







# ESnet

ENERGY SCIENCES NETWORK

## Advanced Scientific Computing Research Network Requirements Review

Final Report

April 22-23, 2015



# 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.

# **Advanced Scientific Computing Research Network Requirements Review Final Report**

Office of Advanced Scientific Computing Research, DOE Office of Science  
Energy Sciences Network (ESnet)  
Germantown, Maryland  
April 22–23, 2015

ESnet is funded by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research. 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 U.S. 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 and Facilities Divisions.

This is LBNL report LBNL-1005790.

# Contents

<b>Participants and Contributors</b>	<b>5</b>
<b>Executive Summary</b>	<b>6</b>
<b>Findings</b>	<b>8</b>
<b>Action Items</b>	<b>10</b>
<b>Review Background and Structure</b>	<b>11</b>
<b>Office of Advanced Scientific Computing Research Overview</b>	<b>13</b>
<b>Case Studies—ASCR Facilities</b>	<b>17</b>
1 <b>Argonne Leadership Computing Facility</b>	<b>17</b>
2 <b>National Energy Research Scientific Computing Center</b>	<b>27</b>
3 <b>The Oak Ridge Leadership Computing Facility</b>	<b>38</b>
<b>Case Studies—ASCR Research</b>	<b>49</b>
4 <b>Advanced Network Services</b>	<b>49</b>
5 <b>Bulk Data Transfer</b>	<b>64</b>
6 <b>Remote Analysis and Visualization Services</b>	<b>70</b>
7 <b>Remote Analysis and Visualization at Sandia National Laboratories</b>	<b>77</b>
8 <b>Workload Management Systems</b>	<b>80</b>
9 <b>Streaming Workflows: Fusion Experimental Data Processing Workflow</b>	<b>87</b>
10 <b>Scientific Workflows</b>	<b>91</b>
<b>References</b>	<b>103</b>



# Participants and Contributors

Katie Antypas, LBNL, NERSC Facility  
Greg Bell, ESnet, Networking, ESnet Director  
Wes Bethel, LBNL, Remote Analysis and Visualization  
Jody Crisp, ORISE, Logistics  
Rich Carlson, DOE/SC/ASCR, ASCR Research Program, Next Generation Networking  
Eli Dart, ESnet, Networking, Review Chair  
Vince Dattoria, DOE/SC/ASCR, ESnet Program Manager  
Kaushik De, University of Texas at Arlington, Data and Job Management  
Ian Foster, ANL, Scientific Workflows  
Barbara Helland, DOE/SC/ASCR, ASCR Facilities Division  
Mary Hester, ESnet, Networking  
Scott Klasky, ORNL, Streaming Data Transfers, Storage System I/O  
Alexei Klimentov, CERN, Data and Job Management  
Tom Lehman, University of Maryland, Advanced Network Services (SDN)  
Miron Livny, University of Wisconsin, Data and Job Management  
Joe Metzger, ESnet, Networking  
Ryan Milner, ANL, ALCF Facility  
Ken Moreland, SNL/NM, Remote Analysis and Visualization  
Lucy Nowell, DOE/SC/ASCR, ASCR Research Program, Computer Science  
Daniel Pelfrey, ORNL, OLCF Facility  
Adrian Pope, ANL, ALCF Facility  
Prabhat, LBNL, NERSC Facility  
Rob Ross, ANL, Storage System I/O  
Lauren Rotman, ESnet, Networking  
Brian Tierney, ESnet, Networking  
Jack Wells, ORNL, OLCF Facility  
John Wu, LBNL, Streaming Data Transfers, Data and Job Management  
Wenji Wu, FNAL, Bulk Data Movement  
Ben Yoo, University of California at Davis, Advanced Networking Services (SDN)  
Dantong Yu, BNL, Bulk Data Transfer  
Jason Zurawski, ESnet, Networking

## Report Editors

Eli Dart, ESnet: [dart@es.net](mailto:dart@es.net)  
Mary Hester, ESnet: [mchester@es.net](mailto:mchester@es.net)  
Jason Zurawski, ESnet: [zurawski@es.net](mailto:zurawski@es.net)

# Executive Summary

The Energy Sciences Network (ESnet) is the primary provider of network connectivity for the U.S. Department of Energy Office of Science (SC), 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 25 years.

