

Lawrence Berkeley National Laboratory

Lawrence Berkeley National Laboratory

Title

BER Science Network Requirements

Permalink

<https://escholarship.org/uc/item/8184z610>

Author

Dart, Eli



Publication Date

2010-11-01

BER Science Network Requirements

Report of the Biological and Environmental Research
Network Requirements Workshop

Conducted April 29-30, 2010



DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

BER Network Requirements Workshop

Office of Biological and Environmental Research, DOE Office of Science
Energy Sciences Network
Rockville, MD — April 29 and 30, 2010

ESnet is funded by the US Department of Energy, Office of Science, Office of Advanced Scientific Computing Research (ASCR). Vince Dattoria is the ESnet Program Manager.

ESnet is operated by Lawrence Berkeley National Laboratory, which is operated by the University of California for the US Department of Energy under contract DE-AC02-05CH11231.

This work was supported by the Directors of the Office of Science, Office of Advanced Scientific Computing Research, Facilities Division, and the Office of Biological and Environmental Research.

This is LBNL report LBNL-4089E

Participants and Contributors

Kiran Alapaty, DOE/SC/BER
(Atmospheric System Research)

Ben Allen, LANL (Bioinformatics)

Greg Bell, ESnet (Networking)

David Benton, GLBRC/University of
Wisconsin (Informatics)

Tom Brettin, ORNL (Bioinformatics)

Shane Canon, NERSC (Data Systems)

Rich Carlson, DOE/SC/ASCR (Network
Research)

Steve Cotter, ESnet (Networking)

Silvia Crivelli, LBNL (JBEI)

Eli Dart, ESnet (Networking)

Vince Dattoria, DOE/SC/ASCR (ESnet
Program Manager)

Narayan Desai, ANL (Networking)

Richard Egan, ANL (ARM)

Jeff Flick, NOAA (Networking)

Ken Goodwin, PSC/NLR (Networking)

Susan Gregurick, DOE/SC/BER
(Computational Biology)

Susan Hicks, ORNL (Networking)

Bill Johnston, ESnet (Networking)

Bert de Jong, PNNL (EMSL/HPC)

Kerstin Kleese van Dam, PNNL (Data
Management)

Miron Livny, University of Wisconsin
(Open Science Grid)

Victor Markowitz, LBNL/JGI
(Genomics)

Jim McGraw, LLNL (HPC/Climate)

Raymond McCord, ORNL (ARM)

Chris Oehmen, PNNL
(Bioinformatics/ScalaBLAST)

Kevin Regimbal, PNNL
(Networking/HPC)

Galen Shipman, ORNL (ESG/Climate)

Gary Strand, NCAR (Climate)

Brian Tierney, ESnet (Networking)

Susan Turnbull, DOE/SC/ASCR
(Collaboratories, Middleware)

Dean Williams, LLNL (ESG/Climate)

Jason Zurawski, Internet2 (Networking)

Editors

Eli Dart, ESnet — dart@es.net

Brian Tierney, ESnet — bltierney@es.net

Table of Contents

1	Executive Summary.....	6
2	Workshop Background and Structure	8
3	Office of Biological and Environmental Research.....	10
4	B-6, Biosciences Division, Los Alamos National Laboratory.....	12
5	Joint Global Change Research Institute.....	16
6	Climate Science—Enabling Worldwide Access to Petascale Climate Data.....	26
7	Network Requirements of the DOE Systems Biology Knowledgebase	42
8	Joint Genome Institute, Walnut Creek, CA	51
9	EMSL, PNNL	55
10	Great Lakes Bioenergy Research Center	62
11	Joint BioEnergy Institute, Emeryville, CA.....	70
12	BioEnergy Science Center.....	77
13	NOAA’s Climate Modeling Program.....	84
14	Atmospheric Radiation Measurement (ARM) Climate Research Facility.....	87
15	Metagenomics, Argonne National Laboratory.....	96
16	Findings	99
17	Requirements Summary and Conclusions	102
18	Glossary	103
19	Acknowledgements	104

1 Executive Summary

The Energy Sciences Network (ESnet) is the primary provider of network connectivity for the US Department of Energy Office of Science, the single largest supporter of basic research in the physical sciences in the United States. In support of the Office of Science programs, ESnet regularly updates and refreshes its understanding of the networking requirements of the instruments, facilities, scientists, and science programs that it serves. This focus has helped ESnet to be a highly successful enabler of scientific discovery for over 20 years.

In April 2010 ESnet and the Office of Biological and Environmental Research, of the DOE Office of Science, organized a workshop to characterize the networking requirements of the science programs funded by BER.

The requirements identified at the workshop are summarized below, and described in more detail in the case studies and the Findings section.

Common Themes

- As in other areas of science, many BER programs have data sets that are growing exponentially. BER scientists are also typical in that a number of them are having trouble fully utilizing the network to handle their data transfer requirements, and transfers often take longer than expected based on available network capacities.
- There is likely to be a significant increase in the use of commercial clouds by BER-funded scientists over the next few years, as a number of case studies mentioned this growing need. In particular the Genomics community is seriously looking into using commercial cloud services
- Another common theme was that the ARM project, EMSL, and others discussed the problems of cross-site authentication and authorization.
- Participants also pointed out the need for middleware that would transfer a specified group of files from system A to system B with *minimal* oversight from the researcher.

Climate Science

- The biggest BER user of ESnet in the future is likely to be the climate community. Climate sites (e.g. LLNL, ORNL) will host huge data sets that will be accessed from all over the world. By 2020, it is estimated that the climate data repositories at these sites will exceed hundreds of exabytes in size. Several key sites for distributing this data are the CMIP-5 Data Centers at LLNL, BADC (U.K.), DKRZ (Germany), ANU (Australia), and JAMSTEC (Japan). NCAR and LBNL will distribute data from the CCSM model, and ORNL will distribute observational and other data. Data will be replicated across these and other locations, and bandwidth requirements will likely reach 100 Gbps in the next few years.
- NOAA is likely to be a significant source of data for climate researchers.

- JGCRI anticipates increasing use of supercomputing resources at ORNL, with data export to PNNL and UMD.
- ARM data will be increasing as new ARRA funded instruments come online.

Genomics

- The cost of genome sequencers is falling dramatically, and sequencers are continuing to increase in data output exponentially (one estimate is for a 12x increase in data volume over the next 5 years). Sequencers will no doubt become more common and exist at a lot more locations. Data management for these instruments will likely become a difficult issue for many institutions. These include the use of commercial clouds for genomics data analysis
- The Genomics community requires good connectivity to the National Center for Biotechnology Information (NCBI) in Bethesda, MD.

Biosciences

- Data transfer is a bottleneck for LANL B6.

Bioenergy

- JBEI has significant need for more efficient distribution of large visualization data sets.

The DOE Office of Science sites with the largest future bandwidth requirements for BER research are LLNL and ORNL. Other sites with significant projected growth are LBNL, NERSC, and PNNL.

2 Workshop Background and Structure

The strategic approach of the Office of Advanced Scientific Computing Research (ASCR – ESnet is funded by the ASCR Facilities Division) and ESnet for defining and accomplishing ESnet’s mission involves three areas:

1. Work with the SC community to identify the networking implication of the instruments, supercomputers, and the evolving process of how science is done
2. Develop an approach to building a network environment that will enable the distributed aspects of SC science and then continuously reassess and update the approach as new requirements become clear
3. Keep anticipating future network capabilities that will meet future science requirements with an active program of R&D and Advanced Development

Addressing point (1), the requirements of the Office of Science science programs are determined by:

A) Exploring the plans and processes of the major stakeholders, including the data characteristics of scientific instruments and facilities, regarding what data will be generated by instruments and supercomputers coming on-line over the next 5-10 years. Also by examining the future process of science: how and where will the new data be analyzed and used, and how the process of doing science will change over the next 5-10 years.

B) Observing current and historical network traffic patterns and trying to determine how trends in network patterns predict future network needs.

The primary mechanism of accomplishing (A) is the Office of Science (SC) Network Requirements Workshops, which are sponsored by ASCR and organized by the SC Program Offices. SC conducts two requirements workshops per year, in a cycle that repeats every three years:

- Fusion Energy Sciences (2008)
- Nuclear Physics (2008)
- Advanced Scientific Computing Research (2009)
- High Energy Physics (2009)
- Biological and Environmental Research (2010)
- Basic Energy Sciences (materials sciences, chemistry, geosciences) (2010)

The workshop reports are published at <http://www.es.net/hypertext/requirements.html>.

The other role of the requirements workshops is that they ensure that ESnet and ASCR have a common understanding of the issues that face ESnet and the solutions that ESnet undertakes.

In April 2010 ESnet and the Office of Biological and Environmental Research, of the DOE Office of Science, organized a workshop to characterize the networking requirements of the science programs funded by BER.

Workshop participants were asked to codify their requirements in a case study format that included a network-centric narrative describing the science, the instruments and facilities

currently used or anticipated for future programs, the network services needed, and the way in which the network is used. Participants were asked to consider three time scales in their case studies — the near term (immediately and up to 12 months in the future), the medium term (two to five years in the future), and the long term (greater than five years in the future). The information in each narrative was distilled into a summary table, with rows for each time scale and columns for network bandwidth and services requirements. The case study documents are included in this report.

3 Office of Biological and Environmental Research

The Office of Biological and Environmental Research (BER) within the Department of Energy's Office of Science advances world-class biological and environmental research and provides scientific user facilities to support our missions in scientific discovery and innovation, energy security and environmental responsibility. BER comprises two Divisions, Climate and Environmental Sciences Division (CESD) and Biological Systems Science Division (BSSD) to address national needs in four key areas (Sustainable Biofuels, Climate Science, Subsurface Biogeochemistry and the Interface of Biological and Physical Sciences)

The Climate and Environmental Sciences Division (CESD) focuses on developing a predictive, systems-level scientific understanding of climate change and the subsurface environment. This includes research and modeling efforts to (1) improve understanding of factors affecting the Earth's radiant-energy balance; (2) predict accurately any global and regional climate change induced by increasing atmospheric concentrations of aerosols and greenhouse gases; (3) quantify sources and sinks of energy-related greenhouse gases, especially carbon dioxide; and (4) improve the scientific basis for assessing both the potential consequences of climatic changes, including the potential ecological, social, and economic implications of human-induced climatic changes caused by increases in greenhouse gases in the atmosphere and the benefits and costs of alternative response options. Research in this area is focused on understanding the basic chemical, physical, and biological processes of the Earth's atmosphere, land, and oceans and how these processes may be affected by energy production and use, primarily the emission of carbon dioxide from fossil fuel combustion. Activities within the CESD Division include:

- Atmospheric System Research activity seeks to resolve the two major areas of uncertainty in climate change model projections: the role of clouds and the effects of aerosol emissions on the atmospheric radiation balance.
- Environmental System Science activity supports research that provides scientific understanding of the effects of climate change on terrestrial ecosystems, the role of terrestrial ecosystems in global carbon cycling, and the role of subsurface biogeochemical processes on the fate and transport of DOE-relevant contaminants including heavy metals and radionuclides.
- Climate and Earth System Modeling focuses on development, evaluation, and use of large scale climate change models to determine the impacts of climate change and mitigation options.
- Two scientific user facilities—the Atmospheric Radiation Measurement Climate Research Facility (ARM) and the Environmental Molecular Sciences Laboratory (EMSL)—provide the broad scientific community with technical capabilities,

scientific expertise, and unique information to facilitate science in areas integral to the BER mission and of importance to DOE.

The Biological Systems Science Division (BSSD) provides the scientific and analytical technologies needed to translate genomic sequence into a profound and comprehensive understanding of the myriad of processes carried out by biological systems. Research within this program leverage the genomic code as a starting point to understand systems biology through (1) Systems analysis of the collective ‘omics (e.g. transcriptomics, proteomics and metabolomics) of plants and microbes, (2) Development of new methods for characterizing and imaging molecular systems, and (3) Development of new approaches to synthesize and redesign biology processes.

Achieving a predictive understanding of systems biology requires the integration of an immense amount of highly coupled and correlated types of data and information and this represents a grand challenge in biology. To address this challenge DOE is developing the Systems Biology Knowledgebase which is a cyber-infrastructure for integrating, accessing, analyzing and modeling of complex biological systems and behaviors. Specific research activities within the BSSD Division include:

- Genomic Science Research - this program develops genome-scale technologies that are needed to understand the function of microbial and plant systems, from proteomics to metabolomics to regulatory networks to ecogenomics. The Genomic Science goal is to develop the computational capabilities and systems needed to predicatively design and model complex biological systems.
- DOE Bioenergy Research Centers – will accelerate genomics-based systems biology research to achieve the transformational breakthroughs in basic science needed for the development of cost-effective technologies to make production of next-generation biofuels from lignocellulose, or plant fiber, commercially viable on a national scale.
- Two scientific user facilities- the Structural Biology and the Joint Genome Institute (JGI) provide the broad scientific community with technical capabilities in biological structure determination and in genomic sequencing in support of the DOE missions in energy, climate and the environment.

A significant scientific asset for BER is the national laboratory system which maintains our research leadership in climate and the environment. In addition to national laboratory led activities, BER engages larger scientific activities that are academic and university partnerships including the Earth System Grid (ESG) and the Bioenergy Research Centers (BRC). A key feature of these activities and facilities include the integration of data and models and the broadcasting this information to their respective communities. ESnet provides a vital component to the transfer of data and the connectivity of scientists to these programs. This workshop has identified a number of critical strengths and gaps that ESnet can fulfill and build upon for the coming years as BER looks towards a stronger future in climate research, sustainable bioenergy and stewardship in the environment.

4 B-6, Biosciences Division, Los Alamos National Laboratory

Background

LANL is the second largest partner institution of the Department of Energy's Joint Genome Institute (DOE-JGI), and specializes in high throughput genome finishing and analysis in support of DOE missions in energy, bioremediation and carbon sequestration. B-6 is the Genome Sciences Group at LANL in the Bioscience Division. Its comprised of molecular biology labs and computational staff who together focus on the high-throughput DNA sequencing of whole microbial genomes, computational finishing and bioinformatics.

Key Local Science Drivers

Instruments and Facilities

B-6 currently utilizes two server rooms. The internal network consists primary of standard 1000Base-TX or fiber equivalent. B-6 owns roughly 50 servers ranging from five-year-old Sun Sparc hardware to new Dell and Sun hardware. There are a number of fiber channel or directly attached storage devices, with a rough total of 100TB of usable storage space. Most of the storage is shared via NFS to various computational servers.

B-6 has a number of genome sequencers attached to the network, which typically dump data to a local workstation. The data are subsequently batch transferred to a storage server. Another configuration with the genome sequencers is to have the controller workstation, which is attached to the instrument, save data directly to one of the storage servers. This is done with the Illumina. Currently, B-6 has 1st and 2nd generation sequencing technology. 1st generation gives 70 KiloBases of sequence per run (8 runs possible per day). 2nd generation gives 500 MegaBases (2 runs possible per day) to 10 GigaBases of sequence per run (1 run every 3-7 days). Note: base is a measure of base pairs sequenced.

Process of Science

Once data is onsite from collaborators, it enters the B-6 pipeline. The pipeline is a set of automated processes meant to process genomic data. Collaborator's data goes through an automated startup process and is then handed off to the Finishing Team. Once in the manual finishing process, a staff member works with the data and has a number of options on how to proceed. Typically a Finisher can request further sequencing be made on instrumentation at B-6, or request the data be run through various applications within the pipeline. The typical project while onsite ranges from 40GB to 250GB of data. Throughout the B-6 finishing process the data is stored in-house, and largely never leaves LANL. However there is current work in progress to automate the uploading of a draft version of the project at various steps during our process to allow our collaborators at JGI

and elsewhere to monitor and work with the data at intermediate steps of completion. Initially this will include only 30-40MB of data per project, but there is a possibility that B-6's collaborators will want all the current raw data of each project at these steps.

Key Remote Science Drivers

Instruments and Facilities

B-6 mainly collaborates with other entities within the Joint Genome Institute (JGI), the Defense Threat Reduction Agency (DTRA), and the Department of Homeland Security (DHS). As far as data transfer, data is moved to and from Lawrence Berkeley National Lab on a regular basis. In addition B-6 is starting to move data to the Air Force Research Laboratory (AFRL) on a more regular basis.

Process of Science

The current automated processes at B-6 download and upload roughly 10-15GB daily of compressed data to or from the Joint Genome Institute, specifically servers located at Lawrence Berkeley National Laboratory (LBNL). A standard Rsync daemon is setup at LBNL for B-6 to automatically pull down any new or changed files. Project download sizes range from 500MB to 7GB, and average 3.1GB based on the last 11 projects downloaded. On average five projects are downloaded a week. However this rate is starting to increase to around 10 projects per week. Typical download speeds seen using Rsync are 1.0 MB/sec.

Once projects are finished with a completed whole genome, they are manually uploaded to a JGI server using secure copy (SCP). The typical upload speed for these transfers are 1.0 MB/sec. Since the beginning of 2010, finished projects sent out averaged 2.9GB in size, and at least 4 projects are uploaded a week.

Local Science Drivers – the next 2-5 years

Instruments and Facilities

As genome sequencing is a fast developing field, it is likely that B-6 would bring in new instruments as time progresses. These new instruments will produce greater volumes of data, and subsequently increase our storage needs, data transfer needs, and computational needs. B-6 will be running 3rd generation sequencing technology in the 2-5 years timeframe. 3rd generation tech at first release will be able to generate 150MegaBases every 15 minutes, and this will likely increase rapidly. These machines produce so much raw data that currently the attached servers will be unable to store the streamed data. As a result the instrument itself converts the raw data into more usable and storable formats before writing the data to disk.

In the 2-5 year timeframe we will still be using the same physical facilities, with an upgraded 10 Gbps local network.

Process of Science

Process will largely stay the same in the next 2-5 years. B-6 informatics systems will be adapted and updated to new sequencing technology as required. Its likely these updates will not require any great increase in the use of ESnet as most our data moves internally. In general, B-6's goals have the group doubling its output (in all areas of sequencing, finishing, and analysis) every 18 months.

Remote Science Drivers – the next 2-5 years

Instruments and Facilities

The future outlook from collaborators is variable. However the general consensus is the amount of data for projects will generally increase. In addition, it seems that B-6 at LANL is slowly diversifying its work, bringing in new collaborators.

Process of Science

The expected amount of daily data transferred between collaborators in the 2-5 year timeframe is double or triple of what it is now. Other remote process changes in the 2-5 year timeframe are currently unknown.

Beyond 5 years – future needs and scientific direction

As genome sequencing and analysis is a very rapidly developing field, the outlook for future needs beyond 5 years is variable at best. However, its fair to say the amount of data transferred from collaborators, data generated in-house, and data moved back from offsite will only increase. The rate of this increase is unknown at this point. However, it is expected that within 5 years technology will allow the sequencing of a human sized genome in 15 minutes. A human genome is estimated to be 3 billion base pairs, so in 5 years we should have the technology that allows for 12 fold increase of data generation over the new (Summer 2010) 3rd generation sequencing.

Middleware Tools and Services

The current data transfer speed to and from JGI and other collaborators is considered a bottleneck in B-6's current processes. Any common file transfer tools or methods that could increase the speed of file transfer would be of great interest. However these tools would need to be easy to setup, use, and maintain. If a standardized method to transport large datasets could be introduced across ESnet, likely it would alleviate some of the administrative issues of working with new collaborators. In other words, if an agreed upon standard for moving files between institutes was decided, there would be no need to setup new services each time a new collaboration was made.

Outstanding Issues

If B-6 was called upon to start transferring whole projects with all the data ever generated at LANL for each project, the upload transfer amount would suddenly spike to about 1TB a week at the current number of finished projects a week. Most if not all collaborators do

not ask for this amount of data. However, if it was timely and easy to move, B-6 would likely be asked to move this amount of data, about 250GB raw per project, more often.

Summary Table

Feature	Key Science Drivers		Anticipated Network Requirements	
	Science Instruments and Facilities	Process of Science	Local Area Network Bandwidth and Services	Wide Area Network Bandwidth and Services
Near-term (0-2 years)	<ul style="list-style-type: none"> • 2 Server Rooms • 100TB of Storage • 1Gbps LAN • ~50 Intel and Sparc Servers • Genome Sequencers 	<ul style="list-style-type: none"> • High Throughput Finishing of Whole Genomes 	<ul style="list-style-type: none"> • 10 Gbps LAN • 25-50TB of additional storage 	<ul style="list-style-type: none"> • 40-50GB Transferred a Week • Common Large Data Transfer Tools
2-5 years	<ul style="list-style-type: none"> • Additional Genome Sequencing, 3rd Gen. Instrumentation • 10 Gbps LAN • Additional Storage • Additional Computation Servers 	<ul style="list-style-type: none"> • Largely Similar to Near Term • Adaptation to New Sequencing Technology and Techniques 	<ul style="list-style-type: none"> • Expanded 10 Gbps Service to Onsite Data Centers 	<ul style="list-style-type: none"> • 120GB-150GB Transferred a Week
5+ years	<ul style="list-style-type: none"> • Largely Unknown due to high rate of development in genome sequencing and analysis field • Expected 12x Increase of Data Output over 3rd Gen Instruments 	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> •

5 Joint Global Change Research Institute

Pacific Northwest National Laboratory and the University of Maryland

Background

The U.S. Department of Energy (DOE), in collaboration with other agencies and institutions worldwide, conducts research to improve and extend capabilities for integrated assessment of climate change and undertakes assessments to address climate change issues and questions, with current emphasis on mitigation, adaptation, and vulnerability. Such assessments span the natural phenomena of the climate system and the human factors that affect drivers of climate change, primarily emissions of greenhouse gases into the atmosphere. Thus, a key feature of integrated assessment research is its comprehensive nature—integrating concepts, data, models and results from a number of disciplines involved in climate change research to enable assessments of broad issues and questions.

The Joint Global Change Research Institute (JGCRI), a joint activity of the Pacific Northwest National Laboratory (PNNL) and the University of Maryland is located in College Park, Maryland. JGCRI is a lead organization within the DOE Climate Change Integrated Assessment Research Program, and maintains a widely recognized integrated assessment model, the Global Change Assessment Model (GCAM). In addition to GCAM, JGCRI uses models of agricultural productivity and of terrestrial ecosystem hydrology and biogeochemistry to investigate aspects of land use, maintenance and utilization of natural resources and other ecosystem services, and food production. Many studies focus on options for mitigation of greenhouse gas increases through carbon sequestration from the atmosphere and storage.

With increasing interest in climate change mitigation, adaptation, and vulnerability, JGCRI activities and the DOE Climate Change Integrated Assessment Research Program are expanding with focus on six major challenges:

- Linking and integrating mitigation, transformational science, and complex interactions;
- Incorporating impacts, adaptation and vulnerability into integrated assessments;
- Improving spatial and temporal resolutions;
- Developing better methods for representing risk and uncertainty, and improving diagnostic methods;
- Building interoperable and accessible modeling frameworks and collaborations; and
- Expanding data development and accessibility.

Key Local Science Drivers

The GCAM model, associated data, and analyses are developed and maintained by a team of JGCRI scientists and engineers working with collaborators at a number of other

institutions. The model is implemented within an object-oriented structure that enables research and development by different investigators with expertise in particular model components such as specific energy resources or technologies, economics of certain sectors, or aspects of the Earth system that determine climate consequences of emissions. The diverse data required by the model are similarly organized to permit coordinated data acquisition and management by specialists in each area.

