizations

High-speed wide-area networks enable new modes of working based on on-demand access to remote resources (scientific instruments, data, computers), innovative applications that use resources at multiple locations, and new forms of interpersonal collaboration that were not previously possible. However, new technologies are needed for these resource federation and collaboration methods to function securely, reliably, and conveniently.

DOE computer scientists have been leaders in the development of the concepts and technologies that enable users to define, configure, and operate within distributed virtual organizations. These so-called grid technologies allow for the discovery of available resources, secure authentication of users and resources, and definition and enforcement of policies that govern access to resources. They have motivated the development of a wide variety of innovative applications and tools worldwide.

Within the DOE context, grid technologies underpin two major infrastructures that have had a significant impact on science: the Open Science Grid [43], which provides much of the U.S. data processing capabilities for major experiments at FermiLab and the Large Hadron Collider in Geneva, Switzerland; and the Earth System Grid [4, 10], the vehicle by which the climate simulation data used by the Intergovernmental Panel on Climate Change was distributed to scientists worldwide.

In the rest of this section, we detail three specific DOE contributions to collaborative science: Grid Security Infrastructure, the DOE public key infrastructure, and Access Grid.

Grid Security Infrastructure: Single Sign-on for Collaborative Science

The need for users to authenticate themselves when accessing remote resources is arguably the single biggest obstacle to collaborative science. Users used to have to acquire and use a separate credential for every site they wished to access — a process that was often insecure, error prone, and impractically complex when using multiple resources at once. The ASCR-sponsored research that defined and implemented the protocols and software called the Grid Security Infrastructure (GSI) [24, 50] overcame this problem. Thanks to GSI, scientists and engineers can now authenticate once and then access computers, storage systems, and other resources worldwide, all using a single credential. They can also delegate authority to programs to perform operations on their behalf, and resource owners can express and enforce policies concerning which credentials are accepted and which sets of users are allowed access.

Security is hard for many reasons. First, there are the technical challenges: developing methods and systems that permit rapid and convenient single sign on, and delegation of authority (essential in distributed systems) without comprising security. Then, there are the sociological and organizational challenges: getting site administrators, security experts, application developers, and users to agree on an approach that addresses the technical challenges while also meeting diverse requirements for performance, security, ease of use, audibility, and so forth.

DOE computer scientists have addressed these challenges via a mix of research, standards development in the Internet Engineering Task Force and Open Grid Forum, and high quality software development. Software that implements GSI standards is widely distributed and used via packages such as OpenSSL and Globus. GSI security methods are widely accepted, making GSI one of a handful of authentication schemes that have gained widespread support. (SSH is another.)

GSI is at the core of scores of collaborative science projects in the U.S. and abroad. It is GSI that has allowed projects such as LHC, OSG, ESG, TeraGrid/XSEDE [15], LIGO [16], caBIG [46], CVRG [51], and many others to integrate computers, storage systems, and other resources in a way that permits shared use. Collaborative science as it exists today would not exist without GSI.

GSI has received ASCR support via several projects over the years.

DOE Public Key Infrastructure

There is no aspect of modern collaborative science that does not involve trust among distributed collaborators. The mechanisms for establishing trust vary among science communities, as do the requirements for proving that a collaborator who wants to gain access to the collaboration resources is, in fact, part of the trusted community. A fundamental aspect of proving membership in a community of mutually trusted entities is that of establishing identity.

One common way of establishing identity is to use the approach of Public Key Infrastructure. PKI provides a mechanism to both identify an individual and to verify that the individual is who he, she, or it claims to be. (Systems and services, as well as people, have PKI managed identities.)

PKI as implemented by the ESnet DOEGrids Certificate Service (www.doeagrids.org), and during 2012 to be transitioned to the Open Science Grid (OSG) Certificate services, is tailored to accommodate the variety of science communities and approaches to trust found in DOE science collaborations. In particular, this service provides for different, and potentially incompatible, trust models. Each science community's approach to trust is rooted in the culture of that community, and different communities have different cultures, which leads to different trust models. In PKI parlance, the trust model is represented by a Certificate Policy. Different trust models require different Certificate Policies and therefore different classes of PKI certificates.