Analyses using the GCAM model are tied to global climate or Earth system models. Integrated assessment models such as GCAM are used to devise scenarios of greenhouse gas emissions that drive climate and Earth system model solutions. In turn, integrated assessment models must be consistent with the representations of natural phenomena within climate models and the simulations of climate change those models produce. Thus, there is coupling between integrated assessment models and climate models through exchange of data and model results.

Much of the data used in integrated assessment research is geospatial, and in many instances involves a temporal dimension as well. The management and utilization of geospatial data is computationally demanding. Coordinate transformations and operations to georeference data are commonplace. Analysis and interpretation require mapping geospatial information and visualization of multiple data layers. Global simulations over climatic time periods produce large volumes of geospatial data that must be transformed into common coordinates and georeferenced.

Within the U.S., there are three major integrated assessment models, and several others are maintained and applied by foreign institutions. A key goal of the DOE integrated assessment research program is to coordinate the development and application of these models much more closely. Rigorous model inter-comparison studies are expected, and these will require the assembly of common input data, development of comparison protocols, and a coordinated evaluation of model results. Gaining a better understanding of error and uncertainty in both input data and in model results is also a high priority for the DOE program. These multi-institutional studies will require transfer of substantial amounts of data over the Internet and capabilities for investigator collaboration in coordinated model runs occurring at JGCRI.

The DOE is working with the integrated assessment research community toward the development of open-source, community-developed, community-maintained models. This activity will require collaboration in programming and data management over the Internet typical of that underway within the Earth system modeling community, and DOE is asking JGCRI to foster these collaborations.

Instruments and Facilities

Currently, JGCRI model development and applications rely primarily on the desktop computing systems of institute investigators. Additionally, two servers provide shared storage for data and code management, respectively, and for data exchanges. Recently, JGCRI acquired a high-end server/computational platform for solutions to high-spatial resolution models of agricultural and ecosystem biogeochemistry required for assessment of options for biomass production to meet energy needs.

Recognizing the expanding role for integrated assessment of climate change, the DOE is providing a high-performance computing and storage system to JGCRI that will be a community resource for integrated assessment research and applications. This system will comprise a 240-node computing cluster coupled to a 1 Petabyte mass storage system. The system will be available to users by October 1, 2010.

While JGCRI investigators will make substantial use of this high-end resource, it is intended to serve the broader integrated assessment community, especially for collaborative activities and to enable more demanding computations required to increase model realism. The mass storage system will allow the community to assemble commonly used data sets and model results in consistent formats, including those at increasingly high spatial resolutions.

To enhance user capabilities and options for collaboration, JGCRI is establishing a data analysis and visualization laboratory with six high-end workstations. A similar laboratory is being established at PNNL in Richland, Washington. One focus in the configuration of these laboratories is to extend capabilities for management, manipulation, analysis, and display of geospatial data, including datasets with a high-resolution temporal dimension. A wall-scale display will allow visualization of multiple geospatial views with global coverage. Laboratory computers at JGCRI and at PNNL in Richland, Washington will be configured to enable close collaboration over the Internet, including simultaneous view and manipulation of graphics and documents coupled with desktop videoconference.

Process of Science

DOE climate change integrated assessment research and applications make use of models that represent aspects of the Earth system and incorporate human factors, processes, and activities that are responsible for the drivers of climate change. Although there is a tight relationship between integrated assessment models and the full Earth system models used to simulate climate change in response to scenarios of radiative forcing due to changes in atmospheric greenhouse gas levels, integrated assessment models treat natural processes in a more approximate manner to allow far more rapid solutions and coupling to detailed representations of energy and economic systems, demography, land use, agriculture, and natural resources.

An important product developed using integrated assessment models is scenarios of future drivers of climate change, primarily changes in Earth's radiation balance due to greenhouse gas emissions into the atmosphere. These scenarios are used within major assessments such as those of the Intergovernmental Panel on Climate Change (IPCC) as inputs to full Earth system models that simulate future climate change expected to result under each scenario. Recently, JGCRI completed development of a representative concentration pathway that stabilizes radiative forcing of climate at 4.5 W/m^2 . This scenario is one of a group that will be the basis for the 5th IPCC assessment.

Scrutiny of the details of IPCC scenarios will continue over the next several years, and as components of the 5th assessment are completed, numerous questions and issues regarding relationships between scenarios and the implied climate changes and their impacts will be addressed through further studies using the GCAM model.

The DOE integrated assessment research program will be engaged in similarly organized assessments at different scales and to address a variety of climate change questions and issues, including a U.S. national assessment now in planning stages.

Key Remote Science Drivers

JGCRI integrated assessment research and applications involve data acquisition from numerous and diverse sources, distribution of numerical results, and collaboration with investigators at other institutions working on the GCAM model or undertaking studies using GCAM or other models developed and maintained by JGCRI. Additionally, JGCRI coordinates integrated assessment model inter-comparisons involving multiple institutions.

Applications of climate change integrated assessment models are tied to simulations of climate change by Earth system models, and research is underway to fully couple these models. A major DOE activity involves scientists from multiple national laboratories. As a result, JGCRI utilizes supercomputers located at the Oak Ridge National Laboratory (ORNL), and connectivity between JGCRI and ORNL is particularly critical to this work.

Instruments and Facilities

The JGCRI makes continual use of climate simulations derived using Earth system models. The Community Climate System Model (CCSM) developed and maintained by the National Center for Atmospheric Research (NCAR) in Boulder, Colorado is particularly important. DOE is funding JGCRI and groups at several other national laboratories (ANL, LBNL, ORNL, PNNL) to couple the GCAM integrated assessment model with the CCSM. Most of the development work to enable this coupling is occurring at ORNL, and supercomputers located at ORNL are being used for model solutions.

JGCRI relies on the Earth System Grid for access to major climate simulation results, especially those prepared for assessments by the Intergovernmental Panel on Climate Change (IPCC). The IPCC Data Distribution Centre is also an important distribution site for data used by IPCC assessments.

JGCRI researchers are making increasing use of satellite remote sensing data to characterize land surface attributes, land use, and other human activities. Remote sensing data products are accessed from NASA's Earth Science Distributed Active Archive Centers (DAACs). Larger acquisitions are from the Land Processes DAAC in Sioux Falls, South Dakota, the Global Land Cover Facility in College Park, Maryland, and the DAAC for Biogeochemical Dynamics at ORNL.

Because of the interdisciplinary nature of integrated assessment research, data are drawn from numerous data centers and portals. As the diversity and volume of data have increased, it is becoming difficult to manage data sets from many sources to insure that latest versions are incorporated. Reformatting and in the case of geospatial data coordinate transformations have become a major computational demand. JGCRI intends to assemble many commonly used data sets on new mass storage provided by DOE and to automate many of the processes of data ingest to enable more robust and reliable applications of models and integrated assessment approaches.

Process of Science

Continuous model development and revision is a fundamental aspect of JGCRI integrated assessment research. Within the institute, this development is highly collaborative between team members who specialize in particular areas. Collaboration between JGCRI researchers and investigators located at PNNL in Richland, Washington is increasing, and data analysis and visualization laboratory sites are being equipped and configured to facilitate this interaction. Additionally, JGCRI is involved in joint projects with the Oak Ridge National Laboratory, Lawrence Berkeley National Laboratory, Argonne National Laboratory, and a number of universities. These joint projects involve data exchanges and joint model development and analysis.

JGCRI investigators maintain up-to-date data sets for a number of factors represented within integrated assessment models and to describe drivers of climate change. Additionally, climate simulations derived using Earth system models are accessed frequently. The institute makes results of analyses and assessments available through its Web site but also transfers results to other data centers for distribution within the climate science community and to the general public.

Local Science Drivers—the next 2-5 years

An important focus during the next 2–5 years will be the completion and application of a fully-coupled integrated assessment and Earth system model, built upon versions of the GCAM model and the CCSM climate or Earth system model. Full solutions to this coupled model will be derived using supercomputers located at ORNL, but setup of the integrated assessment component of the model will occur from JGCRI and investigators at JGCRI will be working with their colleagues at other national laboratories to interpret model results.

During this period, the 5th assessment of the IPCC will be completed, and JGCRI investigators expect major roles as authors and contributors. This work will require access to large datasets of climate model simulations as well as underlying data on model assumptions and inputs.

Utilizing high-performance computing resources being acquired now, JGCRI will undertake a variety of modeling studies to investigate options for climate change mitigation and adaptation. Of particular interest will be global simulations of agricultural and forest systems to investigate the potential for purpose-grown biomass to meet energy demands and concurrent impacts on carbon storage in the world's forests due to land use competition under a changing climate. The land cover and land use components of the GCAM model are being enhanced to support such analyses with more geographic detail in the underlying data and model functions. The use of satellite remote sensing data within this component of GCAM is increasing.

The attributes of climate change vary substantially over the Earth's surface, and policy makers and natural resource managers require regional data and information at higher spatial resolutions than are typically produced by global analyses. PNNL and collaborating institutions are developing extended capabilities for regional assessments of climate change, particularly for parts of the U.S. This activity involves close

collaboration between JGCRI and PNNL investigators in Richland, Washington. The computational requirements of these regional activities parallel those of global-scale studies, but typically with higher spatial and temporal resolutions and correspondingly higher storage and computational demands where large regions are considered.

Instruments and Facilities

During the next 2–5 years, JGCRI expects to make use of the high-end computing system for integrated assessment, located in College Park, Maryland, as well as supercomputers at ORNL in a closely coupled manner to enable coupling of the GCAM integrated assessment model and the CCSM climate system model. Coupling of these models will be complicated by the requirements to merge their underlying data. Data required by the integrated assessment component of the coupled model will be diverse and derived from a variety of sources outside those commonly accessed for climate modeling. Thus, we expect to assemble and process data on JGCRI computing systems and transfer data sets as required to ORNL to be incorporated into specific solutions of the fully-coupled model.

Joint work by JGCRI staff and PNNL investigators located in Richland, Washington to extend capabilities for regional analysis and assessment of climate change will require collaborative work to develop models and interpret results. Large data transfers between the two sites are expected, and staff members located in Richland, Washington will make use of high-end computing systems for integrated assessment located in College Park, Maryland.

Process of Science

Global climate and Earth system models are computationally demanding. Solutions for the cases considered, for example, by the IPCC assessments challenge the nation's supercomputer resources. While integrated assessment models are less demanding computationally, the underlying data are, in many respects, more diverse, and data structures underlying these models are complex. Coupling integrated assessment and Earth system models requires research and development to resolve inconsistencies in scales and other attributes of variables shared by both classes of models and the formulation of a programming and data management framework for implementing solutions to such inconsistencies within a fully-coupled model.

Research to couple integrated assessment and Earth system models is related in many respects to work to enable regional assessments of climate change. Both activities require coupling model components from the climate and integrated assessment domains and resolution of scale inconsistencies.

Coupling of integrated assessment and Earth system models is progressing in a step-wise fashion. Initially, data sets derived from one model are used as input to the other. As coupling issues are resolved, selected variables are coupled directly. Thus, the development effort requires substantial data transfers by its nature, and collaboration between the national laboratories involved in this activity adds to networking requirements.

Remote Science Drivers – the next 2-5 years

Requirements for increasingly robust and accurate information about climate change with regional specificity demand corresponding increases in computational resources. Capabilities for solving fully-coupled Earth system and integrated assessment models will reside at only a few locations. JGCRI expects to make extensive use of supercomputers located at ORNL to continue development of coupled models and to undertake applications. In addition to this trend requiring supercomputer class resources, the demand for regional-scale information will come from widely distributed stakeholders. JGCRI expects to work closely with PNNL colleagues in Richland, Washington to extend capabilities for regional analyses and to develop methods for making regional-scale information widely available.

JGCRI investigators expect to be involved in various aspects of the 5th IPCC Assessment. This international assessment of climate change requires exchanges of data with both U.S. and foreign institutions and coordinated model and data analyses. Additionally, JGCRI is involved in planning for a U.S. national assessment of climate change to be executed over the next several years. This assessment activity will require close collaboration between participating U.S. institutions, and JGCRI expects to provide data management and modeling support to the national effort, particularly in terms of energy-economics and ecosystem impacts.

Instruments and Facilities

During the next two to five years, climate change integrated assessments will rely increasingly on fully-coupled integrated assessment and Earth system models. JGCRI is participating in the development of a system coupling the GCAM model and the NCAR Earth system model. Supercomputers located at ORNL will provide the primary resources for solution of these coupled models. The high-end computing system located at JGCRI will be used to solve and evaluate the integrated assessment component of the model. Diverse data sets required by integrated assessment models will be assembled on the mass-storage system at JGCRI. Subsets required for particular solutions to the fully-coupled model will be extracted at JGCRI and transferred to computing systems at ORNL.

PNNL investigators in Richland, Washington and at JGCRI in College Park, Maryland are collaborating to extend capabilities for integrated assessment of climate change at regional scales. This activity requires daily sharing of data, model results, and simultaneous visualization of both. The work will make heavy use of the new high-end computing system for integrated assessment at JGCRI.

Process of Science

The major activities of integrated assessment research involve collaboration between national laboratories and other institutions. Particular institutions have been involved in the development of various model components or the assembly of data sets, and these institutions continue to focus on these as models are coupled and data are transformed or scaled to allow model coupling or application of models at higher resolutions and regional domains.

In the case of coupled integrated assessment and Earth system models, a fully-coupled model will be solved at a supercomputer center (ORNL), and the multi-institutional team will increasingly focus on the codes and data installed at that center. But interpretation and diagnosis of this complex model will occur with specialization at each participating institution.

Beyond 5 years – future needs and scientific direction

As climate change continues, DOE programs to address this challenge, as well as those of other agencies and institutions, will almost certainly expand. Depending on trends in policies and their success, emphasis may well shift steadily from mitigation and adaptation, which are paramount now, into more effort to address impacts and vulnerability. As humans contend with climate change, observations will likely play a larger role in integrated assessment, particularly in terms of diagnosing the details of change as it occurs and characterizing its impacts on natural and human systems.

Unless climate change mitigation is unexpectedly successful, significant changes within the Earth system and human enterprises will occur, and a variety of measurement platforms and networks likely will be deployed to track and understand these changes. Full interpretation of this expanding volume of Earth observations will require assimilating measurements into integrated assessment and Earth system models to diagnose ongoing changes and to predict their further evolution.

Integrated assessment modeling and analysis groups, including JGCRI and its direct collaborators, can expect to contend with far larger volumes of data and to assimilate data increasingly near to real time. Because of the global nature of climate change, data collected by space-borne remote sensing instruments will become more relevant and better suited to addressing steadily evolving climate change issues and questions. As attention to climate change impacts increases, JGCRI will likely be assimilating satellite remote sensing data and systematic measurements by surface networks within days of their collection. Daily volumes of data transfers will grow to terabytes, and these will be matched by even large volumes of model outputs.

Integrated assessment and Earth system models are so intimately linked that these modeling systems will be fully coupled over the next five years. The ongoing process of coupling such models will accelerate if critical feedbacks within the Earth system and the human factors and activities that drive climate change become more evident or significant. Such fully coupled models, along with increased capabilities to analyze climate change impacts, will be critical to assessing and mitigating particularly dangerous positive feedbacks of these coupled systems.

These directions in integrated assessment and climate change research will require substantial increases in network capacity and speed. The need to couple global observations more closely to models, somewhat as observations are now incorporated into weather forecast systems, will require more flexibility in terms of networking instrument networks globally and improved connectivity between major high-performance networks to permit access to data collected by all relevant institutions.

Middleware Tools and Services

Most data used in the integrated assessment of climate change and most model variables have a geospatial context. Where such data are used in combination, they must be transformed into common coordinates and georeferenced to a standard. Evaluation and interpretation of geospatial data usually requires mapping, often in conjunction with other standard geographic information such as geopolitical boundaries, water features, and terrain.

Management and analysis of geospatial data involving numerous variables requires a geographic information system (GIS), and collaborative studies require exchanges of data between such systems, often requiring reformatting.

Summary Table

Feature	Key Science Drivers		Anticipated Network Requirements	
	Science Instruments and Facilities	Process of Science	Local Area Network Bandwidth and Services	Wide Area Network Bandwidth and Services
Near-term (0-2 years)	<ul style="list-style-type: none"> • GACM integrated assessment model; • CCSM Earth system model; • Earth science data centers; • Satellite remote sensing data; • Surface climate observations; • High-end computing and storage system. 	<ul style="list-style-type: none"> • GCAM model development and revision • EPIC model applications to carbon management • Coupling of GCAM to CCSM to form iESM • U.S. National assessment of climate change • IPCC 5th assessment 	<ul style="list-style-type: none"> • 1 Gbps 	<ul style="list-style-type: none"> • 1 Gbps
2-5 years	<ul style="list-style-type: none"> • Coupled Earth system and integrated assessment models (iESM); • Regional analyses linked to global models and data. 	<ul style="list-style-type: none"> • Inter-comparison of integrated assessment models; • U.S. National assessment of climate change. • IPCC 5th assessment. 	<ul style="list-style-type: none"> • 10 Gbps 	<ul style="list-style-type: none"> • 10 Gbps
5+ years	<ul style="list-style-type: none"> • Data assimilation; • Access to global observation networks • Climate change impacts 	<ul style="list-style-type: none"> • Continuous assessment of climate change impacts and the effectiveness of mitigation and adaptation at global and regional scales. 	<ul style="list-style-type: none"> • 10 Gbps 	<ul style="list-style-type: none"> • > 10 Gbps

6 Climate Science—Enabling Worldwide Access to Petascale Climate Data

Background

Insightful analysis in climate science depends critically on the choices of hardware, software, and network tools to discover, access, manipulate, and visualize the often large and unwieldy data sets of interest. Requiring advanced networking for large-scale access and movement of data, these data exploration tasks can be complex and time-consuming, and they frequently involve many resources spread throughout the modeling and observational climate communities. As climate science moves beyond the physical climate to include detailed carbon cycle and other biogeochemical feedback processes within new Earth System Models (ESMs), the increased complexity generates much more output from simulations than ever before. Aggregations of these processes are used to specify and predict future climate change for time-evolving interactions between the atmosphere, ocean, land surface, sea-ice, and the biosphere. This complexity requires the expertise of a wider group of climate scientists, analysts, educators, private industry, and government officials (both domestic and foreign) to understand the risk of human-induced climate change, its observed and projected impacts, and options for adaptation and mitigation. To these ends, integrating distributed data and computers, high-bandwidth wide-area networks, and remote computing using climate data analysis tools in a highly collaborative problem-solving environment is essential to predict the consequences of climate change due to fossil fuel use and other human activities.

In one of the largest-ever worldwide collaborative efforts in computational science, climate scientists participated in the Intergovernmental Panel on Climate Change (IPCC) Fourth Assessment Report (AR4). The report, released in 2007, included work contributing to the completion and analysis of hundreds of coupled climate simulations. The simulations produced hundreds of terabytes of data spread around the globe. Charged with facilitating access to a unified database of IPCC model output known as the Coupled Model Intercomparison Project, Phase 3 (CMIP-3), the Program for Climate Model Diagnosis and Intercomparison (PCMDI) at Lawrence Livermore National Laboratory (LLNL) used the Earth System Grid (ESG) to house, manage, search, and disseminate data critical to the success of the IPCC AR4.




Figure 1: Simulation of CMIP-3 historical and projected climate data shows a strong warming trend in the lower atmosphere underneath a cooling layer in the upper atmosphere. (Image courtesy of Peer-Timo Bremer, LLNL and Valerio Pascucci, University of Utah.)

The Earth System Grid (ESG) is designed to address the management and use of extremely large and diverse data sets. Begun with DOE support in 1999, ESG has since become a Scientific Discovery through Advanced Computing (SciDAC) Center for Enabling Technologies (CET) project. The ESG-CET is building “science gateways” to these critical climate resources. Open to everyone, the ESG houses and provides access to and management of climate data, information, models, analysis tools, and visualization tools. ESG is built on the concept of grid computing, an approach that pools hardware and software assets of many distinct networks and locations into a federated virtual system and provides single access to distributed resources. The resulting virtual system is more powerful and has much greater capacity than any of the individual systems or networks that comprise it.

The goals of the ESG project are to: (1) make data more useful to climate researchers by developing grid technology that enhances data usability; (2) meet the specific needs of national and international climate projects for distributed databases, data access, and data movement; (3) provide a universal and secure web-based data access portal for broad-based multi-model data collections; and (4) provide a wide range of grid-enabled climate data analysis tools and diagnostic methods. In short, ESG-CET works to integrate distributed data and computers, high-bandwidth wide-area networks, and remote computing using climate tools in a highly collaborative problem-solving environment.

In 2004, the CMIP-3 formed a comprehensive output of both oceanic and atmospheric data on a monthly, daily, and 3-hourly basis—holding over 35 terabytes (TB, where 1 TB is 10^{12} bytes) of data at the LLNL central site. Today, the Coupled Model Intercomparison Project, Phase 5 (CMIP-5) is assembling the most-comprehensive archive yet. It is projected to be 10 petabytes (PB, where 1 PB is 10^{15} bytes) and contain many model simulations from more than 30 national and international modeling centers to form the basis for the next periodic assessment report of IPCC (i.e., the Fifth Assessment Report (AR5)). The multi-model data sets and corresponding observational sets will not reside at any one site. Instead, these resources will be distributed at many sites spanning six continents, requiring that researchers have remote access to the hardware, networks, data, and tools for ultra-scale data management and analysis.

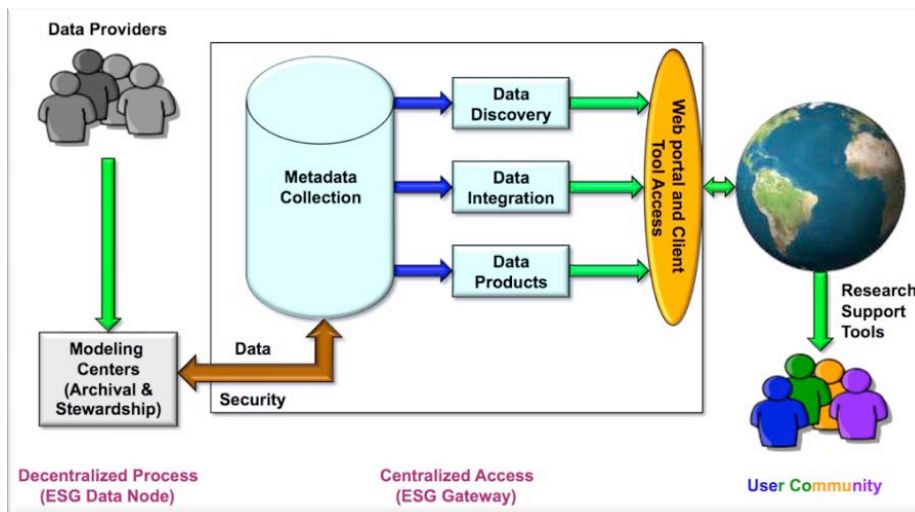


Figure 2: ESG data providers are distributed globally, publishing their data to metadata collection centers, at which users search and discover and get returned data products from a centralized access point (i.e., web portal or client analysis tool).




Figure 3: ESG makes data a community resource, accessible worldwide. This image shows the institutions that accessed ESG data from 2004 to present day. Delivering over 1 PB of data to 2,700 sites in 120 countries, some 20,000 users have registered to access ESG data holdings.

As these past trends indicate, climate data presents a real challenge for data management, analysis, and networks—outpacing storage capacity and network needs for the foreseeable future. In fact, the 2009 ASCR Network Workshop¹ and the 2007 BES Network Workshop² projected that by 2020, climate data will exceed hundreds of exabytes (1 XB, where 1 XB is 10^{18} bytes)—a staggering growth rate. Because of this data explosion and the limitations of networks and computing resources, climate and computational scientists have been compelled to analyze and visualize large-scale data close to where it is formed, moving data only when absolutely necessary.

¹ <http://www.er.doe.gov/ascr/ProgramDocuments/Docs/ASCR-Net-Req-Workshop-2009-Final-Report.pdf>.

² <http://www.es.net/pub/esnet-doc/BES-Net-Req-Workshop-2007-Final-Report.pdf>.

As with CMIP-3, PCMDI at LLNL is charged with managing and storing petabytes of CMIP-5 observational and climate model output for scientists contributing to AR5. The upcoming CMIP-5 will use new ESMs to perform more experiments and create far more data than for the AR4. As an example, the community activity that involved several DOE laboratories and the National Center for Atmospheric Research (NCAR)—as equal partners—participated in the Community Climate System Model version 3 (CCSM3) which provided roughly 9 TB of the total 35 TB model output to the CMIP-3 archive, located at LLNL. For the forthcoming CMIP-5 archive, which will be fully populated in 2011, CCSM version 4 (CCSM4) and the Community Earth System Model version 1 (CESM1) will submit roughly 300 TB of output out of the 1 PB of data generated. (Note, CCSM data will reside at ORNL and LBNL as well as NCAR—particularly the high resolution runs that will be executed by a four-laboratory partnership led by ORNL—that are part of the CCSM CMIP-5 contribution. Those data will be housed at ORNL.) All though it is not required for the AR5, the additional data will also be of interest to the science community. The two-dozen other major modeling groups (e.g., Japan, U.K., Germany, China, Australia, Canada and elsewhere) will create similar volumes of data with merely a fraction of the data migrating to LLNL to form the CMIP-5 *replica centralized archive* (RCA) (see **Table 2**). To accommodate this important multi-PB data archive, LLNL is building a new data repository for anticipated heavy use in the next 1–2 years. Serving this new system with multiple 10 Gbps network connections organized in a link aggregation group (LAG) would be operationally complex and challenging. For effective access to the climate data by a large number of remote users (in the 10s of thousands), LLNL and ORNL will require 100 Gbps network access before it becomes commercially available.

At present, Lawrence Berkeley National Laboratory (LBNL) and Lawrence Livermore National Laboratory (LLNL) are currently involved in developing one of the Advanced Network Initiative (ANI) research projects, called Climate 100. Climate 100 is a research and development project to test end-to-end capabilities of the DOE Advanced Network Initiative (ANI) testbed in collaboration with the ESG community. At this time, LLNL is not directly affiliated with ANI. However, a proposal is being evaluated to include LLNL as part of the ANI 100 Gbps backbone network connection—a connection that will be needed for petascale climate data movement and accessibility in the future.

As stated earlier, ESG is building a “science gateway” that enables a researcher-friendly infrastructure for large-scale climate resources. This next generation problem-solving environment will allow access to distributed federated data, information, models, analysis, visualization tools, and computational resources. It is based on the Internet and the need to access and analyze large-scale data anywhere in the world. There will be a persistent need to move vast amounts of data to and from sites for scientific purposes. For example, large-scale data that reside at various sites around the world will need to be assembled at one location for ensembles or multiple model data intercomparisons. Initially, LLNL, NCAR, and ORNL will provide ESG-CET Gateways. Each Gateway will act as a broker, handling data requests that serve specific user communities. Services deployed on a Gateway include a user interface for searching and browsing metadata, requesting data (including analysis and visualization) products, and orchestrating complex workflows.




Figure 4: The y-axis represents the cumulative download volume in gigabytes (GB) and the x-axis represents yearly time values. Overall, the demand for the ESG data holdings has grown exponentially over the last five years. ESG main portals are recording data download volumes in the tens of terabytes per month, resulting in a cumulative download total of just over 1 PB.




Figure 5: Illustration of two ESG community Gateway portals serving CMIP-3 and observational climate data from the PCMDI portal and CCSM3 and NARCCAP data for the NCAR portal.

Under the leadership of PCMDI at LLNL and with the help of NCAR, ORNL and others in the national and international community—including U.K., Germany, Japan, and Australia to name a few—the data management needs of CMIP-5 will be met by an internationally federated, distributed data archival and retrieval system—referred to as the *Earth System Grid Federation* (ESGF). This federation will exploit software primarily developed by the ESG-CET consortium in the U.S. composed of Argonne National Laboratory (ANL), Los Alamos National Laboratory (LANL), Lawrence Berkeley National Laboratory (LBNL), Lawrence Livermore National Laboratory (LLNL), National Center for Atmospheric Research (NCAR), Oak Ridge National Laboratory (ORNL), Pacific Marine Environmental Laboratory (PMEL), and the University of Southern California (USC) Information Sciences Institute.




Figure 6: The ESG Federation included both the Data Nodes and Gateways, with the Gateways located at CMIP-5 Data Centers. The PCMDI Data Center will archive the “replicated data” and will serve data from modeling centers choosing not to host a Data Node.

In addition to hosting the CMIP-5 archive, ESG will be archiving many other important data sets for access and use by the international community, including data produced from: the Community Climate System Model version 4 (CCSM4), the North American Regional Climate Change Assessment Program (NARCCAP), the NASA Earth Observation Satellites, the Atmospheric Radiation Measurement (ARM) program, the Carbon Dioxide Information Analysis Center (CDIAC), the Carbon-Land Model Intercomparison Project (C-LAMP), and the Advanced Very High-resolution climate model simulation data resulting from Earth System Modeling (ESM) and Regional and Global Modeling (RGM) programs.

Key Local Science Drivers

Instruments and Facilities:

LLNL’s planned archive is based on the scalable PB storage system called Storage Scalable Units (SSU). These units consist of very reliable commodity servers and HPC disk drawers, available at costs lower than conventional redundant array of independent disk (RAID) systems. Working with the ESG-CET enterprise system, virtualized servers would run both the ESG-CET Data Node software stack and CloudFS back-end storage. (ESG-CET Data Nodes reside with the data and publish to the ESG-CET system for client retrieval. Each SSU contains 1.2 PB of raw disk space with two-copy redundancy capability. Other important features include embedded advance compression (>2x), 4-GB/sec random read bandwidth, and 1 Gbps and 10 Gbps storage area network (SAN) infrastructure. The hardware architecture of the SSU consists of 128 ports to 10 Gbps Ethernet switches and 44 external SANs connected to multiple 10 Gbps routers (*Figure 7*).

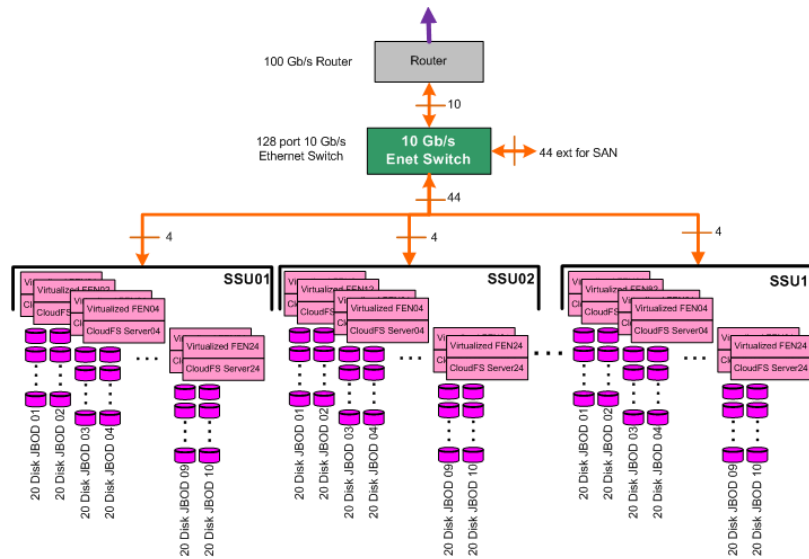


Figure 7: Depiction of the system architecture for the PCMDI using multiple SSU. Commodity servers and HPC disk drawers provide maximum density and the lowest cost. Virtualized servers run both ESG front-end and CloudFS back-end.

Funding from DOE’s Office of Biological and Environmental Research (BER) and LLNL is being used to purchase five SSUs containing at total of 10 PB of raw disks with 20 GB/sec of random Input/Output (I/O) bandwidth—a system capable of easily supporting over 20,000 users worldwide.

Process of Science:

The climate change community faces the difficult challenge of managing the distribution of massive data sets to thousands of collaborators worldwide. LLNL’s PCMDI already hosts the CMIP-3 Multi-Model archive comprising of 35 TB of data from 25 global coupled ocean-atmosphere models that were run for the IPCC AR4. For the forthcoming CMIP-5 Multi-Model archive, the RCA (housed at LLNL) is expected to contain between 1 to 2 PB of data with the remaining distributed archive reaching upwards of 10 PB. This data is expected to begin its online appearance in 2010 and will reach its full archival size by 2011. To accommodate and support such a large and diverse community, the RCA will be replicated at four other sites with LLNL as the hub (shown in **Figure 6**).

To move data replicas efficiently, CMIP-5 must be capable of utilizing advanced networks such as the ANI 100G prototype. **Table 1** illustrates the impact on network performance when moving PB of data instead of TB of data. The fastest we could hope to move only 1 PB of data from PCMDI to one of the RCA data centers is essentially one day at 100 Gbps, whereas with a peak of 10 Gbps, it would take almost 1.5 weeks.

Transfer Rate	Time to Transport 1 TB of Data	Time to Transport 1 PB of Data
10 Mbps	9.7days	27.20 years
50 Mbps	1.94 days	5.44 years
100 Mbps	23.3 hours	2.72 years
1 Gbps	2.28 hours	97.1 days
10 Gbps	13.65 minutes	9.7 days
100 Gbps	81.9 seconds	23.3 hours

Table 1: Time displayed in International System of Units to transport 1 TB and 1 PB of data from one site to another using various transmission rates.

With the multiple CMIP-5 archives expected to begin operations by mid-2010, transport vehicles to move data from PCMDI to the four RCA sites should already be in place to help move the vast amounts of data. Since the networks at this time are not equipped to move such data volumes, the PCMDI team has elected to move data via physical TB disks to bootstrap the ESG distributed enterprise system. Once the initial PB distributed archives are in place and access to faster networks becomes available, it is expected that the transport mechanism of choice will be fast 100 Gbps networks to move data for ESG data versioning control and data replication.

The usual procedure for transporting TB disks are described on the CMIP website (i.e., <http://cmip.llnl.gov/cmip5/submit.html?submenuheader=2>)—relying heavily on multi-terabyte portable hard disks sent to PCMDI via a parcel shipping service (e.g., FedEx). Using reliable 2 TB disks, it will take roughly 500 FedEx shipments to populate the Data Center at PCMDI (i.e., *one way shipment*). To populate the other CMIP-5 Data Centers (i.e., BADC, Germany, Australia, and Japan) it will take a total of 2,000 additional FedEx shipments to meet the archival requirements of the ESG Federation (shown in **Figure 6**). If you include shipment of the initial blank disks to the modeling centers (500 shipments), the data shipment back to PCMDI (500 shipments), the shipment of disks to the other four CMIP-5 Data Centers (20,000 shipments), and the final return of the disks to PCMDI (500 shipments), the grand total of parcel shipments comes to 35,000!

As climate research is a global endeavor, there is not only a strong need for large trans-Atlantic bandwidth as described in the case above, but there is also a need for large trans-Pacific bandwidth to reach researchers in Asia and Australia (see **Figure 3**). (Note: in general it is more difficult and costly to get connectivity between continents and the circuits tend to be more failure-prone with longer outages). To match globally distributed PB archives and data demands, 100 Gbps trans-oceanic bandwidth is needed. **Table 2** shows the trans-oceanic connectivity need by showing some of the countries and their projected archive data volume for CMIP-5.

Primary Group	Country	Estimated Total Data Volume (TB)	Requested Output Data Volume (TB)	Replicated Output Data Volume (TB)
NorClim	Norway	500	30	16
Hadley Centre	U.K.	1000	89	47
GFDL	U.S.A.	1000	151	62
IPSL	France	300	45	20
NIES & U. Tokyo & JAMSTEC	Japan	500	115	33
CCCMA	Canada	300	29	13
MPI	Germany	1000	710	269
INGV	Italy	300	50 (?)	20 (?)
EC-Earth Consortium	Europe	300	50 (?)	20 (?)
CSIRO	Australia	500	20	13
NASA GSFC	U.S.A.	300	50 (?)	20 (?)
CAWCR	Australia	300	21	8
CCSM	U.S.A.	1000	309	107
MRI	Japan	500	312	75
METRI (with Hadley Centre)	Korea	300	13	6
LASG IAP	China	500	63	19
NASA GISS	U.S.A.	300	50 (?)	20 (?)
BCC	China	500	42	20
INM	Russia	300	32	14
CNRM & CERFACS	France	400	64	25
U. Reading (with Hadley Centre)	U.K.	400	63	16
Combined Group Total	CMIP-5 Simulation Data	10,500	2,308	843

Table 2: Shows the participating international CMIP-5 groups, their home countries, and their projected total data volumes. Not all data will be replicated in the PCMDI CMIP-5 RCA, but the majority of the data will be accessible to users from the ESG federated Gateways as shown in **Figure 6**.

Key Remote Science Drivers

Instruments and Facilities

The hardware, storage, network, and software requirements for the ESG Gateway and Data Node are:

- Modern Hardware running the latest version of CentOS or Red Hat Enterprise Linux (RHEL). For example:
 - Dell PowerEdge 2850: 2 core x 3.2 GHz Intel, 4 GB RAM, RHEL 5.3, 26 TB disk.
 - Dell R610 2-CPU E5540, 32 GB RAM, CentOS 5.3, 24 TB disk, direct attachment to 250 TB RAID archive.
 - Other flavors of Linux/Unix will probably work (if they support TCP autotuning and modern congestion control algorithms such as cubic), but will not be supported in 2010 and 2011—as time progresses, additional operating systems will likely be added to the support list.
- For CMIP-5, spinning disks (rather than tapes) are required for storage to ensure fast access. The choice of specific hardware devices is up to the host. For less urgent data, tape archives such as High Performance Storage System (HPSS) are adequate storage facilities.
- The faster the network the better: Some groups will have 2 x 10 Gbps network access while others will have less than 1-Gbps. Initially, PCMDI will have 2 x 10 Gbps network access and is currently working to be ready to connect to future ESnet BAMAN at speeds up to 100 Gbps.
- ESG Gateways can be hosted on a single machine, or decomposed into two servers—one hosting the database back-end, and the other hosting the web applications. The software stack includes:
 - Postgres database, version 8.3+, Java 1.6 or above, Tomcat 6.x.
 - ESG Gateway web application (war file).
 - MyProxy (latest version from Globus toolkit).
 - The Gateway application, including the database, should have access to approximately 100GB of disk space.
- ESG Data Nodes typically will have between 100 to 1,000 TB of disk storage and multiple systems with 10 Gbps network connections. Nodes are typically Linux front-end systems running a minimal software stack consisting of the following:
 - *CDAT/CDMS* – CDAT is based on the Python object-oriented language. CDMS is a sub-module of CDAT used by ESG to scan data.
 - *Python ESG package* – publication modules, scripts, and GUI.
 - *Apache Tomcat (Java servlet container)* – web application server that runs the THREDDS application.

- *THREDDS Data Server (TDS)* – a servlet for cataloging ESG-CET Data Node files and data resources used for reading catalogs and publishing to the PCMDI Gateway.
- *Live Access Server (LAS)* – a servlet for generating visualization products. (Note: LAS product service is optional and not yet ready for distribution).
- *Globus Toolkit/MyProxy* – in our distributed system it is important to provide data security and integrity as well as user authentication. The MyProxy client is used for Single Sign-On to the PCMDI Gateway.
- *GridFTP server* – GRID-enabled FTP server.
- *Test publication* – publish a test data file.

Scenario 1: Generation of Data and Replication and Movement

Data generation sites, with computers specialized for running models: LLNL, NCAR, ORNL, Japan, UK, Germany, Australia, and Canada. Authorized user logs onto LLNL's Gateway and issues a request to collect large-scale data from multiple source nodes (i.e., LLNL, NCAR, ORNL, Japan, UK, Germany, Australia, and Canada) in order to generate a temperature ensemble of the global models. The target node initiates the data transfer by pulling data from the multiple source nodes. A requestID is returned to the user, and transfer starts asynchronously. The user can check status of transfer requests with requestID. Size: approximately 1 PB, depending on the data set.

Scenario 2: Mirroring a Core Dataset

Core dataset needs to be mirrored from source node to target nodes. For example, BADC node needs to mirror all or part of core dataset from LLNL node. A requestID is returned and transfer start asynchronously at a target node by pulling data from the source node. The user can check status of transfer request using requested. Size: approximately 1.2 PB at the maximum depending on what target node requests to mirror. This process could take slightly over 1 day using a 100 Gbps network connection at LLNL.

Scenario 3: Many Users Access Many Data Subsets

Thousands of users log onto the LLNL Gateway to search and browse data. They simultaneously request data subsets consisting of hundreds of thousands of files located on the LLNL, NCAR, and ORNL data nodes. Transferable URLs are returned to each user and users start transfers concurrently. The given size that any one user can access and download is approximately 10 TB. The system manages the I/O requests in parallel and balances loads on transfer servers.

Process of Science

When driven by a 100 Gbps network and extreme-scale computing and the data it generates, scientists' productivity will increase, enhancing their national and international relationships and push their research to new levels of understanding (*Figure 8*). Although the data transfer protocol developed in the ESG project is targeted for access in

middleware and applications, we also want to ensure that end-users—from politicians and decision makers to health officials and students as well as research scientists—can access much of the data that are resulting from our focus on the data transport protocol to deliver large data sets between DOE laboratories over 100 Gbps network connections. In addition, at the two Magellan systems at ALCF and NERSC, we will provide researchers at remote locations with the capability to explore climate comparative analyses with data transport methods over 100 Gbps network connection, allowing them to run realistic experiments at the scale required for climate modeling and intercomparison.




Figure 8: This schematic depicts a use-case scenario supporting all the capabilities and complexities needed for rich data exploration and manipulation, as well as a fault-tolerant, end-to-end system for moving and integrating extreme-scale data sets. The yellow-shared area is where advanced networks can be used to deliver data sets and analysis results over 100 Gbps networks.

An example scenario illustrates how advanced networks might provide a critical method for users. Several outbreaks of malaria have occurred in different regions of central Africa, and health officials believe that local climate conditions are affecting the spread of the disease. Domain experts (i.e., climate scientists) in different international locations are asked to provide crucial inputs for policy decisions to control the outbreaks. From their remote sites, the scientists search an ESG climate portal containing exabytes of high-resolution regional data from central Africa. When their data search proves unproductive, the scientists run several models in real time to generate ensemble simulations of the African climate. Using advanced networks and the cloud-computing resources, they are able to simultaneously transport and calculate results of ensemble climate statistics. The climate scientists then save the analysis results in the ESG data-management system. A policy maker in another region is working to prevent malaria outbreaks and in studying possible prevention tactics discovers this saved session. Working with the scientists, the policy maker conducts further assessment and re-analysis of the derived data sets. In the process, the new results are reduced from 20 petabytes to 100 terabytes and moved to a local storage area for further study. The scientists then

integrate the climate model ensemble data with the African malaria data for potential future use.

Local Science Drivers—the next 2–5 years

Instruments and Facilities

The goal of the ESG project is to research, develop, and deliver end-to-end capabilities in collaboration with the climate community. This project will need to address scalability issues of the software, services, and applications associated with massive climate data sets over the anticipated 100 Gbps infrastructure (**Figure 9**). ESG plans to deploy similar systems on multiple 10 or 100 Gbps connections to move massive data sets between its major Data Node sites in the U.S. With demonstrated success at the U.S. sites, ESG will then extend high-performance data transfers to smaller Data Nodes in the U.S. as well as to foreign partner sites. This scaled testing of the data transfers and communication connections is important because climate researchers worldwide use many different network service providers. Continued success will enable a production-scale system that will empower scientists to try new and exciting data exchanges that could lead to breakthrough discoveries.




Figure 9: The envisioned federated topology of the ESG-CET enterprise system utilizing 100 Gbps network connections. A network of geographically distributed Gateways and Data Nodes is built into a globally federated “built-to-share” scientific discovery infrastructure. By federating these Gateways using a fast network, independent data warehouses deliver seamless access to vast data archives for scientists and their specialized client applications. Experts (e.g., model developers, climate researchers) and non-experts alike need fault-tolerant end-to-end system integration and large data movement, and benefit from rich data exploration and manipulation—in the process moving vast amounts of data to and from sites around the world.

Remote Science Drivers—the next 2–5 years

Instruments and Facilities

Unfortunately, access to this new system by potentially 20,000 global users could be constrained by limited local network bandwidth. Currently, LLNL is connected to ESnet's Bay Area Metropolitan Area Network (BAMAN), a shared 20-Gbps protected Bay Area network that is resilient to human and environmental threats. In a highly collaborative, decentralized problem-solving environment like ESG, a faster network—utilizing individual connections 10 to 100 times faster than what is available today—will be needed to make efficient use of the data and tools available to scientists. Climate researchers want the ability to quickly and seamlessly combine multiple data sets, up to 300 TB each, for analysis—one which is not feasible using today's 10 Gbps networks. The only way to ensure the ESG architecture scales to meet its users' requirements is to utilize next generation network speeds of 100 Gbps. As a result, the ESG project is proposing to leverage a new 100 Gbps prototype network to make high-speed data federation a reality.

(Note: ORNL is connected to ESnet via a 10 Gbps layer-3 and 10 Gbps layer-2 (Science Data Network) through Nashville.)

Process of Science

Same as above.

Beyond 5 years—future needs and scientific direction

See *Summary Table* below.

Outstanding Issues

LLNL would like to become one of the primary sites on the experimental ANI 100 Gbps prototype network in order to begin high-speed data transfer testing with ORNL. This request means the ANI project would need to include a 100 Gbps optical wave service from the network backbone node in Sunnyvale to LLNL.