Since the PKI formalism does not allow for a single Certification Authority (CA) to issue certificates with different trust policies, ESnet set up a hierarchy of CAs with the function of the top-level ("root") CA limited to signing the certificates for a set of lower-level CAs that issue identity certificates to users in different science communities. Each lower-level CA can have a different trust policy and issue certificates to the science practitioners in particular science communities.

A second aspect of science communities is that identity and trust are based on group or virtual organization level relationships. Therefore, the people who can certify the identity of an individual and that individual's right to participate in the science community are people elected by the particular science communities. For each different community/group/virtual organization, a Registration Agent is designated to authorize issuing a certificate in the trust domain of the community — that is, causing the community CA to issue an individual a certificate. This approach makes it possible for ESnet to operate a collection of reliable and highly secure CAs for multiple communities that encompass large populations of users.

This process has resulted in two major ESnet accomplishments that are embodied in the DOEGrids Certificate Service: (1) a set of formal agreements on common, PKI identity based, trust agreements, and (2) establishment of a hierarchy of PKI user identity certificate-issuing CAs that use a scalable methodology for authorizing certificate issuance.

The largest community with a common trust model is the High Energy Physics (HEP) community and their use of the Globus software, which in turn uses PKI certificates to authenticate both people and services/systems. The DOEGrids PKI service trust model for this community is the result of several years of negotiation with the European parts of the HEP community to come up with a common certificate policy that is accepted by the U.S. and European institutions that are home to the HEP community of scientists.

The DOEGRIDS Certificate Service currently provides and manages certificates of some 2,500 humans and 6,500 service instances (computing systems and servers) in DOE science collaboration communities.⁴

Access Grid

In our increasingly connected world, teams in both science and industry are increasingly often distributed across multiple institutions. In this context, tools to enhance communication and

⁴ Additional comment in 2012:

In an example of how services are transitioned from the research environment to the sustained cyberinfrastructure, the DOE CA capabilities will be transitioned to be supported by OSG Certificate Services during 2012. OSG is exploring whether a commercial CA provider can meet the special needs of the DOE Science Community. At the moment this seems possible, but the development and agreements are still in progress.

collaboration can have a tremendous impact on productivity. To meet this need, DOE computer scientists developed Access Grid [17]: a collaborative environment in which advanced video, networking, and collaboration technologies are combined to enable many simultaneous participants to interact in a natural manner via the use of a mix of collaboration tools.

Access Grid innovations include scalable bandwidth management and the use of multicast protocols and multicast bridges to permit Access Grid collaborations to scale to many sites; a tool integration approach that permits the concurrent use of many collaboration tools, including video, audio, chat, and presentations; integrating security permitting authentication of users and sites; and scalability of platform from phone to laptop and fully configured collaboration room. In addition, Access Grid's packaging approach emphasizes the use of commodity hardware, software, and networking, facilitating adoption and widespread deployment.

The Access Grid is the only collaboration system built on open, standards-compliant technologies. Competing technologies are closed source and rely on proprietary protocols and data formats, limiting their applicability to only the function set that their developers decide to enable. In contrast, the Access Grid relies on unencumbered video and audio encodings and on well-established Internet technologies. The Access Grid is collaboration software, but it is also a platform for developing collaborative applications and performing further collaboration research. The Access Grid project has met the difficult tasks of building production-ready collaboration software, exposing key components as a foundational infrastructure, exploring collaborative applications thereon, and establishing and maintaining a worldwide community of users and researchers.

Access Grid users can be found in all areas of industry, academia, and research institutions worldwide. The number of downloads exceeds 50,000. Applications have included such diverse areas as computer science, surgical instruction, biology, climatology, and game development.