Summary Table

Feature	Key Science Drivers		Anticipated Network Requirements	
	Time Frame	Science Instruments and Facilities	Process of Science	Local Area Network Bandwidth and Services
Near-term (0-2 years)	<ul style="list-style-type: none"> ESG gateways at LLNL, NCAR, and ORNL with scalable front ends—multiple 10 Gbps ESG nodes at MPIM/DKRZ, BADC, ANU, Japan, and 45 other sites worldwide 	<ul style="list-style-type: none"> CMIP-5, CCSM3, NARCCAP, POP, C-LAMP data access 	<ul style="list-style-type: none"> Expand 10 GigE domestically 1 GigE internationally 	<ul style="list-style-type: none"> 100 Gbps network connections 10 Gbps network internationally
2-5 years	<ul style="list-style-type: none"> ESG gateways at LLNL, NCAR, and ORNL with 100 Gbps connection Extend ESG gateways to NOAA, NASA, other national and international agencies serving observational data, at 10 to 100 Gbps 	<ul style="list-style-type: none"> Continuing CMIP-5 and CCSM3 data access CCSM4/CESM1.0 data access ARM, CDIAC, ArmeriFlux NASA satellite observations Very high-resolution regional and global model output 	<ul style="list-style-type: none"> Expand 100 GigE domestically 10 GigE internationally 	<ul style="list-style-type: none"> 10x100 Gbps network in U.S. 100 Gbps network internationally

5+ years	<ul style="list-style-type: none"> • All U.S. ESG gateways and nodes at 10x100 Gbps connections • International nodes to 100 Gbps 	<ul style="list-style-type: none"> • Continuing CMIP-5 and CCSM3 data access • New versions of CESM data access • ARM, CDIAC, ArmeriFlux • NASA satellite observations • Very high-resolution regional and global model output • CMIP-6 and CCSM4 and observational data access 	<ul style="list-style-type: none"> • Expand 10x100 GigE domestically • 100 GigE internationally 	<ul style="list-style-type: none"> • 1-Tbps network in U.S. • 10x100 Gbps network internationally
----------	---	---	---	---

7 Network Requirements of the DOE Systems Biology Knowledgebase

A Pacific Northwest National Laboratory Perspective

Background

Many systems biology researchers today do not engage directly with their research object, but do so via digitally captured, reduced, calibrated, analyzed, synthesized and, at times, visualized data. Advances in high-throughput experimental and computational technologies have led to an exponential growth in the volumes, variety and complexity of these data. For example, the submission rate at the European Molecular Biology Laboratory Nucleotide Sequence Database (EMBL-Bank – comparable to the US GenBank) is currently growing at a rate of 200% per annum, requiring a doubling in its storage capacity every year (5 PB in 2009). EMBL's web pages, which give access to 63 distinct databases, received ~ 3.5 million hits per day, and programmatic access to their data via web services is at ~ 1 million requests per month and growing (Southan, 2009¹). The number of databases reported in *Nucleic Acids Research* jumped from 218 to 1,170 in the years 2001-2009 (Goble, 2009²).

A central goal in systems biology is to use diverse research results from various experimental and computational methods, stemming from diverse and distributed sources, to study the complex interactions of a biological system in a holistic manner. This is a particularly fast-developing and evolving research field, where rapidly developing research methodologies are utilized for scientific discovery, and existing data sources and analysis tools are combined in ever changing ways to facilitate scientific progress. These research results will often be locally integrated and synthesized with the scientists' own findings, and published back into the community as yet another valuable source of information. Over the years, the community has created a wealth of outstanding data resources, as well as access, integration and analysis tools. The size of these data resources varies greatly from local databases with a few hundred megabytes of data to large-scale experimental data resources with potentially hundreds of terabytes of raw data output, such as that generated by 3rd generation synchrotron light sources. Data access for systems biology is still largely ad hoc, driven by individual scientific needs; however the usage of data mining algorithms and workflows has grown significantly in the past few years, and often involve the movement of larger volumes of data. At the higher end of the data volume range systems biology, like many other science domains, suffers from the inability to move data reliably between sites and applications, in particular for ad hoc transfers. Much data is either never fully accessed or distributed on hard drives, strongly limiting the community's ability to fully exploit its inherent scientific value.

Current data access request numbers are limited by the fact that many of the data resources available to the community are incompatible with each other, making discovery and integration a formidable challenge. In response, the DOE Office of Biological and Environmental Research is currently supporting the development of the Systems Biology Knowledgebase (<http://www.systemsbiologyknowledgebase.org/>), a focal point and cyberinfrastructure to enable a new level of scientific inquiry by facilitating the

integration of modeling, simulation, experimentation, and bioinformatic approaches through community agreed sets of shared terms, methodologies and technologies. The implementation of this vision over the next few years will make data resources more accessible to the scientific community. Furthermore, the provision of integrated application tools and scientific workflows coupled with cloud computing resources will further increase the communities demand for data access and the transfer of data between different resources. The nature of the knowledgebase implies that network transfers will be ad hoc, strongly vary in size and duration.

Key Local Science Drivers

Process of Science

The process by which systems biology researchers utilize facilities and resources such as those available at EMSL typically involves the researcher providing a biological sample (e.g., cells or tissues) for analysis, generation of the data at the facility (e.g., mass spectra for proteomics), analysis of the raw data, and return of processed data (e.g., peptide identifications) to the researcher. The biology community has so far been very good at reducing complexity and working with reduced problem sets to overcome challenges in data analysis, transfer and integration, therefore utilizing processes that obviate the need to transfer large volumes of data (processed data often being at least an order of magnitude smaller in size). While this has been a workable model in the early years of systems biology research, scientists are now requiring more frequently access to complete data sets to carry out more complex, specialized and sophisticated analyses of the research results. These data sets are not only composite results from different experimental and computational methods on one particular sample, but more importantly also include comparative analysis of the same type of results across a wide range of related species and species strains. Thus, transfer of large volumes of raw data is likely to become an issue in the near future.

PNNL Instruments and Facilities used by Systems Biology

The Environmental Molecular Sciences Laboratory (EMSL), a national scientific user facility at Pacific Northwest National Laboratory (PNNL), houses an unparalleled collection of state-of-the-art capabilities including over 100 major instrument systems and supercomputing capabilities. Researchers from around the world are encouraged to use EMSL's unique capabilities in combination with each other with an emphasis on merging computational and experimental instruments. Specifically for systems biology EMSL offers a wide range of capabilities for the investigation of biological interactions and dynamics. The EMSL is further described in its own case study in this report.

Key Remote Science Drivers

Process of Science

The process of knowledge discovery in systems biology demands access to diverse and distributed data sources and analysis tools. Typically, to study the complex interactions of

a biological system in a holistic manner, a researcher needs access to existing knowledge in the research literature (PubMed), access to archives of existing data related to the biology system under study, as well as a wide range of facilities for new data generation.

For example, large archives of data are available for:

- DNA sequence – GenBank, EMBL, DNA Data Bank of Japan (DDBJ)
- Gene expression – Gene Expression Omnibus (GEO), EST
- Protein structure and function – Protein Data Bank (PDB), the Universal Protein Resource (UniProt), Structural Classification of Proteins (SCOP), Pfam
- Taxonomy and phylogeny – the Ribosomal Database project (RDP), Greengenes, ARB
- Metabolism – Biocyc, KEGG
- Organism-specific data – e.g., EcoliWiki, FlyBase

And facilities for new data generation include:

- DNA sequencing (e.g., JGI) for genomics and transcriptomics data,
- Mass spectrometry (e.g., EMSL) for proteomics and metabolomics,
- Advanced spectroscopy and microscopy for biological sample visualization,
- Magnetic resonance for molecular structure studies,
- High-performance computing for data analysis and storage.

This leads to an inherent data integration problem. Because there is no centralized source, a researcher will spend considerable time searching for existing data and information from the multitude of distributed biology databases, then downloading the relevant data to local compute resources or moving them to remote resources where suitable capacity and applications are available. In addition, the researcher may have designed an experiment and utilized one or several of the types of facilities listed above, and these data are also typically downloaded to facilitate integration. Furthermore, data analysis and integration of these rich data sources often requires many different computational tools, which may also need to be downloaded for local use, or run remotely, requiring transfer of the data to the computer resources.

Instruments and Facilities

The systems biology community utilizes a wide selection of research facilities for its work, however their key drivers in terms of network transfer rates are Experimental facilities, High Performance Computational Resources and associated workflows e.g. for data mining.

Experimental facilities:

- There are a wide range of relevant medium scale facilities available in support of systems biology, with technologies becoming more widely/cheaply available the number of these facilities has strongly increased including those for genomics,

proteomics, microarray (transcriptomics) and metabolomics. There are currently well over 100 of these in the USA alone, but much more around the world would be of interest, in particular in fields where data coverage is sparse. Data sets typically are in the range of 10's to 100's of GB.

- In addition there are a range of large-scale experimental facilities of interest both nationally and internationally including neutron, synchrotron light sources and free electron laser facilities. These in particular can produce data sets in the 10's to 100's of TB range.

Computational simulations and analysis are widely used by the community, from the atomic and molecular level to full systems, fluid dynamics and environmental studies. Equally the analysis of experimental data requires increasingly high throughput or high performance computing resources. Scientific workflows are used to automate some of these processes, frequently involving high data transfer rates between varying resources (Figure 1).




Figure 1 – Gene Sequence Annotation

Computational facilities:

- Petascale Compute Facilities (e.g. Jaguar at ORNL)
- Regional Supercomputer (e.g. Chinook at PNNL)
- Campus or Community Grids (i.e. CONDOR pools)
- Cloud Computing (e.g. Amazon)

Typical output sizes e.g. on Chinook range from a few MB's to several TB, or 10's of TB's on Petascale facilities.

Networking:

Overall the community has a wide range of network requirements, many of which are problematic to reliably support at present. The complexity of the network infrastructure, ad hoc nature of the transfers, data sizes and the interaction of end user applications and networking are all contributing factors to this situation.

The severity of the networking challenges faced by the users varies depending on the size and rate of the data to be transferred and the regularity of the transfer. Small scale transfers (MBs to a few GBs) are relatively well supported today, although data collection in the 'field' is still a challenge. Medium range transfers (few 10's GBs) can be unreliable (e.g. lost connection), even more so when these are used for data streaming (sequence of experimental or observational measurements).

For large-scale data transfers it can be very difficult and time consuming to resolve network problems between any two sites. There are a number of things that cause performance problems including:

- bottlenecks that are not obvious (for example, security appliances that operate in stealth mode),
- bit errors on interfaces and links,
- queues that overflow.

There are usually multiple carriers that are participating in the end-to-end network path, and it is difficult to get any one carrier to take ownership of the problem. Experiences have shown that to "clean-up" a connection can take in the worst case several months. So if a connection is not of useful quality, it is usually going to take days if not weeks to resolve the problem. In this case the researcher would probably either find a work-around (i.e. send it in the post) long before the problem was resolved or give up, if this was an ad hoc requirement. For frequent/regular transfers between two sites, once a clean connection is established, it will be easy to maintain, though it will still require regular management and intervention to keep it working. Overall the effort required to set up transfers between two sites often outweighs the time required for the actual transfer, in particular for ad hoc requirements.

In summary it is very difficult for scientists or end user application developers to get the best performance out of the available network capacity, e.g. there are no firm figures or statistics on which size of transfers are best supported by which protocols, which sizes can be reliably transferred ad hoc or reliably transferred on a regular basis (and what is regular). In case of problems, support is difficult and time consuming given the number of different parties involved in any transfer, complicated even more by the communication (and software support) gap between the research scientists and networking personal, preventing either side to provide the best information to each other. As an example whilst there is some material available on highly performing protocols (bbcp) or how to tune the TCP stack for various OS these are often at a level that most scientists and application developers will not understand, leading to end user applications that counteract any network optimization efforts. Few sites have experts that can provide

hands on support to application developers and scientists to close the gap between application requirements and network performance.

Science Drivers – the next 2-5 years

Process of Science

The ultimate goal of achieving a predictive understanding of biological systems requires the integration of large volumes of methodologically, structurally and semantically heterogeneous information—functional descriptions assigned to DNA sequence, molecular interactions, images of physical structures within an organism, and details about the environment in which an organism lives. This mix of experimental, analyzed and modeled data will span diverse environmental conditions, spatial scales (nanometers to kilometers), and temporal scales (nanoseconds to decades).

Instruments and Facilities

The DOE will develop a Systems Biology Knowledgebase (Kbase) to provide researchers with an integrated infrastructure (experiments, modeling, data, analysis, workflows) to produce, access, analyze and synthesize a wide range of scientific research results collectively to further the predictive understanding of microbial and plant systems.

The knowledgebase will have many different data producing and consuming components (Figure 1), that will be distributed across not only the US, but worldwide. Users would be able to ‘publish’ their data, applications and workflows into the knowledge ‘cloud’ making these available for access and consumption by the wider community, adding to the continuous growth of the number of research resources. Links between people, data, applications and workflows could be explicitly recorded and published, or discovered following the semantic description ‘trail’. This would lead to a linked network of resources not dissimilar to the Linked Data Web as started by Tim Berners Lee (Figure 2), however extended with more domain specific information about the data and in particular applications and other resources to aid directed scientific discovery and experimentation. Furthermore it is anticipated that much of the computing needs of the community will be satisfied through flexible cloud computing resources, making it unpredictable where these will be physically located, and thus where data needs to be transferred to and collected from.

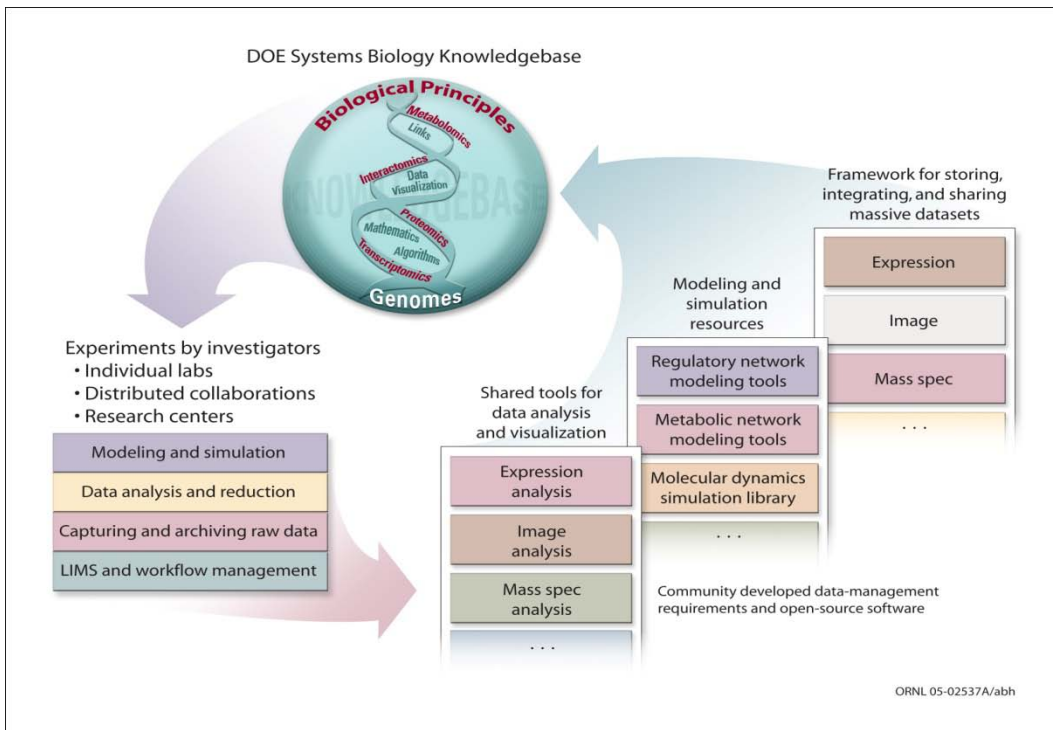


Figure 1 – DOE Systems Biology Knowledgebase Data Generating Components

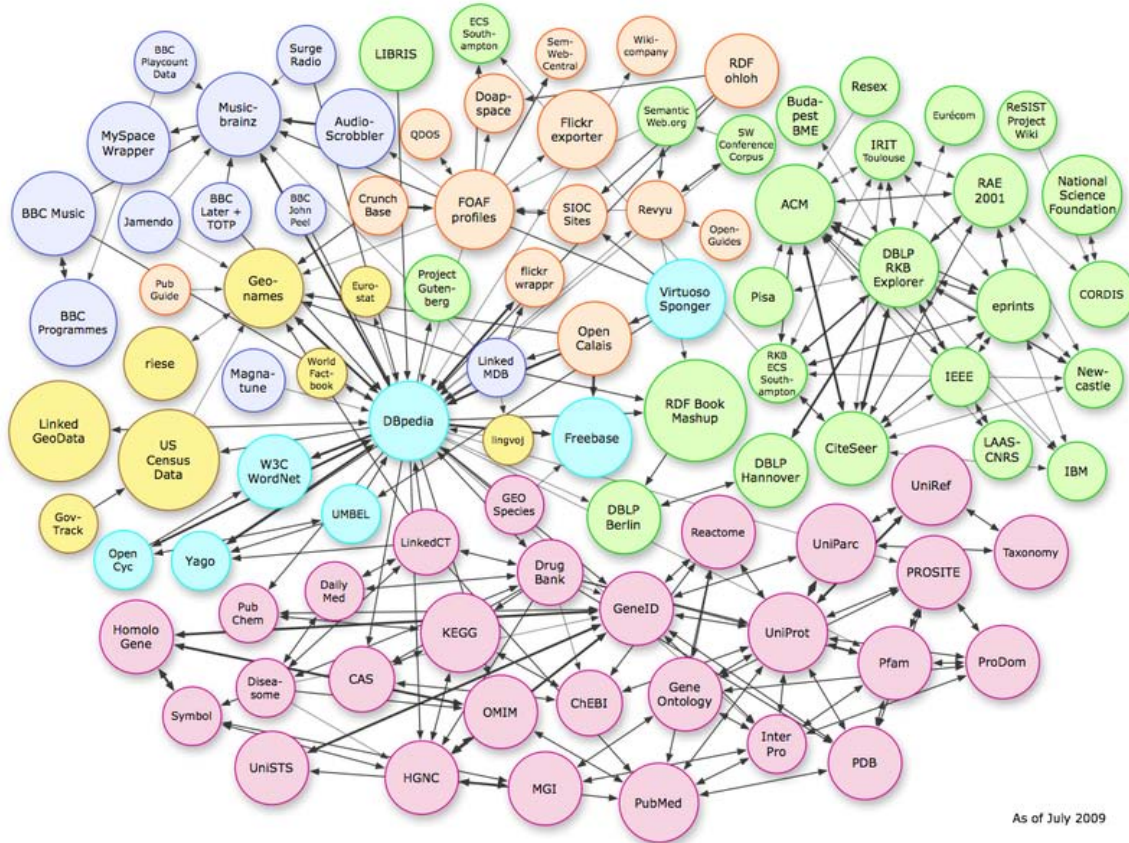


Figure 2 – Linked Data Web

Having enabled the discovery of the semantic, scientifically required connections between data, computing and experiments, scientists will want to combine these in ever changing ways to support their research. Leading to the following networking requirements:

- Reliable high capacity pipelines between key large scale data providers/producers and consumers (computing/analysis centers) where physical locations are fixed.
- Support for reliable and fast ad hoc data transfer between data providers/producers and cloud resources (data and computing).
- Easy access and transparent access to different transfer protocols, depending on end user application requirements.
- Training and support for end user application and transport protocol integration and optimization.

Summary Table

Feature	Key Science Drivers		Anticipated Network Requirements
Time Frame	Science Instruments and Facilities	Process of Science	Wide Area Network Bandwidth and Services
Near-term (0-1 years)	<ul style="list-style-type: none"> Access to Genomics, proteomics, microarray (transcriptomics) and metabolomics raw and derived data 	<ul style="list-style-type: none"> Access to key data centers for limited data sets, access to selected cloud based services. 	<ul style="list-style-type: none"> Reliable ad hoc transfer of 10s of GBs. Reliable transfer of 100s of GB between key sites.
1-3 years	<ul style="list-style-type: none"> Access to the above and new compute (cloud) and experimental services (incl. synchrotron, X-Ray free electron laser etc. high data rate facilities) 	<ul style="list-style-type: none"> Access to increasingly available data and compute resources for both in-depth and comparative analysis. Usage of key sites for data and computing. 	<ul style="list-style-type: none"> Reliable ad hoc transfer of 100s of GBs. Reliable transfer of 10's to 100s of TB between key sites. Training and support for end user application and transport protocol integration and optimization
3-5 years	<ul style="list-style-type: none"> Access to full range of (including researcher managed resources) data, application, workflow, computing and experimental services. 	<ul style="list-style-type: none"> Full utilization of the DOE Systems Biology Knowledgebase - an integrated network of data, computing, analysis and experimental facilities, combining different components in ever new ad hoc ways. Increased usage of automated workflows for high throughput, high data rate repetitive work processes. 	<ul style="list-style-type: none"> Support for reliable and fast ad hoc data transfer between data providers/producers and cloud resources (data and computing). – up to 10s of Terabytes Reliable high capacity pipelines between key large scale data providers/producers and consumers (computing/analysis centers) where physical locations are fixed. – up to Petabyte level Easy access and transparent access to different transfer protocols, depending on end user application requirements. Training and support for end user application and transport protocol integration and optimization.

Please note that for the science community it is less important what network bandwidth is available between two sites, but it is crucial to them how much data they can reliably transfer in a reasonable time both ad hoc and regularly. So quality of service and hands on support for end users to utilize the available capacity are key issues to address. The figures in the summary table and throughout the report reflect that notion.

8 Joint Genome Institute, Walnut Creek, CA

Background

The Joint Genome Institute (JGI) supports the Department of Energy's mission in the areas of bioenergy, carbon cycling, and biogeochemistry. JGI is DOE's sole production sequencing center. It sequences and annotates genomic data for nearly 1800 users from a variety of scientific community and DOE projects including the Bioenergy Research Centers (BRCs). JGI provides access to various genome data sets through a number of portals (Genome Portal, IMG, Phytozome, etc). JGI also supports the submission and analysis of sequence data from external users. These activities result in data movement both to and from JGI.

Key Local Science Drivers

Instruments and Facilities

JGI operates a variety of sequencing platforms (Sanger, 454, and Illumina) and requires significant computing resources to process and analyze the data. Currently, the Illumina GA-IIx sequencers are capable of generating the largest data volumes at around 20 Gb (billion base pairs) during a week long run. There are 12 Illumina GA-IIx in operation today. JGI expects to upgrade these to Illumina HiSeq in the coming year. The HiSeq platform is expected to output 150-200 Gb. The total data output (raw and processed) from these platforms is expected to exceed 1 PB/year within the next year. JGI has also recently increased its computing capacity through the help of ARRA funding. Clusters to support assembly and analysis will expand to roughly 4000 cores. In addition, there are around 20 specialized large-memory nodes to support the most demanding analysis workloads. Finally, JGI is expanding its storage to approximately 1.5 PB. Critical data is archived in NERSC's HPSS system.

In April 2010, JGI announced that the NERSC Division at LBNL would take over responsibility of operating the IT infrastructure for JGI. This includes relocating most of the IT infrastructure (including the analysis clusters and storage) to the Oakland Scientific Facility (OSF). This transition will require several months to complete. The production sequencing will continue to be performed at the Walnut Creek Facility but the data will be transferred over ESnet to NERSC for storage and analysis. Some computing resources will remain at Walnut Creek to insure sequencing runs smoothly and to provide temporary caching of data.

JGI and NERSC recently demonstrated running in a distributed fashion by working with ESnet to provision a 9Gbps layer 2 virtual circuit between the two facilities. Computing systems located at OSF were connected to an extension of the JGI network. As the transition of the IT systems continues, we expect to use a different model but will continue to require at least 10G network bandwidth between the sites. JGI is in the process of upgrading its IDS and standard network connections to support an upgrade of its general-purpose network connection from 1G to 10G. Based on current estimates, this should sustain JGI's data transfer needs between NERSC and JGI at least for the next 4 years.

Process of Science

JGI supports and performs genomic science by sequencing and analyzing data, as well as utilizing the end product data for scientific insight. JGI accepts Community Sequencing Proposals (CSP) for plant, microbial, fungal, and other organisms. These sequencing proposals include both single organism as well as metagenome sequencing. Accepted proposals ship wet samples to JGI for sequencing, assembly, and analysis. Users access the data for completed sequences through data portals such as the Genome Portal. In addition, the data is typically uploaded to GenBank for broader access. Data can also be analyzed through various Pipelines such as IMG and Phytozome. The results of this analysis are then made available through data portals specific to the Pipelines. In addition to analyzing sequence data produced at JGI, external users can submit sequence data to these Pipelines for analysis and publishing. Finally, these data products are used by research groups at JGI for scientific insight.

Key Remote Science Drivers

Instruments and Facilities

The largest users of JGI are the Bioenergy Research Centers (BRCs) at LBNL, ORNL, and University of Wisconsin-Madison. In addition to the BRCs there are a wide variety of Universities, Government Agencies, and Research Institutes who utilize the sequencing resources at the JGI.

Process of Science

JGI projects ship wet samples to JGI for sequencing and analysis. The data is then accessed through portals such as the Genome Portal. Depending on the type of sample and the preferences of the customer, the data may also be uploaded to NCBI.

In addition to shipping wet samples for sequencing, remote users also submit sequenced data to JGI for analysis by the various pipelines. This data is typically submitted through a web portal. The resulting analysis is either viewed through the portal or potentially downloaded. Today these data sets are typically less than 1GB. However, there may be hundreds to thousands of submissions. Since these are submitted via the web, the connection may vary from 100Mb to 1Gb at the remote site.

Local Science Drivers

Instruments and Facilities

Next generation sequencing platforms from Illumina and Pacific Biosciences are expected to increasing sequencing rates while also reducing the cost to sequence. In the past two years, sequencing cost has dropped by over an order of magnitude as measured on a \$/base cost. This trend is expected to continue. Consequently, facilities such as JGI are grappling with the expected growth for data generated locally as well as the volume of data submitted from external users. JGI expects to upgrade its Illumina sequencers in the coming year. This will significantly increase the data volumes that must be assembled, analyzed and eventually made public.

JGI in collaboration with the NERSC Division will operate computing systems to support this analysis. However, as the ingest rate increases, data rates may outgrow JGI's computing resources. Therefore, JGI is also evaluating how it can leverage HPC resources at NERSC and the LCF's (ANL and ORNL) to help meet this need. Furthermore, members of JGI in collaboration with the Systems Biology Knowledgebase program are exploring how cloud computing may be used to augment local resources. If cloud computing becomes a critical component of the analysis infrastructure for JGI then the network patterns could change significantly.

Process of Science

The overall process of sequencing, assembly, analysis and annotation will remain the same.

Remote Science Drivers – the next 2-5 years

Instruments and Facilities

It is anticipated that the lower sequencing cost will lead to a rapid growth in remote sites operating sequencers and submitting data to JGI for analysis. This will translate into more data being transferred across the wide-area network. As the computational demands grow, it is possible that sites will start to utilize cloud computing. This would impact how data flows between the remote sites, JGI, and Cloud providers.

Process of Science

It is not anticipated that the overall process will change.

Beyond 5 years – future needs and scientific direction

If the current trends continue in sequencing technology, then there could be increasing demands on the network. It is expected that more sequence data will come from external facilities and, consequently, JGI will be required to import and process more data.

Middleware Tools and Services

JGI currently relies primarily on in-house developed portals. Some of these portals do use third-party developed tools (i.e. Genome Browser) but most of these do not have a significant network component. Data transfers are typically performed through a web interface or through an FTP site. It may be beneficial to see how data-transfer tools can be integrated into these portals.

Summary Table

Feature	Key Science Drivers		Anticipated Network Requirements	
Time Frame	Science Instruments and Facilities	Process of Science	Local Area Network Bandwidth and Services	Wide Area Network Bandwidth and Services
Near-term (0-2 years)	<ul style="list-style-type: none"> • Illumina GA-IIx sequencers (12) • Roche 454 • GA-IIx expected to be upgraded to HiSeq in the 2011 	<ul style="list-style-type: none"> • Data are sequenced and possibly annotated at JGI • Results are uploaded to NCBI • Sequence data is submitted remotely for analysis 	<ul style="list-style-type: none"> • Upgrading to 10 Gbps 	<ul style="list-style-type: none"> • Upgrading to 10 Gbps
2-5 years	<ul style="list-style-type: none"> • Follow-on platforms from Illumina and new systems from PacBio • Use of clouds and DOE HPC systems to deal with surge computing needs 	<ul style="list-style-type: none"> • Same process, but growth of external data submission 	<ul style="list-style-type: none"> • 10 Gbps 	<ul style="list-style-type: none"> • 10 Gbps in addition to occasional use of 10 Gbps on SDN
5+ years	<ul style="list-style-type: none"> • Growing use of Clouds and DOE HPC systems for computation 	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> • 40 Gbps 	<ul style="list-style-type: none"> • 40 Gbps

9 EMSL, PNNL

Background

EMSL is a national user facility that provides world-class fundamental research capabilities for scientific discovery and the development of 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. 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 is located in Richland, WA and is operated by PNNL for the DOE's Office of Biological and Environmental Research.

Key Local Science Drivers

Instruments and Facilities

EMSL houses an unparalleled collection of state-of-the-art capabilities that are used to address scientific challenges relevant to the U.S. Department of Energy and the nation. Researchers from around the world are encouraged to use EMSL's unique capabilities in combination with each other with an emphasis on merging computational and experimental instruments. At present EMSL is developing data storage and archiving capabilities that will provide a central storage facility for scientific data from all instruments. The capabilities are described below.

Cell Isolation and Systems Analysis is used to isolate cells from complex populations or environmental samples for subsequent, integrated 'omics and imaging analyses. EMSL specializes in high-throughput genomics and proteomics studies as well as electron and fluorescence microscopy characterization at high spatial and temporal resolutions. These capabilities provide the foundation for attaining a molecular-level understanding of protein-network and microbial-community dynamics and enable the pursuit of systems biology, including the new field of systems microbiology. Instruments include cell isolation and fractionation resources, fluorescence microscopes and spectrometers, electron microscopes, and transcriptomics instruments to perform massively parallel next-generation sequencing.

Deposition and Microfabrication is used to tailor surfaces, atom by atom. With an emphasis on oxide mineral films and interfaces, users apply these tools to design and construct materials of various size distributions, ranging from high-quality, single-crystal thin films to nanostructures, with real-world applications. These materials are subsequently characterized and studied in detail using EMSL's mass spectrometry, microscopy, as well as spectroscopy and diffraction capabilities. Instruments include a

unique oxygen plasma-assisted molecular beam epitaxy system, a metalorganic chemical vapor deposition system, ballistic deposition, spin coating, and wet chemical synthesis systems, ion beam synthesis and ion beam-assisted deposition systems, and microfabrication equipment and clean-room capabilities.

Mass Spectrometry enables high-throughput, high-resolution analysis of complex mixtures. These resources are applied to a broad range of scientific problems from proteomics studies with applications to human health and environmental remediation to aerosol particle characterization, as well as fundamental studies of ion-surface collisions and preparatory mass spectrometry using ion soft-landing. Instruments include FT mass spectrometers, including 6-T, 7-T and 12-T FTICRs, five LTQ-Orbitraps, and three Orbitraps, five linear ion traps, three triple-quadrupole spectrometers, three ion mobility spectrometry (IMS) - time-of-flight (TOF) spectrometers, seventeen custom HPLC, three Agilent HPLC, a field-deployable second-generation single-particle laser-ablation time-of-flight mass spectrometer, and an ion soft-landing deposition instrument.

Microscopy has wide variety of sophisticated microscopy instruments, including electron microscopes, optical microscopes, scanning probe microscopes, and computer-controlled microscopes for automated particle analysis. These tools are used to image a range of sample types with nanoscale—and even atomic—resolution with applications to surface, environmental, biogeochemical, atmospheric, and biological science. Each of the state-of-the-art instruments and customized capabilities is equipped with features for specific applications. Instruments include electron microscopes with tomography, cryo, scanning, photoemission, and high-resolution capabilities, a NMR microscope, a dual Raman confocal microscope, optical microscopes, single-molecule fluorescence tools, spectroscopy tools with visible, near-, mid-, and far-infrared capabilities, atomic force microscopy, and scanning probe microscopes.

Molecular Science Computing provides an integrated production computing environment supporting a wide range of computational activities in environmental molecular research, archive storage, scientific expertise, and the Molecular Science Software Suite (including NWChem, ECCE, and the Global Array Toolkit). Systems include a 2310 node supercomputer (with a peak of 163 Tflops, consisting of dual quad-core AMD Opteron processors, 37 Tbyte of memory, a 300 Tbyte global file system), a 4.5 Pbyte hierarchical archive storage system called Aurora, and a graphics and visualization system.

NMR and EPR has NMR instruments with frequencies up to 900 MHz, as well as an electron paramagnetic resonance (EPR) spectrometer. EMSL staff is active in developing a variety of probes and techniques to complement their collection of state-of-the-art magnets. Instruments include eleven NMR spectrometers, one pulsed EPR spectrometer, an NMR with radiological capabilities, an NMR with metabolomics capabilities, extreme-temperature probes, and a combined optical and magnetic resonance microscopy instrument.

Spectroscopy and Diffraction has a suite of spectroscopy and diffraction instruments in EMSL allow users to study solid-, liquid-, and gas-phase sample structure and composition with remarkable resolution. Ideal for integrated studies, spectrometers and diffractometers are easily coupled with EMSL's computational and modeling capabilities,

allowing EMSL users to apply a multifaceted research approach for experimental data interpretation and to gain a fundamental understanding of scientific problems. Instruments include electron spectrometers, various electron microscopes, fourier transform infrared spectrometers, an ion accelerator system, five Mössbauer spectroscopy systems, optical spectroscopy tools including confocal-Raman, time-resolved fluorescence, and second harmonic generation capabilities, and multiple X-ray diffraction instruments.

Subsurface Flow and Transport provides a variety of flow cells, including column, batch, radial, wedge, and rectangular flow cells as well as microfluidics instrumentation. Flow cells are used in coordination with high-precision, high-sensitivity analytical tools to generate data about sample characteristics by detecting the presence of carbon, trace metals, ions, nonvolatile compounds, thermally labile chemicals, and more. EMSL users have the benefit of designing experiments using the predictive subsurface flow and transport simulator.

In 2009 EMSL received \$60M from the American Recovery and Reinvestment Act (ARRA) to fund the acquisition of capital equipment, and to refresh the facilities capabilities. EMSL has invested ARRA funds in 33 new and enhanced capabilities in microscopy and imaging, materials and bio characterization, nuclear magnetic resonance, mass spectrometry, radiological characterization, computational capabilities, and infrastructure and support. Some examples of instruments are a Helium ion microscope, a microbial transcriptome sequencer, and a 15T FT-ICR.

EMSL's internal network core is built around a fiber infrastructure that provides redundant connectivity to a standard 8 port network switch in each office and laboratory space. The capability exists to directly attach gigabit and ten gigabit instrumentation and computational resources directly to the EMSL core network. Isolated instrumentation networks can be created using the building fiber to interconnect lab spaces throughout the EMSL building, and in some cases extending into other buildings: over 400 systems are on isolated networks in EMSL to allow the use of antiquated operating systems and to allow instruments to acquire data without the potential interruption of network scanning or other cyber attacks. The EMSL core network has multiple connections into the PNNL core network through redundant fiber paths.

Process of Science

EMSL's capabilities are available to researchers through a peer-reviewed proposal process, at no cost, if research results are published in the open literature. Users access the facility to use one or multiple capabilities, and work with EMSL's expert staff to gain insight and knowledge into their scientific problem. A large majority of instruments at EMSL require hands-on work, and access to scientific experts from EMSL. Data is stored close to the instrument, and small data sets and results are moved to the user's home institution through email, and sometimes by media such as CD and thumb drives. In extreme cases, hard drives are shipped to the user's home institution, owing to the quantity of data, and the uncertainty of reasonable bandwidth between EMSL and the home institution. The Aurora storage archive at EMSL contains 450 terabytes of user data from the proteomics capability and the Chinook supercomputer. Presently, EMSL produces 1-2 terabytes of data weekly.

Key Remote Science Drivers

Instruments and Facilities

EMSL's users have remote access the Chinook supercomputer, the Aurora data storage archive, and use virtual NMR tools to remotely access the spectrometer systems, saving time and travel costs.

EMSL has multiple and redundant connections to ESNET (and the internet) via 10 Gigabit links through PNNL to Seattle (Primary) and Boise (Failover). The network was engineered with reliability and performance in mind: should one of the 10 Gigabit links suffer an incident that disrupts primary service, traffic is automatically failed-over to the redundant link.

Process of Science

EMSL's Chinook supercomputer and Aurora data storage archive access is obtained through ssh and sftp and requires each user to have an EMSL/PNNL specific one-time password token. BBCP is also available from both systems, but appears to be seldom used. The virtual NMR tools allow users to control these world-class instruments from their home institutions.

Local Science Drivers – the next 2-5 years

Instruments and Facilities

In addition to the ARRA projects EMSL is also in the early stages of developing a high resolution and mass accuracy capability (HRMAC). EMSL's current high-field mass spectrometry (MS) capability, although world-leading, is reaching technical performance limits, revealing a capability gap that precludes the solution of challenging scientific questions regarding chemical and biochemical structure and composition, particularly for mixtures with high mass and/or diverse composition. A next-generation, high-resolution and mass accuracy capability is needed to ensure that EMSL will continue to provide leading-edge resources to serve national and international users who are addressing critical DOE mission needs. Next-generation capability, taking advantage and applying newly developed technologies and approaches, will address this need, significantly enhancing overall analytical and characterization performance in terms of sensitivity, dynamic range, accuracy, resolution, and speed or throughput, and enable previously intractable types of applications. In addition, EMSL will be commissioning a new quiet building that will house (among others) an OPA-MBE, a nano-SIMS, and a FE-CTEM instrument with the capability to generate large high-resolution data sets.

EMSL's process for a planned upgrade of the molecular science computing capability will start in FY10. The Molecular Science Computing Greenbook, generated by EMSL's users, outlines a need for data-intensive computing. Hence, in addition to scientific computing the supercomputer will increasingly be used to provide real-time analysis of the experimental data streaming of the scientific instruments. In addition, in the next five years the archive storage system will be upgraded to reach 20 petabytes of data storage.

The common denominator is increased resolution and increased data rates coming off of the instruments. New data management policies and processes will improve EMSL's ability to make unique data available to the scientific community. This should make EMSL a supplier of at least hundreds of terabytes of data to ESnet's users. It is anticipated that the archived proteomics data will be accessed with increasing frequency as its use in gene annotation becomes more common. Thus the volume of accessed data should increase by 2-5 fold in the next several years. Peak bandwidth to allow timely access to remote instrumentation will also continue to be a factor.

Process of Science

No major changes in the way scientists use the instruments in the EMSL facility are anticipated. One change that will occur in this timeframe is the development and use of a centralized data storage and access system. The storage, analysis and dissemination (often on a variety of media with disparate formats) of EMSL-produced data are an issue for EMSL users, management and advisory boards. The Aurora storage archive at EMSL contains 450 terabytes of user data from the proteomics capability and the Chinook supercomputer. Presently, EMSL produces 1-2 terabytes of data weekly. With the addition of new equipment, the output is expected to increase to tens of terabytes each week. To address this issue, EMSL is implementing an EMSL wide scientific data management system (called MyEMSL) that accurately and reliably stores data acquired from EMSL's experimental and computational instruments and the output of analysis software together with relevant metadata and allows users a simple way to find, retrieve, visualize, and mine their stored data. The data management system will provide a simple and consistent interface for the EMSL staff operating the instruments and for the users accessing and sharing their data. Once MyEMSL reaches its full potential, it is likely to be a catalyst for creation of new data and discoveries derived from existing data.

Remote Science Drivers – the next 2-5 years

Instruments and Facilities

At present the expansion of remote access to capabilities in EMSL will be limited. Many of the instruments at EMSL are not well suited for automation, or require hands-on work, or a scientific expert from EMSL to work with the user. In the 2-5 year timeframe the next-generation supercomputing capability will become available to the EMSL user community. Increased size and number of raw data sets will make it more difficult for users to move the data to their home institutions for analysis, which drives the need for increased access to remote analyses and visualization capabilities at EMSL.

Process of Science

No major changes in the way scientists use the remote instruments in the EMSL facility.

The storage, analysis and dissemination (often on a variety of media with disparate formats) of EMSL-produced data are an issue for EMSL users, management and advisory boards. The Aurora storage archive at EMSL contains 450 terabytes of user data from the proteomics capability and the Chinook supercomputer. Presently, EMSL produces 1-2

terabytes of data weekly. With the addition of new equipment, the output is expected to increase to tens of terabytes each week. To address this issue, EMSL is implementing an EMSL wide scientific data management system (called MyEMSL) that accurately and reliably stores data acquired from EMSL's experimental and computational instruments and the output of analysis software together with relevant metadata and allows users a simple way to find, retrieve, visualize, and mine their stored data. The data management system will provide a simple and consistent interface for the EMSL staff operating the instruments and for the users accessing and sharing their data in a collaborative fashion.

Beyond 5 years – future needs and scientific direction

The vision that directed the development of EMSL has led to significant scientific progress. EMSL plans to greatly increase its scientific impact during its second decade of operation by focusing attention and capability development in specific areas identified as high-priority science themes. These science themes help define and direct development of key capabilities and collections of user projects that can have significant impacts on important areas of environmental molecular science that are critical to DOE and the nation.

As EMSL's user research expands and matures, guided by the science themes, new and enhanced capabilities will be developed. Additionally, existing systems will be modified to support the needs of the user community driven by the EMSL Science Themes. Beyond 5 years we should see growth in the body of biological data that EMSL and PNNL will maintain for searches and analysis. The increased focus on systems biology, as well as chemical imaging capabilities will significantly increase the data volumes for complex samples analyzed by EMSL. A new generation of mass spectrometers for proteomics applications are being developed that should increase sample throughput and data output by multiple orders of magnitude. Access to the massive sets of data generated by these new instruments will significantly increase network requirements. Much of this data will be transferred in from offsite, combined with existing data, curated, and shared back with the scientific community. Strong integration of data from multiple scientific domains to allow users to address systems-level problems will require EMSL to manage and integrate multi-petabyte-scale datasets. This will require the development of complex workflows accessing data generated and stored at EMSL with data from other user facility and research laboratories, which will significantly impact network requirements.

Middleware Tools and Services

Middleware and services will likely need to be adopted in order to facilitate data sharing and remote analysis and visualization for offsite EMSL users and collaborators. It has yet to be determined what middleware and services might be required.

Outstanding Issues

Remote access to EMSL's computing capabilities require a site-specific one-time password token. With users having access to many facilities, they end up with a plethora of tokens. The adoption of a complex wide one-time password token will greatly enhance the experience of users.

Over 50% of the EMSL users are from academic institutions. To enable data transfer and good quality of service to remotely control EMSL instruments, robust connectivity and high bandwidth to sites both on and off ESnet are essential.

Summary Table

Feature	Key Science Drivers		Anticipated Network Requirements	
	Science Instruments and Facilities	Process of Science	Local Area Network Bandwidth and Services	Wide Area Network Bandwidth and Services
Near-term (0-2 years)	<ul style="list-style-type: none"> Broad suite of instruments available to users at EMSL 33 new instruments will become available through ARRA funding 	<ul style="list-style-type: none"> Primarily onsite access to instruments Remote access to Chinook computer, Aurora archive, and Virtual NMR MyEMSL data management system will store data (and metadata) of all instruments 	<ul style="list-style-type: none"> Currently 450 terabytes of data in archive Data accumulation of 50-100 terabytes per year 	<ul style="list-style-type: none"> Remote access to data stored at EMSL Remote control of instruments at EMSL Reliable transfer of 10-100 gigabytes
2-5 years	<ul style="list-style-type: none"> Next-generation supercomputer Next-generation mass spectrometer Quiet wing with hi-res chemical imaging capabilities Systems biology 	<ul style="list-style-type: none"> Enhanced integration of data from multiple instruments Collaborative data access and analysis MyEMSL with search and first set of workflow and analysis capabilities 	<ul style="list-style-type: none"> Significant data increase due to new capabilities, with data accumulation of 500 to 1000 terabytes per year Onsite access to full data sets 	<ul style="list-style-type: none"> Remote access to data to share, mine, and analyze remotely Limited expansion of remote control access Reliable transfer of tens of terabytes
5+ years	<ul style="list-style-type: none"> Next-generation instruments driven by user community and Science Themes 	<ul style="list-style-type: none"> Strong integration of data across capabilities Comprehensive problem-solving environment on top of MyEMSL 	<ul style="list-style-type: none"> Data accumulation expected to grow to multiple petabytes per year Onsite real-time analysis, manipulation, and visualization of full data sets 	<ul style="list-style-type: none"> Remote access to data with potential for moving large (multi-petabyte) data sets Remote real-time analysis, manipulation, and visualization of data sets Reliable transfer of 10-100 terabytes, potentially petabytes

10 Great Lakes Bioenergy Research Center

(University of Wisconsin-Madison, Michigan State University, Illinois State University, Iowa State University, Lucigen, University of Missouri, Oak Ridge National Laboratory, Pacific Northwest National Laboratory)

Background

Great Lakes Bioenergy is one of three Bioenergy Research Centers (BRCs) established by the U.S. Department of Energy Office of Science to provide basic science breakthroughs for a cellulosic biofuels sector. Activities of this academically-based center are carried out by integrated teams of university, national laboratory, and industrial scientists and staff. The Center approaches cellulosic biofuels research with programs having near-, mid-, and long-term goals or impacts. To achieve its mission, the Center assembles leading biological, physical, and computational scientists into the intellectual core of systems-driven, genome-informed, high-throughput basic science programs. Individual Great Lakes Bioenergy research activities are led by experts in plant biology, microbiology, molecular or cell biology, biochemistry, protein design, engineering, computer sciences, systems analysis, and ecology. In keeping with the roadmap of the DOE Genomic Sciences program, the Great Lakes Bioenergy basic science programs analyze cellulosic biofuels production across a variety of scales, including molecular, genomic, organismal, population, and global levels.

The Center's discovery research thrusts seek to improve (a) the characteristics of plant biomass (Thrust 1); (b) the procedures for processing plant biomass (Thrust 2); (c) the processes used to convert biomass into fuels (Thrust 3); and (d) the economic and environmental sustainability of the biomass-to-biofuel pipeline (Thrust 4). The four Research Thrusts are integrated with the Enabling Technologies (ET) and Informatics and Information Technologies (IIT) groups, which together provide access to the Center's high-throughput analytics and integrated data management and analysis tools.

GLBRC currently comprises some 44 inter-related research projects, 8 enabling technology cores, and supporting informatics and operations functions at 8 institutions.

Key Local Science Drivers

Instruments and Facilities

The following table lists major data producers by experiment type.