References

1. Allcock, B., Bresnahan, J., Kettimuthu, R., Link, M., Dumitrescu, C., Raicu, I. and Foster, I., The Globus Striped GridFTP Framework and Server. SC'2005, 2005.
2. Allcock, W. GridFTP: Protocol Extensions to FTP for the Grid. GFD-R-P.020, Global Grid Forum, 2003.
3. Allen, B., Bresnahan, J., Childers, L., Foster, I., Kandaswamy, G., Kettimuthu, R., Kordas, J., Link, M., Martin, S., Pickett, K. and Tuecke, S. Software as a Service for Data Scientists. *Communications of the ACM*, 55(2):81-88, 2012.
4. Ananthakrishnan, R., Bernholdt, D.E., Bharathi, S., Brown, D., Chen, M., Chervenak, A.L., Cinquini, L., Drach, R., Foster, I.T., Fox, P., Fraser, D., Halliday, K., Hankin, S., Jones, P., Kesselman, C., Middleton, D.E., Schwidder, J., Schweitzer, R., Schuler, R., Shoshani, A., Siebenlist, F., Sim, A., Strand, W.G., Wilhelmi, N., Su, M. and Williams, D.N. Building a global federation system for climate change research: the earth system grid center for enabling technologies (ESG-CET). *Journal of Physics: Conference Series*, 78, 2007.
5. Babich, R., Brannick, J., Brower, R. C., Clark, M. A., Manteuffel, T. A., McCormick, S. F., Osborn, J. C., Rebbi, C., Adaptive multigrid algorithm for the lattice Wilson-Dirac operator, *Phys. Rev. Lett.* **105**, 201602, 2010. arXiv:1005.3043 [hep-lat].
6. Babich, R., Clark, M. A., Joo, B., Shi, G., Brower, R. C., Gottlieb, S., Scaling Lattice QCD beyond 100 GPUs, arXiv:1109.2935 [hep-lat].
7. Baggerly, K.A. and Coombes, K.R. Deriving Chemosensitivity from Cell Lines: Forensic bioinformatics and Reproducible Research in High-Throughput Biology. *The Annals of Applied Statistic*, 3(4):1309-1334, 2009.
8. Baraniuk, R.G. More Is Less: Signal Processing and the Data Deluge. *Science*, 331(6018):717-719, 2011.
9. Barros, J K., Babich, R., Brower, R., Clark, M. A., and Rebbi, C., Blasting through lattice calculations using CUDA, PoS **LATTICE2008**, 045, 2008. arXiv:0810.5365 [hep-lat].
10. Bernholdt, D., Bharathi, S., Brown, D., Chanchio, K., Chen, M., Chervenak, A., Cinquini, L., Drach, B., Foster, I., Fox, P., Garcia, J., Kesselman, C., Markel, R., Middleton, D., Nefedova, V., Pouchard, L., Shoshani, A., Sim, A., Strand, G. and Williams, D. The Earth System Grid: Supporting the Next Generation of Climate Modeling Research. *Proceedings of the IEEE*, 93(3):485-495, 2005.
11. Bose, R., Foster, I. and Moreau, L. Report on the International Provenance and Annotation Workshop (IPAW'06). *SIGMOD Record*, 35(3):51-53, 2006.
12. Bresnahan, J., Link, M., Kettimuthu, R., Fraser, D. and Foster, I., GridFTP Pipelining. TeraGrid Conference, Madison, WI, 2007.
13. Brower, R. C. Rebbi, C., and Vicari, E., Projective multigrid method for propagators in lattice gauge theory, *Phys. Rev. D* 43, 1965, 1991.
14. Buneman, P., Khanna, S. and Tan, W.-C., Why and Where: A Characterization of Data Provenance. International Conference on Database Theory, 2001.
15. Catlett, C., Beckman, P., Skow, D. and Foster, I. Creating and Operating National-Scale Cyberinfrastructure Services. *CTWatch Quarterly*, 2(2), 2006.
16. Chervenak, A., Schuler, R., Kesselman, C., Koranda, S. and Moe, B., Wide Area Data Replication for Scientific Collaborations. 6th IEEE/ACM Int'l Workshop on Grid Computing 2005.
17. Childers, L., Disz, T., Olson, R., Papka, M.E., Stevens, R. and Udeshi, T. Access Grid: Immersive Group-to-Group Collaborative Visualization. 4th International Immersive Projection Technology Workshop, 2000.

18. Clark, M. A., Babich, R., Barros, K., Brower, R. C., Rebbi, C., Solving Lattice QCD systems of equations using mixed precision solvers on GPUs, *Comput. Phys. Commun.* **181**, 1517–1528, 2010. arXiv:0911.3191 [hep-lat].
19. Cummings, J.N. and Kiesler, S. Collaborative Research Across Disciplinary and Organizational Boundaries. *Social Studies of Science*, 25(5):703-722, 2005.
20. CVS: <http://savannah.nongnu.org/projects/cvs/>.
21. Dubey, A., Antypas, K., Ganapathy, M., Reid, L. B., Riley, K., Sheeler, D., Siegel, A., Weide, K. Extensible component-based architecture for FLASH, a massively parallel, multiphysics simulation code. *Parallel Computing* 35 (10-11), 512–522, 2009.
22. Fonck, R., et al., Remote Operation of the TFTR BES Experiment From an Off-Site Location, *Rev. Sci. Instrum.* 63, 4803, 1992.
23. Foster, I. Globus Online: Accelerating and democratizing science through cloud-based services. *IEEE Internet Computing*(May/June):70-73, 2011.
24. Foster, I., Kesselman, C., Tsudik, G. and Tuecke, S., A Security Architecture for Computational Grids. 5th ACM Conference on Computer and Communications Security, 1998, 83-91.
25. Fryxell, B., Olson, K., Ricker, P., Timmes, F., Zingale, M., Lamb, D., MacNeice, P., Rosner, R., Truran, J., Tufo, H.. FLASH: Adaptive mesh hydrodynamics code. *The Astrophysical Journal Supplement Series* 131, 273–334, 2000.
26. Horne, S., et al., Remote Control of Alcator C-Mod from LLNL, *Fusion Technol.* 32, 52, 1997.
27. Hudson, R., Norris, J., Reid, L.B., Weide, K., Jordan, G., Papka, M. Experiences Using Smaash to Manage Data-Intensive Simulations. *Proceedings of HPDC*, 2011.
28. Jordan, G.C., Fisher, R.T., Townsley, D.M., Calder, A.C., Graziani, C., Asida, S., Lamb, D.Q. and Truran, J.W. Three-Dimensional Simulations of the Deflagration Phase of the Gravitationally Confined Detonation Model of Type Ia Supernovae. *The Astrophysical Journal*, 681(2):1448, 2008.
29. Keahey, K., Fredian, T., Peng, Q., Schissel, D.P., Thompson, M., Foster, I., Greenwald, M. and McCune, D. Computational Grids in Action: The National Fusion Collaboratory. *Future Generation Computer Systems*, 18(8):1005-1015, 2002.
30. Lingerfelt, E., Messer, O.E.B., Osborne, J.A., Budiardja, R.D., Mezzacappa, A. A multitier system for the verification, visualization and management of chimera. *Procedia Computer Science* 4 (0), 2076–2085, 2011.
31. LSST Science Requirements Document. Large Synoptic Survey Telescope Corporation, 2007.
32. Lustig, M., Donoho, D.L., Santos, J.M. and Pauly, J.M. Compressed Sensing MRI. *Signal Processing Magazine, IEEE*, 25(2):72-82, 2008.
33. Maia, F., MacDowell, A., Marchesini, S., Padmore, H., Parkinson, D., Pien, J., Schirotzek, A., Yang, C. Compressive Phase Contrast Tomography. *Proc. SPIE Vol. 7800, 78000F* (2010).
34. McHarg, B.B., et al., Tools for Remote Collaboration on the DIII-D National Fusion Facility, *Fusion Eng. & Design* 42, 343, 1999.
35. MDSplus: <http://www.mdsplus.org/>.
36. Messer, O.E.B., Bruenn, S.W., Blondin, J.M., Hix, W.R., Mezzacappa, A., Dirk, C.J. Petascale supernova simulation with CHIMERA. *Journal of Phys.* 78, 2049, 2007.
37. Messer, O.E.B., Bruenn, S.W., Blondin, J.M., Chertkow, M.A., Hix, W.R., Lee, C., Lentz, E.J., Marronetti, P., Yakunin, K.N. Core-Collapse Supernova Simulations with CHIMERA. In: Martinez-Pinedo, G., et al. (Eds.), *Proceedings of Nuclei in the Cosmos XI. SISSA Proceedings of Science*, p. 027, 2011.
38. Moreau, L., Clifford, B., Freire, J., Futrelle, J., Gil, Y., Groth, P., Kwasnikowska, N., Miles, S., Missier, P., Myers, J., Plale, B., Simmhan, Y., Stephan, E. and Bussche., J.V.d. The open provenance model core specification (v1.1). *Future Generation Computer Systems*, July, 2010.