Experiment type	Instrument	Samples (current)	Samples (1-2 years)	Approximate Data quantity
Metagenome sequencing	(See JGI)	3 communities	9 communities	500MB/metagenome = 1.5 - 4.5GB
Genome sequencing	(See JGI)	2 plant genomes, 6 microbial		10 GB per plant genome, 10 MB per microbe = 21 GB total
Genome re-sequencing	(See JGI)	2 corn strains, switchgrass		20 GB
EST Sequencing	454 pyrosequencing, Illumina (see JGI)	~ 7.5 x 10 ⁶ ESTs		2 GB (1 TB raw trace data)
Restriction site adjacent sequencing		1000 genotypes		3.5 x 10 ⁶ bp per genotype (7 GB total)
Transcriptomics profiling	DNA probe microarray, RNAseq	50-100 arrays/yr		80 GB/yr
Proteomics profiling	LC-MS/MS (see also PNNL)	300 samples, 3000 runs / yr.		100 MB/yr (1 TB raw data)
Biomass characterization	Sugar, protein concentration	~ 500k/yr		< 50 MB/yr
Metabolomics and Cell wall profiling	NMR			1-5 GB/yr
Cell wall composition analysis				90 MB/yr
Transcriptional start & stop site mapping	RNA end sequencing (see JGI)	~ 117 x 10 ⁶ reads	~ 200 x 10 ⁶ reads	< 20 GB
Mapping DNA binding sites of proteins (e.g., transcription factors)	ChiPSeq (see JGI)	~ 60 x 10 ⁶ reads	~ 500 x 10 ⁶ reads	< 50 GB
Biogeochemistry, biodiversity, trace gases, weather, etc. measurements				< 20 MB
Economic, biophysical & biogeochemical modeling & simulation				?

Over the initial five-year funding period of the GLBRC, while several terabytes of raw data (e.g., chromatograms, mass spectra, sequencer trace files) will be produced, the entire Center will generate on the order of 500 GB of finished experimental data and

metadata. Data sets produced by sustainability modeling and simulation (e.g., Life Cycle Analysis) could potentially add significantly to this total.

Central computing infrastructure for GLBRC at UW-Madison is located in two partially redundant data centers. The main site has 10 physical servers containing 95 x86-64 cores typically running in the 2-3 GHz range, as well a total of 400GB of RAM. These servers have SAN access to 60TB of raw disk storage. The second site has 5 physical servers with 24 similar CPU cores, 100GB of RAM, and SAN access to an additional 16TB of raw disk storage. The two sites each have redundant 1Gbps links to the UW campus backbone, although higher speeds are available. The UW backbone is in turn connected to high-speed research networks such as Internet2. Likewise, the GLBRC computing infrastructure at Michigan State University and the Kellogg Biological Station are both connected to Internet2 via 1 Gbps connections.

Process of Science

Thrust 1: Improving plant biomass. Great Lakes Bioenergy has two major long-term goals for improving plant biomass. The first, and most important, is to increase the quantity and accessibility of easily degraded polysaccharides in the cell walls of biofuel crops. The second is to increase the yields of energy rich hydrocarbons (oils) in vegetative tissues. To achieve these long-term goals, the GLBRC will strive to decipher, model, and redirect plant carbon into pathways that increase their value to the bioenergy marketplace.

The first goal is being pursued in two ways: (1) generate fundamental understanding of the pathways mediating plant cell wall polysaccharide biosynthesis and the regulatory mechanisms that control polysaccharide accumulation and (2) modify the lignin in ways that allow the release of sugars with milder chemical pretreatments. This knowledge will eventually be used to produce (1) plants with more easily degraded polysaccharides that contain a higher fraction of easily fermentable sugars and (2) plants that synthesize novel lignin precursors that, when incorporated, produce lignin polymers that can be more easily degraded during pretreatment.

Efforts toward the second goal have focused on using a genome-enabled analysis of the relevant metabolic and regulatory networks that control hydrocarbon synthesis and accumulation.

Thrust 2: Improved biomass processing. Deconstruction of plant biomass to generate a feedstock stream of sugars and other convertible organics is a major bottleneck in the production of cellulosic biofuels. Our activities in this Thrust are guided by the hypothesis that new and better-integrated physical and biological treatments are needed to improve processing of the different biomass types envisioned for the biofuels pipeline.

To date, the Center has focused on developing novel combinations of physical, chemical, and enzyme treatments to process plant biomass and tailoring them to the composition of the biomass. In support of this activity, the Center deployed a moderate- to high-throughput bioprospecting program to query relevant environments for novel biomass processing microbes and enzymes.

The Center's current enzyme efforts focus on identifying and characterizing the proteins required for effectively releasing sugars from plant biomass. Our approach is to use enzymes that flow from our screening platform as the basis for new, defined mixtures of purified catalysts to replace the existing, crude mixtures of enzymes. By allowing each enzyme in a mixture to be replaced with newly identified candidates, our combinatorial approach allows a systematic comparison of whether different native or engineered catalysts are superior to their predecessors. Efforts are also underway to query novel biomass degrading ecosystems to identify new organisms and/or enzymes that will improve sugar release from pre-treated biomass.

Thrust 3: Conversion of biomass into biofuel. The research programs in Thrust 3 are guided by the premise that new basic science knowledge is needed to enable transformative breakthroughs in microbial conversion of lignocellulosic biomass to biofuels. To generate this knowledge, Great Lakes Bioenergy uses genomic, cellular and computational analyses to (1) understand current roadblocks in microbial fuel synthesis and (2) develop strategies for strain improvement, using an iterative program of strain selection, computational predictions, and metabolic engineering.

GLBRC has built on strengths in analyzing the anaerobic lifestyles of model or potential biofuels microbes. These strengths are important since most current or envisioned biofuels are more reduced than the sugars present in the feedstock stream. This knowledge is being used to derive genome-enabled predictions of the transporters, metabolic pathways and regulatory networks that can be altered to increase the flow of carbon skeletons and reducing power into precursors for biofuel production in the absence of oxygen.

Thrust 4: Development of a sustainable bioenergy economy. For bioenergy to become an integrated, positive economic force it must be sustainable – economically, environmentally, and socially. Understanding sustainability outcomes requires modeling informed by measurements and experimental results that expose the interactions among these factors.

Sustainability research in Great Lakes Bioenergy aims to understand the environmental and economic responses of different biofuel production systems in order to develop cropping systems that are optimally configured to maximize ecosystem services and productivity. We have established field studies at multiple spatial scales to provide needed knowledge about novel cropping systems, microbe-plant interactions, and the biogeochemical, biodiversity and economic responses of these systems. A major simulation effort integrates biophysical, economic, and life-cycle models to provide a capacity to test alternate design and deployment scenarios.

At the field scale we have established a replicated set of intensive field experiments in Michigan and Wisconsin. At each location we grow model crops along a diversity gradient that ranges from annual monoculture to native prairie. Within these plots we investigate productivity and biomass quality, nitrogen and carbon balances, greenhouse gas fluxes, water use and water quality, pest dynamics, and microbial and insect diversity. At the landscape scale – in 60 paired fields across Michigan and Wisconsin – we examine the biodiversity of predaceous insects, pollinators, and birds to understand the impact of biofuel crops on organisms that provide important ecosystem services in

larger landscapes. At the subregional scale – in four sets of multiple-county groupings in Michigan and Wisconsin – we apply actual land use and built infrastructure frameworks to model impact and costs of various biofuel scenarios.

Key Remote Science Drivers

Instruments and Facilities

See local instruments and facilities, above.

Process of Science

Most current research workflows follow one of two (or a combination of the two) models with regard to remote drivers:

- A) Large-scale data are produced remotely (e.g., at JGI) and transferred once to the requesting lab and/or a central GLBRC compute and/or storage server, where the analysis and integration with other high-throughput data occurs.
- B) Data are produced locally (e.g., in one of the Enabling Technology high-throughput core facilities) and transferred once to the requesting lab and/or a central GLBRC server where analysis and integration occurs.

In either case, collaborative data sharing is generally on a much-reduced (processed) data set. Where a public data repository (e.g., GenBank) is available, the data are deposited in compliance with DOE data sharing policy. Multi-omics data are archived locally and will be deposited in the DOE Systems Biology Knowledgebase as it becomes available. GLBRC sustainability field measurements are made publicly available through the Kellogg Biological Station Long Term Ecology Research (LTER) database.

Local Science Drivers – the next 2-5 years

Instruments and Facilities

Current compute and storage capacity have been scoped to accommodate projected needs for projects as proposed through 2012, with modest incremental growth. If novel, much higher throughput research pipelines are implemented, these plans will be revised. A more likely scenario is that new compute-intensive analytical or simulation methods will be introduced. In this case, we will make use of computing capacity available through Condor pools and the University of Wisconsin Center for High Performance Computing.

Process of Science

Thrust 1: Improving plant biomass. GLBRC researchers will use knowledge of the genome-enabled blueprints of carbon flow during biomass deposition to perform rational modification of grasses, crops, and trees that can supply the biomass feedstocks for biofuels production.

Thrust 2: Improved biomass processing. We plan to continue bioprospecting efforts as long as new organisms with superior enzymes are found. If, as we move forward, it

appears that our pipeline is no longer identifying candidate enzymes with improved features at a reasonable rate, we will direct more of our efforts to improving existing enzymes or mixtures of enzymes. When this point of diminishing returns from bioprospecting is reached, the amount of high-throughput metagenomic and genomic sequencing will be reduced. It is arguable, however, that it is unlikely that this point will be reached within the next five years.

Thrust 3: Conversion of biomass into biofuel. At present most investigations are focused on gram-negative bacteria, especially *Escherichia coli*, and on brewers yeast, *Saccharomyces cerevisiae*. However, our bioprospecting studies (see Thrust 2, above) are identifying organisms that appear to have novel paradigms for biomass degradation, and we may initiate focused study of one or more of these organisms to dissect the features that make these microbes superior to existing strains.

Developing consolidated bioprocessing microorganisms (that both produce the enzymes needed for releasing sugars from pretreated biomass and convert the released sugars into ethanol or next generation biofuels) will require integration of results derived from pretreatment and enzyme efforts (Thrust 2) and systems biology and metabolic engineering. Another important goal is applying directed evolution, genome-enabled multi-omic assays, and computational predictions to improve microbial tolerance to toxic metabolites derived from pretreatment or fuel molecules produced during fermentation.

Thrust 4: Development of a sustainable bioenergy economy. We plan to extend the scope of our modeling activities by establishing a network of satellite sites across the Midwest. We will recruit cooperators with established or newly initiated cropping systems identical to several of those in our intensive field plots (e.g. continuous corn, switchgrass, and old field communities), and use measurements in these systems to test and – as necessary – further parameterize our biophysical models of productivity, carbon and nitrogen flow, water balance, and greenhouse gas fluxes. By incorporating a larger range of soils and climates we will significantly improve confidence in our modeling results and thereby the validity of models to predict outcomes in regions different from the upper Midwest.

Remote Science Drivers – the next 2-5 years

Instruments and Facilities

Current and projected rates of data production are essentially budget limited. However, transformative technologies may reduce costs by an order of magnitude or more and enable experimental approaches that are currently infeasible. These technologies will produce proportionately greater quantities of data with corresponding challenges of management, storage, and analysis.

Process of Science

No change.

Beyond 5 years – future needs and scientific direction

As fundamental knowledge is developed in Research Thrusts 1, 2, and 3, bioprospecting for new candidate structures and catalysts will be reduced or eliminated and research efforts will focus intensively on a relatively small number of targets. While we expect that this research phase will involve generation of substantial multi-omics data, the quantity of genomic and metagenomic sequence will probably be reduced.

Middleware Tools and Services

Initiating and sustaining active scientific collaborations across the GLBRC sites has been a challenge. To supplement the limited opportunities for face-to-face meetings the GLBRC has recently installed high-definition video conference rooms at the three principal sites (UW-Madison, MSU, and the Kellogg Biological Station). These rooms are already being heavily used to support scientific discussion among researchers across the campuses. We anticipate that the level of use will increase as more researchers are exposed to the advantages of this mode of interaction.

Summary Table

Feature	Key Science Drivers		Anticipated Network Requirements	
Time Frame	Science Instruments and Facilities	Process of Science	Local Area Network Bandwidth and Services	Wide Area Network Bandwidth and Services
Near-term (0-2 years)	<ul style="list-style-type: none"> • Metagenome sequencing • Genome sequencing • EST & restriction site adjacent sequencing • RNA end sequencing • ChiPSeq • Transcriptomics • Proteomics • Cell wall profiling • Environmental measurements • Biome-BGC & EPIC environmental model simulations 	<ul style="list-style-type: none"> • Transfer of high-throughput experimental datasets from providers (e.g., JGI, PNNL) to GLBRC • Transfer of annotated datasets to public repositories • Exchange of reduced datasets and analytical results among Center investigators and with external collaborators. 	<ul style="list-style-type: none"> • Campus network bandwidth of 1 Gbps 	<ul style="list-style-type: none"> • Total external network bandwidth requirements of 1 Gbps
2-5 years	<ul style="list-style-type: none"> • (Potentially) transformative high-throughput 'omics methods 	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> • Campus network bandwidth of 1 Gbps 	<ul style="list-style-type: none"> • Total external network bandwidth requirements of 1 Gbps
5+ years	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> •

11 Joint BioEnergy Institute, Emeryville, CA

Background

JBEI is one of the three DOE BioEnergy Research Centers dedicated to advance the development of the next generation of biofuels. Major foci of JBEI are to understand the recalcitrance of biomass as well as the mechanistic details of biomass deconstruction, and to reengineer microbes for the efficient production of transportation fuels. Such efforts require a system biology approach that is based on the symbiotic interactions among four main technologies: Microsystems, Omics, multi-scale imaging, and integrated bioinformatics. While the first two are becoming high-throughput technologies, the other ones remain a bottleneck. For this case study, we will focus on multi-scale imaging for it is the technology that can benefit the most from a high-bandwidth, low latency network service, and because it can provide the framework that can be used to integrate the other technologies.

Today, advances in 3D imaging techniques allow scientists to easily collect (tens of) gigabytes of 3D volume data of plant cell walls and microbial communities, including their molecular organization and cellular machine distribution, and we can expect an explosion of 3D volumes over the next decade as advances in hardware and software will allow the automated data collection and 3D reconstruction of several dozens of 3D data sets per day. However, the inspection of the resulting volumes remains a tedious and time-consuming process that requires the slicing and manual segmentation of the volumes as well as the integration of experimental data from many sources (Genomic Encyclopedia of Bacteria and Archaea, MicrobesOnline, metagenomics, mass spectrometry, cryo-electron tomography, and others) to produce a complete picture of a biological system. Typically, different pieces of data are collected by different research groups focused on particular aspects of a common problem. Therefore, principal investigators and their teams need to travel to where the data is analyzed in order to provide their input for comprehension and interpretation of the data. However with increased complexity and hence the need of prolonged analysis, such scenarios will become impossible. Therefore, in order to realize the potential of true collaborative analysis on complex multidimensional imaging data sets, we content that the 3D data must be stored in a common repository to make it available to the scientific community, as is common for Omics type of data. Such repository should provide three levels of access to the data. The first level represents the raw data, the second level comprises the results of data analysis and annotations, and the third level consists of integrating different data sources such that data from multiple datasets can be simultaneously queried.

As the 3D imaging data becomes centralized, an infrastructure must be developed to enable the remote analysis and interactive visualization of the data as well as the interactive collaboration of groups of scientists located at different sites. We believe this goal is well-aligned with DOE's mission to establish a Systems Biology Knowledgebase that promotes a new era in biology.

Key Local Science Drivers

Instruments and Facilities:

Currently, most investigators collect their data in their own laboratories, where they house transmission electron microscopes (cryo-EMs) for single particle and electron tomographic data sets. In addition to such in-house sources, several national centers such as those at the Scripps, the University of Colorado at Boulder, and the Baylor College of Medicine exist, where high quantities of data can be collected. As a matter of fact, microscope prices have risen to the point that top end microscopes are beyond the reach of most of individual investigators, and it seems only a matter of time that most cryo-EM data will be collected in a high-throughput manner by centralized centers. For Focused-Ion-Beam Scanning Electron Microscopic Imaging, JBEI has recently obtained access to suitable machines at the National Center for Electron Microscopy (NCEM) and the Molecular Foundry, allowing 3D data collection of resin-embedded, stained microbial community samples.

Currently most of the data collected by Transmission Electron Microscopes (TEMs) and Focused Ion Beam Scanning Electron Microscopes (FIB/SEMs), in particular in its raw form is stored on local hard drives in the PIs lab. Often such data sets are mined only for a very specific question, and upon publication the data gets stored away and never again sees the light of day. Structural databases may be populated with a selected portion of the original data, but the actual raw data is virtually never uploaded, in part due to its size. This is most unfortunate, as 3D EM images are very rich in content and would likely be of high interest to a larger scientific community and could be mined for different questions, yet, most of the data is eventually buried and hence the opportunity for gaining additional insight is lost.

Process of Science:

JBEI electron microscopists currently generate an average of two 3D-image data sets that are about 5-10 GB in size per day, using either TEMs or FIB/SEMs. The FIB/SEM is of particular interest as it allows data volume of 10 by 10 by 10 microns to be imaged, which is enough for the identification and 3D spatial mapping of microbes at the community level in order to enable realistic modeling of microbial behavior. They store those images locally on their workstation's hard drive, and slice portions of it to fit their desktop's graphics card for analysis and visualization. Segmentation is usually performed by hand on individual data slices, or by placing seed points. Alternatively, automated segmentation algorithms can be employed, but due to the noise in the data and the complexity of the volumes, additional steps of correction of such automated approaches are necessary and usually equally time consuming. Thus, although biologists have access to 3D volumetric datasets, which in turn contain 3D biological structures, they are often constrained to 2D operations since they simply do not have the interaction capabilities or the software functionality to understand the data in 3D, or they are simply overwhelmed by the complexity of the 3D data sets, in particular upon rendering of large volumes. This is in part due to the size, the noise within the data set, and its sheer complexity, which makes comprehension as 3D volume or surface rendering confusing and impracticable.

To complicate matters, the analysis of 3D images also requires collaborations with colleagues with different expertise who usually need to travel to discuss the results in person. Clearly, having the data centralized and enabling remote interaction with the data as well as remote collaborations would enormously accelerate the process of analysis and interpretation.

Key Remote Science Drivers

Instruments and Facilities

None (see below).

Process of Science

Currently, scientists do not take advantage of the tremendous potential of collaborative environments because they do not have the necessary infrastructure that enables interactive visualization of data stored remotely. Data is stored locally and transferred to collaborators via anonymous ftp and collaborators travel when more interactive discussions are needed. Clearly, this modus operandi must change to accelerate the process of discovery.

Local Science Drivers – the next 2-5 years

Instruments and Facilities

Within the next decade we can expect the further centralization of data acquisition as such high-end microscopes are increasingly costly and require substantial maintenance. The exciting recent development in FIB/SEM imaging make it likely that a fair number of such microscopes will become online before too long, and will substantially increase the amount of image data being generated.

Process of Science

In the next 2-5 years we expect the data to be increasingly stored in a centralized repository and therefore, the storage and interaction with the data will be done remotely. We discuss the process of science in the next section.

Remote Science Drivers – the next 2-5 years

Instruments and Facilities

The explosion of 3D images will create the need for storage, backup, and sharing. We believe it is highly necessary to create and maintain a central repository for 3D images. In the next 2-5 years, there will be about 40-50 groups in the world (about 10-20 in the US) developing 10-50 3D images per day, each between 5-10 GB in size. This requires between 2-25TB of storage per day. A repository is needed that can handle this massive downloads as well as servers that run the visualization and analysis tools remotely. The tools needed for 3D imaging analysis must be highly interactive and so, scientists will need a network system that allows for very low latency. In fact, one of the major

challenges of this approach will be to ensure close feedback between human actions and visualization at the human time scale. In other words, if a user interacts with the 3D data, these manipulations must be reflected back to the user with minimally perceived delay. If the delay between an action and perception of its effect is longer than around 0.1secs, then users will subconsciously overcompensate and the interactive process will become cumbersome.

With regards to bandwidth, this cloud-computing approach should run the visualization software remotely and stream the video images at a video rate transmission of about 30 frames/sec. Estimation of the 3D, high-definition, stereo images at 10MB translates to a 5 Gbps bandwidth requirement. Thus the user-repository connections should be low-latency and have 5 Gbps of bandwidth.

Process of Science

Figure 1 illustrates the ideal scenario for handling and exploiting the flood of 3D images that will be generated in the next 2-5 years. Users 1, 2, and 3 are at the 3 DOE BioEnergy Centers located at ORNL, Univ. Wisconsin, Madison, and Emeryville, CA, respectively. Data are in repository R1 and users 1, 2, and 3 are conducting a collaborative session that allows them not only to manipulate, segment, and annotate a 3D data set in real time but also to intuitively interact with the visualized biological entities. These interactions require very low-latency and enough bandwidth to allow for video rate transmissions of about 30 frames/sec. Notice that there will be other users accessing other 3D images in the repository. If such data analysis requires international collaboration (User 4), one requires mirrors repository sites, e.g. in Europe and Asia. Hence communication and version control is needed between the various repositories. And one or several backup-sites must exist to secure the data and keep track of the work process and different versions of the data set (e.g. raw data, noise-reduced data, segmented data, etc...)



Figure 1. Storing and accessing 3D images in the next few years.

Beyond 5 years – future needs and scientific direction

Beyond 5 years, we anticipate the creation of repositories in Europe and Asia. Ideally, these repositories should contain mirror copies, which will add a new level of connection between repositories. In addition to central repositories, a backup system will be needed. The repository-backup connections should support higher bandwidth than the user-repository ones. We estimate that a bandwidth of 160 Gbps should be adequate for this level.

Figure 1 shows the necessary infrastructure that includes the mirror repositories in Europe and Asia as well as a backup system. This infrastructure will require three different levels of latency and bandwidth:

- Level 1. User-Repository: Very low latency and 10 Gbps bandwidth
- Level 2. Repository-Repository: Higher latency and 15 Gbps bandwidth
- Level 3. Repository-Backup system: Higher latency and 160 Gbps bandwidth

Middleware Tools and Services

The availability of sufficient bandwidth will spur development of the interactive collaborative visualization and analysis tools just as cloud computing has become popular for the much easier task of collaborative editing of word processing documents. These tools will allow scientists to take full advantage of the qualitative and quantitative information provided by the vast image space as well as to gain insights at many levels of biological scale.

Summary Table

Feature	Key Science Drivers		Anticipated Network Requirements	
Time Frame	Science Instruments and Facilities	Process of Science	Local Area Network Bandwidth and Services	Wide Area Network Bandwidth and Services
Near-term (0-2 years)	<ul style="list-style-type: none"> • Data collected by TEMs and FIB/SEMs is stored on local hard drives in the PIs lab • JBEI generates an average of two 3D-image data sets (5-10GB) per day 	<ul style="list-style-type: none"> • Data must be sliced to fit the desktop's graphics card for analysis and visualization. • Segmentation and annotation must be done slice by slice • Data must be transferred to collaborators or these need to travel to where the data is analyzed in order to provide their input for comprehension and interpretation of the data. 	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> • No use of wide area network at this point mainly for lack of a remote access, collaborative infrastructure.
2-5 years	<ul style="list-style-type: none"> • Within the next 2-5 years there will be an explosion of 3D images • Data acquisition will be more centralized as high-end microscopes are becoming more costly to purchase and maintain • A centralized repository will be needed to store the 3D images and an infrastructure will be needed for accessing and analyzing the remote data 	<ul style="list-style-type: none"> • Users interact with data as illustrated in Figure 1 • Data is centralized and long-distance collaborations take place. 	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> • With a central repository in place, bandwidth requirements will increase to 10 Gbps between users and repository • Low latency is also required.