39. National Research Council, *An Assessment of the Department of Energy's Office of Fusion Energy Sciences Program*, National Academy Press, 2000.
40. Novotny, J., Tuecke, S. and Welch, V., An Online Credential Repository for the Grid: MyProxy. 10th IEEE International Symposium on High Performance Distributed Computing, San Francisco, 2001, IEEE Computer Society Press.
41. Peng, R.D. Reproducible Research in Computational Science. *Science*, 334(1226), 2011.
42. Plewa, T., Calder, A.C. and Lamb, D.Q. Type Ia Supernova Explosion: Gravitationally Confined Detonation. *The Astrophysical Journal*, 612(1):L37-L40, 2004.
43. Pordes, R., Petravick, D., Kramer, B., Olson, D., Livny, M., Roy, A., Avery, P., Blackburn, K., Wenaus, T., Würthwein, F., Foster, I., Gardner, R., Wilde, M., Blatecky, A., McGee, J. and Quick, R., The Open Science Grid. Scientific Discovery through Advanced Computing (SciDAC) Conference, 2007.
44. QUDA (QCD in CUDA) library and github repository: <https://github.com/lattice/quda>.
45. Romberg, J. Compressed Sensing Creates a Buzz at ICIAM '07. *SIAM News*, 40(8), October 2007.
46. Saltz, J., Hastings, S., Langella, S., Oster, S., Ervin, D., Sharma, A., Pan, T., Gurcan, M., Permar, J., Ferreira, R., Payne, P., Catalyurek, U., Kurc, T., Caserta, E., Leone, G., Ostrowski, M., Madduri, R., Foster, I., Madhavan, S., Buetow, K.H., Shanbhag, K. and Siegel, E. e-Science, caGrid, and Translational Biomedical Research. *IEEE Computer*, 2008.
47. Schissel, D.P. Grid Computing and Collaboration Technology in Support of Fusion Energy Sciences. *Phys. Plasmas* 12(058104), 2005.
48. SciDAC Outreach Center: <http://www.scidac.gov/SDoutreach.html>.
49. Simmhan, Y.L., Plale, B. and Gannon, D. Karma2: Provenance Management for Data-Driven Workflows. *International Journal of Web Services Research*, 5(2):1-22, 2008.
50. Welch, V., Siebenlist, F., Foster, I., Bresnahan, J., Czajkowski, K., Gawor, J., Kesselman, C., Meder, S., Pearlman, L. and Tuecke, S., Security for Grid Services. 12th IEEE International Symposium on High Performance Distributed Computing, 2003.
51. Winslow, R.L., Saltz, J., Foster, I., Carr, J.J., Ge, Y., Miller, M.I., Younes, L., Geman, D., Graniote, S., Kurc, T., Madduri, R., Ratnanather, T., Larkin, J., Ardekani, S., Brown, T., Kolasny, A., Reynolds, K., Shipway, M. and Toerper, M. The CardioVascular Research Grid Project. *AMIA Summits Transl Sci Proc*, 2011.
52. Wuchty, S., Jones, B.F. and Uzzi, B. The Increasing Dominance of Teams in Production of Knowledge. *Science*, 316(5827):1036-1039, 2007.