5+ years	<ul style="list-style-type: none"> • About 2-25TB/day of image data will be generated by different centers from all over the world and stored in various repositories that will keep mirror images of the data • A backup system should be in place. 	<ul style="list-style-type: none"> • Users perform remote analysis and interactive visualizations of the data as well as remote interactive collaborations. 	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> • With the addition of other repositories in Europe and Asia, and the need for a backup system, bandwidth requirements will increase to 15 Gbps for Repository-Repository connections and to 160 Gbps for Repository-Backup System connections.
----------	--	--	---	---

12 BioEnergy Science Center

Department of Energy Bioenergy Research Center
Oak Ridge National Laboratory (headquarters), Oak Ridge, TN

Background

The DOE Bioenergy Research Centers are focused on reducing the US dependence on foreign sources of fuel, enhancing national energy security, and decreasing the harmful environmental effects of traditional energy use. To develop a national biofuel industry will require transformational science involving interdisciplinary research communities across the US and abroad. As the program moves forward, new research communities are emerging as are knowledgebases and new computational resources. Linking together computations as well as human interactions through robust high-speed networks is one critical piece of achieving these goals. In the near term, great progress can be achieved by introducing ESnet Collaboration Services between the DOE bioenergy research centers and within each bioenergy research center. It is the latter, networking within each bioenergy research center that will be the most challenging (see section 10 for details of members of one bioenergy research center - GLBRC).

Key Local Science Drivers

Instruments and Facilities

The Computational Biology facilities managed by the Computational Biology and Bioinformatics Group include those associated with the BioEnergy Science Center, the GTL Shewanella Federation, and the Joint Genome Institute (JGI). These facilities consist of servers developed and maintained by the group along with access to the National Center for Computational Sciences. All of the servers are housed in a 40,000 ft² state-of-the-art data center. The systems include clusters with over 500 processors, 800 GB of memory and 14 TB of storage. A new system is being delivered this month that includes 1056 cores in 66 nodes with 64Gb memory per node and 8TB storage per compute node. There is a knowledgebase system runs over several instances of a 3-tier architecture that includes a MySQL relational database management system, Java Beans and Servlets, and XHTML and JavaScript delivered by an Apache web server. Also there is a Laboratory Information Management System (LIMS) running on Nautilus, a commercial software package that uses an Oracle relational database management system. The investigators also have access to the National Center for Computational Sciences (NCCS) that provides resources for high-performance computational science and computing science research, and offers storage resources as well as computing, networking, and visualization resources.

The Microbiology Laboratory has facilities available for work with aerobic and anaerobic microorganisms including all of the required incubators for the proposed studies. The Molecular Laboratory has modern equipment for the use and development of techniques including Polymerase Chain Reaction (PCR), DNA hybridization analysis, gene cloning, gene deletion, and related techniques. The laboratory has an extensive supply of molecular enzymes and reagents, along with refrigerators and freezers for storage. A

variety of controlled temperature growth chambers, water baths and incubators are available as well as autoclaves, and freeze-drying facilities. Equipment includes 20 PCR thermocyclers, an iCycler Thermal Cycler (Bio-Rad) for real-time quantitative PCR, an AKTA FPLC (Amersham Pharmacia Biotech) for protein purification, a denatured gradient gel electrophoresis system (Bio-Rad), several high-speed centrifuges, numerous microcentrifuges, an electroporator, and a UV crosslinker for DNA hybridization studies.

The Genomics Laboratory is equipped with several automated liquid handling robots pulse field gel electrophoresis unit (CHEF) (Bio-Rad), a two automatic capillary DNA sequencers (ABI 3730xl), and a Hydra HTS workstation (Robbins Scientific) for 96-well sample dispensing. The Microarray facility contains three robotic microarray printers including a Cartesian Technologies printer (model PixSys 5500), a Microgrid II robotic arrayer (GeneMachines), and an Omnigridd300 (GeneMachines). A ScanArray® 5000 Microarray Analysis System (Perkin Elmer) and a microarray scanner with two micron resolution upgrade is available for high-density microarray analysis (Agilent Technologies Inc.). A Nimblegen microarray hybridization, an Illumina GAIIX and HiSeq DNS sequencing instruments, a 454 DNA sequencing instrument (Roche) and necessary support equipment have recently been acquired.

Process of Science

Scientists use the instruments and facilities in a rather unpredictable way. There is little production like activities from which a repeatable process can be defined. A few of those production activities involve routine DNA and RNA sequencing. In this case, data is generated and reduced locally. The reduced data is in the form of base calls and quality scores and is what is typically transferred across the WAN. Downloading nightly updates at the NCBI short read archive (SRA) of the base calls and quality scores generally completes within an hour.

Key Remote Science Drivers

Instruments and Facilities

The three DOE bioenergy centers are:

- DOE BioEnergy Science Center (BESC) led by DOE's Oak Ridge National Laboratory in Oak Ridge, Tennessee. This center will focus on the resistance of plant fiber to breakdown into sugars and is studying the potential energy crops poplar and switchgrass.
- DOE Great Lakes BioEnergy Research Center (GLBRC) led by the University of Wisconsin in Madison, Wisconsin, in close collaboration with Michigan State University in East Lansing, Michigan. This center is studying a range of plants and, in addition to exploring plant fiber breakdown, aims to increase plant production of starches and oils, which are more easily converted to fuels. This Center also has a major focus on sustainability, examining the environmental and socioeconomic implications of moving to a biofuels economy.

- DOE Joint BioEnergy Institute (JBEI) led by DOE's Lawrence Berkeley National Laboratory. This center will concentrate on "model" crops of rice and *Arabidopsis*, in the search for game-changing breakthroughs in basic science, and is exploring microbial-based synthesis of fuels beyond ethanol.

Considering the network within one center, BESC, you have multiple institutional partners (see Figure 4, below). A Similar network of organizations and collaborators exist at JBEI and GLBRC.



Figure 4. BESC represents 18 institutions and over 300 people working together to solve the problems of converting biomass into fuel.

BESC Partners

- **DOE's Oak Ridge National Laboratory (ORNL), Oak Ridge, Tennessee (lead institution):** DOE's largest science and energy laboratory, ORNL features research programs in poplar genomics, computational science, bioenergy, and plant and microbial systems biology. The ORNL Spallation Neutron Source and supercomputers at the ORNL National Leadership Computing Facility will be used to investigate and simulate the activity of enzyme complexes.

- **University of Georgia, Athens (UGA):** UGA's Complex Carbohydrate Research Center maintains state-of-the-art capabilities in mass spectrometry, nuclear magnetic resonance spectroscopy, chemical and enzymatic synthesis, computer modeling, cell and molecular biology, and immunocytochemistry for studying the structures of complex carbohydrates and the genes and pathways controlling plant cell-wall biosynthesis.
- **DOE's National Renewable Energy Laboratory (NREL), Golden, Colorado:** NREL has more than 30 years of experience in biomass and biofuel research and houses premiere facilities for analyzing biomass surfaces. NREL also has a long and successful history of establishing biofuel pilot plants and partnering with industry for commercial development of technologies.
- **University of Tennessee, Knoxville (UT):** UT conducts successful programs in bioenergy-crop genetic and field research (particularly switchgrass) and biotechnological applications of environmental microbiology.
- **Dartmouth College, Hanover, New Hampshire:** Dartmouth's Thayer School of Engineering is a leader in the fundamental engineering of microbial cellulose utilization and consolidated bioprocessing approaches.
- **Georgia Institute of Technology, Atlanta:** Georgia Tech's Institute for Paper Science and Technology provides BESC with expertise in biomass processing and instrumentation for high-resolution analysis of plant cell walls.
- **ArborGen, Summerville, South Carolina:** ArborGen provides expertise in forest genetics research, tree development, and commercialization.
- **Verenium Corporation, Cambridge, Massachusetts:** Verenium is a biofuels-focused biotechnology company and developer of specialty enzymes found in diverse natural environments and optimized for targeted applications.
- **Mascoma Corporation, Boston, Massachusetts:** Mascoma develops microbes and processes for economical conversion of cellulosic feedstocks into ethanol.
- **The Samuel Roberts Noble Foundation, Ardmore, Oklahoma:** This nonprofit research foundation is devoted to improving agricultural production and advancing the development of switchgrass and other grasses through genomic research. The foundation's activities are conducted through programs in agriculture, plant biology, and forage improvement.
- **Ceres, Inc., Thousand Oaks, California:** Ceres uses advanced plant breeding and biotechnology to develop and market low-carbon, nonfood crops for next-generation biofuels and biopower.
- **Individual Researchers.** Specializing in biomass pretreatment, characterization of plant-associated microbes, cellulose and enzyme modeling, consolidated bioprocessing, and lignin biochemistry are researchers from the University of California, Riverside; DOE Brookhaven National Laboratory (Upton, New York); Cornell University (Ithaca, New York); Virginia Polytechnic Institute and State University (Blacksburg); University of Minnesota (St. Paul); North Carolina State University (Raleigh); and Washington State University (Pullman).

Local Science Drivers – the next 2-5 years

Technology is changing; new sequencing machines are on the horizon. Micro-fluidic chips are capable of detecting DNA, development of the DOE Systems Biology

Knowledgebase will involve more tracking and transferring of data and information. It is hard to anticipate growth rates, but we know sequencing instruments and other instruments are evolving and continue to increase in data generation exponentially.

Instruments and Facilities

We would expect to go from a 1 Gbps LAN to 10 Gbps LAN in the coming years. We expect in the coming months to deploy a ~1000 core ~.5PB storage cluster in the open collaborative network. This system will provide for data intensive, data parallel computing as part of a growing DOE knowledgebase initiative.

Process of Science

In the coming years, as is happening today, certain experiments that generate data to be shared will become more predictable. Lab automation and advances in technology are fueling a trend in transforming parts of biological research into data production processes. The growth of systems biology and the 'omics' platforms will continue as will the need to share and analyze the data generated from these platforms. Still, a significant portion of the scientific investigation processes will remain unpredictable from a network requirements perspective.

Summary Table

Feature	Key Science Drivers		Anticipated Network Requirements	
Time Frame	Science Instruments and Facilities	Process of Science	Local Area Network Bandwidth and Services	Wide Area Network Bandwidth and Services
Near-term (0-2 years)	<ul style="list-style-type: none"> • BESC knowledge base • Genome annotation • Laboratory information management system • DNA sequencing 	<ul style="list-style-type: none"> • Sequence data is generated at JGI in CA, transferred to ORNL (some sequence generated locally as well) annotated and deposited in Genbank and sent back to CA. • Laboratory information systems track various experimental data and results throughout the BESC sites, stored at ORNL, and rendered is summarized form. • Knowledge base captures annotation, laboratory information and external information, combines with analyses and tools, and is delivered to the community. • Sequence data is generated at JGI in CA, transferred to ORNL (some sequence generated locally as well) annotated and deposited in Genbank and sent back to CA. • Laboratory information systems track various experimental data and results throughout the BESC sites, stored at ORNL, and rendered is summarized form. • Knowledge base captures annotation, laboratory information and external information, combines with analyses and tools, and is delivered to the community. 	<ul style="list-style-type: none"> • 1-10Gbps LAN • Data archival services 	<ul style="list-style-type: none"> • 1-10 Gbps WAN

2-5 years	<ul style="list-style-type: none"> • Continuation of near term requirements • Expansion of 'omics related data generation and analysis • Continued improvements to instruments and facilities 	<ul style="list-style-type: none"> • Continuation of near-term scientific processes 	<ul style="list-style-type: none"> • 10-100Gbps LAN • Data archival services 	<ul style="list-style-type: none"> • 10-100 Gbps WAN
5+ years	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> •

13 NOAA's Climate Modeling Program

Background

NOAA's Climate modeling program has seen a recent influx of funding (ARRA) which is allowing NOAA to expand into a much larger computational requirement. To make the most effective use of this change NOAA will move to a remote compute model. This is essentially only possible because NOAA has invested in a high performance research network that will provide the connectivity between its primary compute facilities, storage facilities and its users. Additionally NOAA's scientists continue to expand their collaborations with other agencies that provide an incentive to expand NOAA's network peering relationships with other agency and research networks.

Key Local Science Drivers

Instruments and Facilities

NOAA's primary climate modeling facility is located in Princeton, NJ. This 10,000 sq ft. facility currently supports about 8000 cores of SGI Altix 4700 for batch compute, 4.5PB disk supporting scratch and archive files systems and approximately 25PB current tape storage. The LAN currently supports multiple 10Gbs connections. Currently the WAN network connections support multiple 1Gbs connections to ESnet and Internet2.

Process of Science

The NOAA climate modeling process has traditionally been to submit a batch job to conduct the cyclical compute of decadal, centennial and millennial type experiments. The batch output is then analyzed. The average data growth rate from batch to analysis is 4x. Once complete the entire experiment is placed on tape storage.

Key Remote Science Drivers

Instruments and Facilities

In FY10 NOAA has been awarded INCITE grants of about 120 million computational hours on large scale computational systems at ORNL and ANL. Additionally, NOAA has contracted with ORNL to provide the Climate Modeling Research System (CRMS) a large scale computational system for NOAA's use which will be online by October 2010. Within the next 6-8 months NOAA will complete the installation of their new research network which will support 2x10Gbs connections into the network core from ORNL and 4x10Gbs connections from the Princeton facility into the network core.

Process of Science

NOAA's climate modeling process will stay the same with the exception of moving the model output data from the remote site back to Princeton. The data movement will be primarily handled via GridFTP and will be integrated into the Flexible Modeling System runtime environment. Systems are being sized to support 80TB per day from the CMRS and up to 10 TB per day from other remote systems to support INCITE hours.

Local Science Drivers – the next 2-5 years

Instruments and Facilities

NOAA will rely more on remote compute facilities as the cost for power increases and networks costs decrease. The facility at Princeton will move from a compute facility to perform just analysis and storage. It is expected that the tape storage requirements could exceed 500PB in the next 3-5 years.

Process of Science

NOAA's climate modeling process will have to evolve however this has a big dependency on the system architecture and the capability of the analysis node. The I/O requirement for the analysis node is much higher than the batch nodes. As the scientific analysis process evolves and the network bandwidth is less expensive the scientists may not need to move as much data and therefore reduce the data transfer rates.

Remote Science Drivers – the next 2-5 years

Instruments and Facilities

NOAA's climate modeling program will continue to expand its computational requirements with additional facility that will come online in early FY12. NOAA's research network components will support many 10Gbs links and are capable of upgrades to support 100Gbs if the requirements change.

Process of Science

Same as above.

Beyond 5 years – future needs and scientific direction

NOAA currently is estimating a 12x increase in its research high performance compute capability in the next 18months and traditionally NOAA has seen a 2x increase every 2 years so we need to plan for a 48x increase in capability which translates to a 24x in data growth within the next 5 years. If this growth is realized, NOAA will have to upgrade to 100Gbs to support this demand if the system architecture remains the same.

Middleware Tools and Services

NOAA will use GridFTP as the primary data movement tool utilizing other Globus tools. The data movement and overall resource allocation and management will be handled using MOAB. The majority of the infrastructure for this larger process already exists however it will need to be integrated together to achieve the most efficient use of the resources. NOAA also expects to have implemented HSPD 12 (CAC card access)

Summary Table

Feature	Key Science Drivers		Anticipated Network Requirements	
	Science Instruments and Facilities	Process of Science	Local Area Network Bandwidth and Services	Wide Area Network Bandwidth and Services
Near-term (0-2 years)	<ul style="list-style-type: none"> • INCITE hours • NOAA CMRS at ORNL 	<ul style="list-style-type: none"> • Large scale compute • Data transfer 	<ul style="list-style-type: none"> • Multiple 10 Gbps connections 	<ul style="list-style-type: none"> • 1-10 Gbps connection and IP peering at multiple locations (2) • 2x10 Gbps connections to ORNL (Chicago and Atlanta)
2-5 years	<ul style="list-style-type: none"> • INCITE Hours • NOAA CMRS II at ORNL 	<ul style="list-style-type: none"> • Large scale compute • Data transfer 	<ul style="list-style-type: none"> • Multiple 10 Gbps connections 	<ul style="list-style-type: none"> • 10 Gbps connection and IP peering at multiple locations (2) • 4x10 Gbps connections to ORNL (Chicago and Atlanta) • Possibly 100 Gbps peering with ORNL
5+ years	<ul style="list-style-type: none"> • INCITE Hours • NOAA CMRS II at ORNL follow on contracts 	<ul style="list-style-type: none"> • Large scale compute • Data transfer • Evaluation in how the science workflow is managed 	<ul style="list-style-type: none"> • 	<ul style="list-style-type: none"> •

14 Atmospheric Radiation Measurement (ARM) Climate Research Facility

Background

The Atmospheric Radiation Measurement (ARM) Climate Research Facility is a long-term measurement facility funded by CESD of BER in DOE that focuses on measuring:

- Cloud properties: microphysics (phases of water), optical properties, and patterns of occurrence
- Aerosol properties: size, chemistry, optical, and generation and decay pathways
- Cloud and aerosol interactions: absorption of aerosols by clouds and cloud formation triggered by aerosols
- Sunlight energy “fate”: radiative flux transfer, heating rate profiles, components of reflected and absorbed radiant energy, direct and diffuse light
- Atmospheric state: profiles of temperature, water vapor, wind, and aerosols
- Surface properties: Surface fluxes, soil conditions

The ARM climate research facility is building a “climatology” (multi-year record) of these measurements that are related to cloud formation, sunlight energy fate, and aerosol formation / decay, and aerosol interactions with clouds. The measurements are used to improve parameters used to represent these processes in global circulation models (GCMs). The GCM models are used for the prediction of future climate patterns. Parameters for cloud formation, sunlight energy fate, and aerosol interactions are thought to be the source of the largest uncertainties in these models and long-term climate forecasts.

ARM field sites are located in Oklahoma / Kansas, North Slope of Alaska, Tropical Western Pacific (Manus and Nauru Islands and Darwin Australia). The ARM mobile facility is currently located in the Azores, and soon to be in India. Formerly it was in China, Germany, Niger, Africa, and Coastal California. A second, more modular, Mobile facility is nearing completion and will be initially operated in Colorado. This second Mobile facility is designed to be ship-based, if needed. As a user facility, ARM regularly has field campaigns collocated with existing sites that involve collaborations with the entire atmospheric community and their instruments. The program also has ties with NOAA, NASA, and ECMWF.

Field data systems are located at each of the field and mobile sites. Data systems with facility-wide functions are located at PNNL, BNL, and ORNL. A distinct aspect of the ARM data collection is that it is continuous and has essentially the same parameters for its entire history. Most other studies in these aspects of atmospheric science include only short-term case studies of only a few weeks or months.

The users of the ARM data and network resources include: ARM facility personnel (for initial data collection, internal transfer, processing, and storage) and the research community (for access / download / use of documentation and data). The user community is mostly in the US, but is also globally distributed. Users can be divided into

the following categories: working at DOE facilities; not working at DOE but within the US; located at universities; and persons from foreign countries.

Key Local Science Drivers

Instruments and Facilities

Each of the instrumented sites has a local computer system with several terabytes of storage to handle buffering and onsite review of data. Most of the instruments (e.g., radiometers, meteorological sensors, aerosol samplers, etc.) are relatively small data producers (<10MB/day). However, other instruments, such as 3D scanning radars or lidars, produce about 15GB/day per unit. ARM will soon be operating 28 radar systems. The local network connects the site data system with data loggers and instrument computers and facilitates maintaining instrument uptime and quality and the continuous data collection.

The ARM data systems at PNNL, BNL, and ORNL are each involved in the processing functions that create or distribute higher orders of data and each lab has many terabytes of storage to manage the data sets. These systems use the “local” network infrastructure at their DOE lab. Within each of the data systems, some use of private networks occurs to link together multiple system performing similar functions. The ARM Archive at ORNL shares the HPSS Mass Storage System with the super computers at ORNL and has access to significant storage resources.

Process of Science

The use of local networks is dominated by monitoring of instruments, data collectors, data processing, data storage, and data distribution. Quality review of the data products and processes also use the local networks at the ARM sites. A variety of operations, scientific (instrument and quality experts) and systems personnel use the local network. Researchers have very limited and infrequent access to the local networks and this access is typically during field campaigns with a limited duration.

Key Remote Science Drivers

Instruments and Facilities

Because of the globally dispersed nature of the ARM sites and the goal of a continuous record of measurements, the Internet is a critical component for accomplishing the ARM mission. Each site is minimally staffed and has offsite monitoring and maintenance of computer systems. ARM has a large dispersed infrastructure team monitoring all aspects of data quality and system components. Each site uses a local ISP to connect to the global Internet. VPNs are implemented between each measurement site and ANL and the ARM Data Management Facility (DMF) at PNNL.

The ARM infrastructure at ANL provides VPN tunnels to each of the measurement facilities and supports the following services:

- Global and ARM Infrastructure remote access to the measurement sites with access controls;

- Secondary (hidden), recursive arm.gov DNS service accessible to measurement sites;
- Scheduled and random security scans of measurement facilities;
- Measurement Facility Syslog archive;
- VOIP support among measurement facilities and limited access to US POTS lines;
- Measurement Facility traffic monitoring (Snort IDS and related tools);
- User level VPN access to measurement facilities coming near future;
- Secondary ARM Program-wide LDAP service coming in the near future; and
- Measurement Site network device configuration management.

The following documents the current connection to the Internet for each measurement facility and possible bandwidth upgrades, should operations funding become available:

- ARM Southern Great Plains (SGP) site near Lamont, OK – Currently 6Mbps through Pioneer Long Distance Telephone Company. This summer the link will be moved to a 100Mbps connection through Oklahoma OneNet. ARM anticipates that over 120GB/day will be transferred to the DMF over this link.
- ARM North Slope of Alaska (NSA) Barrow, AK site: T1 connection (via ATT satellite) to Internet2 at University of Alaska, Fairbanks. Infrastructure capacity available to upgrade to 3xT1.
- ARM NSA Atqasuk, AK site: 64Kbps through GCI (Alaska North Slope communications company).
- Manus Island, Paupa New Guinea: Satellite Link with IntelsatGeneral – (U.S. Ground station in Riverside, California). 256Kbps/64Kbps (uplink/downlink relative to site). FY10 will install a second ground station to support an off-site radar. The new ground station and the Nauru measurement facility share the satellite bandwidth. ARM infrastructure is available to upgrade the link to 576/192Kbps or 1024/256Kbps (the former more likely considering costs).
- Republic of Nauru: Shares the satellite bandwidth with Manus Island.
- Darwin, AU site: 2Mbps through Telstra Corporation. Potential to upgrade to 10Mbps (symmetrical).
- ARM Mobile Facility: Currently located on Graciosa Island, Azores, Portugal. Using 2Mbps link through PT Comunicacoes (Portuguese government telecommunications agency). Will be re-located to India FY11. Plan on a minimum of 2Mbps link.

Each measurement site implements a VPN with the DMF for hourly data and metadata transfer. The DMF at PNNL provides centralized access for first order data evaluation by the ARM infrastructure. Infrastructure access to hourly updates of data at the DMF helps

ensure optimal data quality and minimizes data gaps. The DMF supports the following ARM-wide services:

- Receives raw measurements and meta-data from measurement sites;
- Performs the “ingest” of raw data (i.e. convert raw data into a standardized NetCDF format for ease of use by ARM Facility users);
- Hosts the ARM Data Quality processing systems;
- Hosts arm.gov DNS and arm-wide e-mail lists;
- Hosts science.arm.gov which provides for ARM user logins and scientific collaboration via shared file resources, Wiki collaboration, and other services;
- Hosts measurement facility-wide local- and wide-area network capacity monitoring;
- Hosts measurement facility-wide compute systems capacity monitoring;
- Manages the reliability, timeliness and completeness of all ARM data streams.
- Receives data/meta-data storage media from measurement sites (large data streams which cannot be delivered by Internet) for ingest and subsequent transfer to ARM Archive at ORNL; and
- Hosts engineering services to design, develop and evolve the ARM data system and data flow processes.

Within the next year, the DMF will be receiving >100GB/day from the sites over the Internet. Because of bandwidth limitations to most of the remote sites (e.g., Alaska and oceanic islands), up to 400GB/day will be sent to the DMF via transportable hard drives.

The DMF sends all of the site data and derivative products to the ARM Archive at ORNL. This has been less than 50GB/day but is expected to grow to over 500 GB/day in the next year. Further growth to 1+ TB/day will occur over the next few years as the secondary data products are implemented or the new ARRA funded instruments.

The eXternal Data Center (XDC), at BNL, manages the receipt of non-ARM instrument data of interest to ARM user community. These include field campaigns and regular data sets produced by other groups (like satellite data). The data volume varies but is on the order of 50GB/day. The XDC also hosts the following ARM-wide services:

- Acquire and ingest external data relevant to ARM measurements (satellite data, climatological data from other agencies, etc);
- Transfer of external data to ARM Archive;
- Reprocesses external data with dedicated system located at ARM Archive at ORNL;
- Host the <http://www.arm.gov> website which is tightly integrated with <http://www.archive.arm.gov> and requires synching of databases between Archive and XDC;

- Hosts secondary arm.gov DNS domain; and
- Hosts arm-wide LDAP service.

The Archive at ORNL maintains the long-term storage for all ARM data and distributes it to scientists who request the data. Requests for data are submitted by way of a variety of web-based user interfaces. These user interfaces include web applications ranging from simple forms to interactive graphical displays. The current volume of stored data at the Archive is 200TB with 1PB expected within 5 years. Data requested for downloads (via FTP) are currently 1.5 - 5.5 TB per month (150-750 K files).

In addition to data storage and distribution, the ARM Reprocessing Center is hosted at the ARM Archive. Large numbers of data files (100's K) are transferred between the Archive and the Reprocessing Center on the local network. Remote access to the Reprocessing Center is enabled for ARM staff who perform application software installation, review processing results and control processing flow. The ORNL login server facilitates this access.

Process of Science

Scientists have typically downloaded data sets relevant to their research from the Archive and performed analysis at their local institution. This paradigm is augmented by ARM's production of Value Added Products (VAPs) that include data values based on higher order integration, analysis, or derivation. VAPs can result in smaller data sets being downloaded. However, this may not be true for the 3D radar data. The 3D gridded products may be larger than the original radial rays of collected data.

ARM has several developments that may reduce the data volumes to be transferred over the network. Researchers often only need a small portion of a data set for their work, and ARM is developing tools to extract and subset data in ways that make small portions of the large sets of ARM data more easily accessible to scientists. A recent survey of ARM scientists showed that a majority would have reservations about downloading more than 100GB of data per task.

Local Science Drivers – the next 2-5 years

Instruments and Facilities

In FY 2009, ARM received \$60 million in capital equipment funds from ARRA to expand its instrumentation and improve its infrastructure. This funding specifically includes the purchase of radar systems that will generate at least 10X the data than was previously being collected. After the ARRA funded equipment is implemented, no significant additional instrumentation is expected for the next several years.

The fundamentally larger volume of data will increase the capacity demand for local networks within both the field sites and the ARM data centers at PNL, BNL, and ORNL. Except for throughput through firewalls, when needed, the capacity of the local networks should be sufficient. The new ARRA instrumentation, especially at SGP, Darwin, Manus and the AMF, is expected to increase the frequency and volume of global remote access to these sites:

- ARM (and external) radar scientists remotely interacting with the on-site radars to observe and tune scanning algorithms. This applies to other new instruments also.
- Remote monitoring of instrumentation in support of, likely, increased number of short-term field studies involving guest instrumentation.

Process of Science

The primary change in the scientific process associated with the much large data volumes will be the relocation of part of the processing from the user's systems to the ARM systems. More preliminary processing for data extraction, summarization and integration, for preliminary data visualization, and preliminary data analysis is likely to be performed on ARM's systems (primarily at the Archive). The primary network implications are the need for reduced latency, the secure transport of authentication processes, and the prevention of untoward activities (accidental or malicious). Many of these data processing instructions will need to originate from users on foreign networks. Acquiring user specifications for data extraction via "finite" user interfaces is relatively straightforward and constraining the results is manageable. However, specifications for more complex summarization and data integration functions are very inefficient to formulate in a user interface and the processing requests generated outside of a user interface are more difficult to constrain or review. Accommodating interactive graphics via the Internet will require limited authentication into limited command domains or proxy processes.

Remote Science Drivers – the next 2-5 years

Instruments and Facilities

In conjunction with the implementation of the new 3D radars funded by ARRA, ARM is also working to provide resources for very large data tasks (10's TB) so that users can perform their work on a system adjacent to the Archive at ORNL. The implementation of large libraries of pre-computed visualizations (data plots and animations) will also be developed. If access and performance issues can be resolved, interactive visualization may also be implemented. While many of these developments may reduce the volume of data flow, the expectations for reduced latency and increased reliability of the network may increase (particularly to sustain the communications needed for interactive analytical or visualizations of very large data.)

Process of Science

Many of new, very high volume instruments being implemented by ARM do not have any usage history for the data products to be generated. Except for a very small number of researchers (a few 10's at most), the usage of these very large data products will be very exploratory (working with small selections, reviewing visualizations, etc.). As the user community better understands the scientific value of these data, their vision for data products to be generated and data products to be transferred to them or analyzed by them will grow rapidly. Between more download usage and more combinations of data products (statistical or visuals), the data flow volume on the network is likely to increase

3-5X within the next 5 years. The exact growth is difficult to predict and depends largely on where the user processing is conducted (on their system or on ARM's systems). Historically, researchers have been much more likely to use data in their processing environment. The overhead to learn the processing environment of numerous other data centers is too high. Also, security issues are frequently more complex (of necessity). The next few years may continue to follow historical patterns of data use (primarily download) or may have new patterns of remote processing. The usage of very large-scale data centers and their analytical tools is also likely to lead to more Internet-based transfer of data (e.g., Earth System Grid; See ESG Use Case for more details).

Beyond 5 years – future needs and scientific direction

Plans for facilities beyond 5 years are difficult to predict for ARM. Many “never before operated continuously” implementations are just beginning. The needs beyond 5 years will be better known after the next 2-3 years. It will take that long for the infrastructure and the research community to develop a next generation view that extends beyond the very large collection of new instruments currently being implemented. A minimal and likely scenario is a continuation of the exponential increase with a significant jump up in the trend due to the new high data volume field instrumentation and secondary data products.

Significant effort is being invested into finding more frequent and better ways to conduct climate research with more joint usage of observational data and simulation results. The observations and simulations will not only continue to become spatially and temporally intense, they are more likely to focus on “short-term” (decades) and regional (sub-continental) analyses. As this occurs, intensive scanning of observations and simulations (individually and combined) to find instances of critical impacts is likely to be common. The impact assessment is more likely to access data from numerous locations with data integration and use patterns that are much less predictable.

The development of data access and analytical tools that focus on simulation results and integrated measurement observations is a final factor that will affect the future of data transfer for climate research. Earth System Grid (ESG) is specific example of this type of tool. ESG will have the potential to use many TB of data (simulated and observed results) for each analytical task. The processing conducted by ESG will be distributed across the Internet (see more details in ESG use case).

Middleware Tools and Services

The program currently uses ftp for data transfer among ARM facilities. Since the program has control over the ftp end-points, the program can effectively take advantage of available bandwidth through tuning of TCP windows and other parameters and performing parallel transfers. However, at some point, the program will be required to initiate transfer connections using secure authentication techniques. The actual transfer process, however, need not be secure (i.e., no need to use compute resources to encrypt/de-encrypt the actual data transfer).

Network-based interactions needed that can enable remote access to processes that have restricted “write” actions by users that are temporarily authenticated by program criteria.

For example, these are needed for interactive visualization or statistical processes that are commonly used during initial data exploration. Procedures currently used that include distributing secure ID tokens and passwords are too slow and too tedious for initial data exploration. Global solutions developed and tested by ESnet that can provide this capability would be very helpful. Allocating sufficient resources for this purpose is difficult even for projects the scale of ARM.

Scientific programs like ARM that are dependent on extensive use of the Internet often opt to use DOE lab-based security plans and firewalls rather invest significant resources for independent security plans and firewalls. When the data flows for scientific programs like ARM become extreme, then the performance and management of security devices like firewalls are very important. Identifying continually improving security hardware and associated policies that can be commonly adopted across DOE for handling large volumes of internal and external data transfers is a valuable contribution by ESnet. This enables the research support programs to remain focused on particular needs of the scientists. Security enclaves may provide benefits and are under discussion and limited implementation. However, implications on performance and configuration management are not yet known.

Summary Table

Feature	Key Science Drivers		Anticipated Network Requirements	
	Science Instruments and Facilities	Process of Science	Local Area Network Bandwidth and Services	Wide Area Network Bandwidth and Services
Near-term (0-2 years)	<ul style="list-style-type: none"> Anticipate little instrument expansion at measurement facilities after the new ARRA instruments are installed this FY 	<ul style="list-style-type: none"> Anticipate increased on-site field campaigns as a result of new ARRA instruments. This implies increase in remote access to both ARM and visitor instrumentation at measurement sites 	<ul style="list-style-type: none"> Upgrade measurement facilities to 1Gbps to support higher bandwidth instruments. Increased compute power at measurement facilities to create “interesting” subsets of very large data streams that are capable of being shipped over available network bandwidth. 	<ul style="list-style-type: none"> Increase network capacity from remote measurement facilities to ISP as possible given operations budget
2-5 years	<ul style="list-style-type: none"> Very few additions are expected. This time period will be used to complete the secondary implementation of the currently ongoing changes. 	<ul style="list-style-type: none"> Visualization and analysis of very large 3D data products maybe executed by the scientists. 	<ul style="list-style-type: none"> Demands are likely to continue to increase as secondary products are developed for the very large 3D data sources. Upgrade local networks to 10 Gbps. 	<ul style="list-style-type: none"> Anticipate increase in volume of data flow from measurement sites to DMF as bandwidth costs decrease. The complexity of required network protocols are likely to increase.
5+ years	<ul style="list-style-type: none"> Changes are likely to continue increases for network capacity. These are very difficult to predict because of the magnitude of changes that are just now being started. 	<ul style="list-style-type: none"> Inter-comparison and assimilation between 3D simulations and measurements are common. 	<ul style="list-style-type: none"> Continued increasing demands for network capacity. 	<ul style="list-style-type: none"> Expanded use of very large transfers between data centers. Expanded use of analysis and visualization software control by users from remote networks.

15 Metagenomics, Argonne National Laboratory

Background

The Metagenomics program at ANL is responsible for analysis of genome sequence data produced by local sequencing facilities and remote users. The data sets are growing ~5-10X annually, with users often requiring remote analysis of collected data sets.

Key Local Science Drivers

Instruments and Facilities

Local sequencing facilities produce datasets at the rate of 10 GB/day per machine. Typically, single runs take several days to complete, with local buffering of results. Data sets are moved (via commodity protocols like http, ftp, scp, etc) to relatively local compute resources.

Process of Science

Samples are sequenced and the results are moved to relatively local computational resources. Analysis is performed, returning the results to the user. Results are considerably smaller than the input data sets.

Key Remote Science Drivers

Instruments and Facilities

Some sequencing is performed at the JGI, DOE's major sequencing facility. In these cases, data may be moved a longer distance before analysis.

Process of Science

Same basic process, but different computation location.

Local Science Drivers – the next 2-5 years

Instruments and Facilities

We expect the output of sequencers to continue to grow at the rate of approximately 10X annually for the next several years. The cost of sequencing systems is currently ~\$500K per machine, however, new machines are being developed that will be considerably cheaper (~\$50K per machine). We expect this to change the distribution of sequencing machines, from being largely centralized today to a largely decentralized layout, with any lab being able to procure inexpensive sequencing machines. We expect this to change the patterns of data movement substantially, particularly as users become unable to locally analyze sequence data due to growing computational costs. (Analysis is growing with data set size, causing a shortfall in local analysis capabilities).

Process of Science

Unchanged from above.

Remote Science Drivers – the next 2-5 years

Instruments and Facilities

The limitations of local analysis are already prompting a scramble for cost effective analysis techniques. Expect to see some adoption of DOE computation centers (NERSC, *LCF) for analysis in the next few years. Moreover, since most of the analysis performed in metagenomics consists of “pleasantly parallel” execution, existing commercial cloud offerings are already practical to use, and will probably be cost effective in the future.

Process of Science

With luck, collaborative techniques (data set co-analysis) will become more popular. These techniques aren’t currently used widely, but will likely be required for future data volumes.

Middleware Tools and Services

Currently, commodity protocols are used, typically by hand. This will need to change to cope with the growing data set sizes and data movement requirements.

Outstanding Issues

Genomics is in a major state of flux; the capabilities of instruments and their associated data projection are growing very quickly, and analysis techniques are rapidly changing as well. As a field, bioinformatics has not built organized collaboration infrastructure (unlike say, high energy physics) to this point.

Summary Table

Feature	Key Science Drivers		Anticipated Network Requirements	
	Science Instruments and Facilities	Process of Science	Local Area Network Bandwidth and Services	Wide Area Network Bandwidth and Services
Near-term (0-2 years)	<ul style="list-style-type: none"> JGI sequencing machines Site sequencing machines 	<ul style="list-style-type: none"> Data analysis in the same location as data production, movement of smaller data sets 	<ul style="list-style-type: none"> Data sets in the 100-1000 GB range 	<ul style="list-style-type: none"> Movement of 20-100 data sets a week
2-5 years	<ul style="list-style-type: none"> Large influx of distributed sequencing capabilities (machines in labs) 	<ul style="list-style-type: none"> Data will need to be transported to analysis (conducted at centers or in clouds) 	<ul style="list-style-type: none"> Data sets in the 1-100 TB range 	<ul style="list-style-type: none"> Data sets in the 1-100 TB range (moderate growth of volume due to commoditization of sequencing)
5+ years	<ul style="list-style-type: none"> 	<ul style="list-style-type: none"> 	<ul style="list-style-type: none"> Data sets in the 100TB-1PB size 	<ul style="list-style-type: none"> Data sets in 100TB-1PB size (same number of sets)

16 Findings

The following issues were reported and discussed at the workshop.

Common Themes

As in other areas of science, many BER programs have data sets that are growing exponentially. BER scientists are also typical in that a number of them are having trouble fully utilizing the network to handle their data transfer requirements, and transfers often take longer than expected based on available network capacities.

There is likely to be a significant increase in the use of commercial clouds by BER-funded scientists over the next few years, as a number of case studies mentioned this growing need. In particular the Genomics community is looking into using commercial cloud services. ESnet has already upgraded its peerings at the commercial exchanges that carry cloud computing traffic to 10Gbps (10Gbps in each of three geographically diverse locations). ESnet will track the traffic load on these peerings, and upgrade them as necessary to support science.

Another common theme was that the ARM project, EMSL, and others discussed the problems of cross-site authentication and authorization, and suggested that ESnet lead an effort to solve this problem. It is difficult for scientific collaborations to navigate the overlapping and sometimes mutually exclusive security requirements of the various National Laboratories.

Participants also pointed out the need for middleware that would transfer a specified group of files from system A to system B with *minimal* oversight from the researcher. For example, this middleware should continue to transfer the files regardless of access interruptions due to network, system shutdown / restarts, and so on. Several projects expressed the opinion that that the effectiveness of data transfer middleware was not just on the transfer speed, but also the time and interruption to other work required to supervise and check on the success of large data transfers.

Several participants expressed the need for videoconferencing capabilities such as those provided by the ESnet ECS service.

Climate Science

The biggest BER user of ESnet in the future is likely to be the climate community. LLNL and ORNL will be hosting huge data sets that will be accessed all over the world. By 2020, climate data at these sites will exceed hundreds of exabytes. The key sites for distributing this data are the CMIP-5 Data Centers at LLNL, ORNL, NCAR, BADC (U.K.), Germany, Australia, and Japan. Data will be replicated across a subset of these locations, and bandwidth requirements will likely reach 100 Gbps in the next few years. ESnet needs to carefully monitor the utilization of these links, and work with the DOE Office of Science to secure funding for upgrading the connections to these sites before they get overloaded.

NOAA is likely to be a significant source of data for climate researchers. In addition, if NOAA receives an INCITE allocation at the ALCF, there will be significant data flows from the ALCF to NOAA storage systems. ESnet and NOAA are currently collaborating to meet this need, and multiple high-bandwidth peerings will soon be established between ESnet and the NOAA network.

JGCRI anticipates increasing use of supercomputing resources at ORNL, with data export to PNNL and UMD. This will result in significant data flows from ORNL to the other two sites. In addition, JGCRI has a need to distribute significant amounts of data between the PNNL and UMD sites – an OSCARS circuit and performance tuning is likely to be needed.

ARM: As the use of ARM data increases and new ARRA funded instruments come online, it is likely that there will be additional performance tuning effort required to ensure efficient use of ARM data by the science community. As an example, the ARM case study points out that “A recent survey of ARM scientists showed that a majority would have reservations about downloading more than 100GB of data per task.” This points to a need for performance tuning assistance for ARM data users, as well as better data management middleware.

Genomics

The cost of genome sequencers is falling dramatically, and sequencers are continuing to increase in data output exponentially (one estimate is for a 12x increase in data volume over the next 5 years). Sequencers will no doubt become more common and exist at more locations instead of being centralized at one facility. However, data management for these instruments will likely become a difficult issue for many institutions (e.g. it might be easy for a small university to get a grant to buy a sequencer, but analyzing the resultant huge volume of data will prove difficult). Several outcomes are possible, and all are likely to happen to at least some degree. These include the use of commercial clouds for genomics data analysis (this is already happening), the transfer of genomics data to centers such as JGI for analysis, and the development of local analysis infrastructures at institutions that purchase genome sequencers as they become affordable. The prospect of significant data inflows will drive changes at JGI, and also drive changes at many sites that use JGI.

A need was expressed by the Genomics community for good connectivity to the National Center for Biotechnology Information (NCBI) in Bethesda, MD. ESnet currently has 10Gbps connectivity to the MAX Gigapop, where the National Library of Medicine and the National Institutes of Health are both present. ESnet will monitor its connection to MAX and upgrade it if needed.

Biosciences

Data transfer is a bottleneck for LANL B6. LANL B6 would like a standardized method for file transfer “implemented across ESnet.” This could be implemented by deploying a dedicated data transfer host at B6 and at the other genomics sites (NERSC already has dedicated Data Transfer Nodes, for example). LANL site networking personnel are currently working on installing dedicated data transfer hosts for general site use – hopefully these hosts could be used by LANL B6.

Bioenergy

JBEI has significant need for more efficient distribution of large visualization data sets. Currently, researchers travel to JBEI because they are unable to use the network for their visualization needs. The “data crunch” at JBEI will increase significantly as data production processes are automated, resulting in terabytes of data per day that must be analyzed by remote scientists. In addition, JBEI is planning to build a distributed data repository at LBNL that is likely to generate significant bandwidth (over 10Gbps) when fully deployed.

17 Requirements Summary and Conclusions

A number of common themes emerged from the case studies and workshop discussions. One is that BER science, like many other disciplines, is becoming more and more distributed and collaborative in nature. Another common theme is that data set sizes are exploding. Climate Science in particular is on the verge of needing to manage exabytes of data, and Genomics is on the verge of a huge paradigm shift in the number of sites with sequencers and the amount of sequencer data being generated.

ESnet will continue to closely monitor the needs of these communities.

Action Items

Several action items for ESnet came out of this workshop. These include:

- Create a “climatedata” email list for discussion of network performance issues for climate (done)
- Create a “biodata” email list for discussion of network performance issues for biology and genomics (done)
- Work with ESG and other members of the climate community on immediate performance tuning and troubleshooting needs in support of data distribution for climate science. This effort is ongoing, and has already had significant success.
- ESnet will continue to develop and update the fasterdata.es.net site as a resource for the community
- Work with PNNL to establish a Data Transfer Node for the PNNL site. This process has begun.
- ESnet will continue to assist sites with perfSONAR deployments and will continue to assist sites with network and system performance tuning

In addition, ESnet will continue development and deployment of the ESnet On-demand Secure Circuits and Advance Reservation System (OSCARS) to support virtual circuit services on the Science Data Network.

18 Glossary

Gb: GigaBase – one billion base pairs

GB/sec: Gigabytes per second – a measure of network bandwidth or data throughput

Gbps: Gigabits per second – a measure of network bandwidth or data throughput

GigaBase: One billion base pairs – a measure of genomics data

Illumina: Illumina is a company that builds and sells genome sequencing equipment and analysis tools.

IPCC: Intergovernmental Panel on Climate Change. See <http://www.ipcc.ch/>

KB/sec: Kilobytes per second – a measure of network bandwidth or data throughput

Kbps: Kilobits per second – a measure of network bandwidth or data throughput

KiloBase: One thousand base pairs – a measure of genome base pairs sequenced

MB/sec: Megabytes per second – a measure of network bandwidth or data throughput

Mbps: Megabits per second – a measure of network bandwidth or data throughput

MegaBase: One million base pairs – a measure of genomics data

PB/sec: Petabytes per second – a measure of network bandwidth or data throughput

Pbps: Petabits per second – a measure of network bandwidth or data throughput

TB/sec: Terabytes per second – a measure of network bandwidth or data throughput

Tbps: Terabits per second – a measure of network bandwidth or data throughput

19 Acknowledgements

This work would not have been possible without the contributions and participation of those who provided information and attended the workshop. ESnet would also like to thank the BER program office for their help in organizing the workshop and providing insight into the facilities supported by the BER program. In addition, the LBNL conference support and logistics staff was very helpful.

ESnet is funded by the US Department of Energy, Office of Science, Office of Advanced Scientific Computing Research (ASCR). Vince Dattoria is the ESnet Program Manager.

ESnet is operated by Lawrence Berkeley National Laboratory, which is operated by the University of California for the US Department of Energy under contract DE-AC02-05CH11231.

This work was supported by the Directors of the Office of Science, Office of Advanced Scientific Computing Research, Facilities Division, and the Office of Biological and Environmental Research.

This is LBNL report LBNL-4089E