In April 2015, ESnet and the Office of Advanced Scientific Computing Research (ASCR), of the DOE Office of Science, organized a review to characterize the networking requirements of the programs funded by the ASCR program office.

Several key findings highlighting the results from the review are noted below.

1. There was a discussion of achieving and documenting a recommended method and expected performance level for large-scale data transfers between the ASCR computing facilities over ESnet. Representatives from the ASCR computing facilities agreed in principle to collaborate together and with ESnet on establishing and documenting methods for users to routinely achieve high levels of data transfer performance between the ASCR computing facilities.
2. As the ASCR computing facilities deploy systems with burst buffer capabilities, it is likely that users will also want to stream data from a burst buffer to a remote location without going through the parallel filesystems.
3. There is a need for stable, functional, hardened, production-quality workflow tools that can easily be supported in production at the ASCR computing facilities. A small number of well-understood, easily-supported workflow packages will be much easier for the facilities to support than one home-grown workflow tool per experiment type. This is an area where a modest amount of focused effort in the near term could bring about significant strategic benefits to all the ASCR computing facilities and the scientists they support.
4. As *in situ* analysis becomes more common, users will push for interactive analysis, so they can steer the analysis and perhaps the simulation. The ability to steer the analysis is important because otherwise *in situ* analysis requires *a priori* knowledge of what the user wants to discover or analyze.
5. Based on current technology and projections, the traditional parallel filesystem may be nearing the end of its growth trajectory in terms of performance and scale. Other storage technologies, such as key/value stores and object stores, are likely to be deployed. The impact of these new technologies on data transfer tools and technologies (e.g. Data Transfer Nodes), and therefore on ESnet, is unknown at this time, but may be significant.
6. The plans for available storage (e.g., filesystem capacity, filesystem I/O performance) at HPC facilities in the exascale era are lower in relation to the scale and performance of the computational systems than has been the case historically. This is likely to have implications for ESnet as user behavior changes in response to scarce storage resources. One possibility is additional data movement as data sets are transferred to other facilities instead of being stored next to the exascale system. Another possibility is increased network load resulting from data transfers directly from high-performance burst buffers rather than from a traditional parallel filesystem which would be much slower in comparison to the burst buffers.
7. There is an open research topic involving the determination of the correct set of interfaces, service abstractions, and metrics to allow for the productive interaction between networks, workflow systems, computing resources, and storage resources.

This report expands on these points, and addresses others as well. The report contains a findings section in addition to the text of the case studies discussed during the review.

# Findings

Below are the findings for the ASCR and ESnet Requirements Review held April 22–23, 2015. These points summarize important information gathered during the review.

- There was a discussion of achieving and documenting a recommended method and expected performance level for large-scale data transfers between the ASCR computing facilities over ESnet. Representatives from the ASCR computing facilities agreed in principle to collaborate together and with ESnet on establishing and documenting methods for users to routinely achieve high levels of data transfer performance between the ASCR computing facilities.
- The collaboration between the LHC/ATLAS experiment and the OLCF has the potential to significantly increase the amount of network traffic at ORNL. The ATLAS experiment, which uses the PanDA workload management system to schedule and run simulation and analysis jobs, is collaborating with the OLCF to run simulation jobs on the Titan system. The simulation workload is expected to result in about 13 Gbps of steady-state traffic load at ORNL for a day at a time. The data transfers will be between ORNL and multiple ATLAS data centers, including the ATLAS Tier-1 center at BNL.
- The poor performance of tape archives sometimes causes ASCR HPC facility users to transfer data to a different facility, or to resources at their home institution. Reading the data back from tape is sometimes slow enough that data written to tape are not analyzed.
- Many projects would benefit from increasing the data transfer performance to tape archives (e.g., HPSS) via Globus. Currently, data transfers to filesystems perform at a higher throughput than data transfers to tape archives. This means users must choose between a single transfer directly to tape (much slower performance), or transferring data to the filesystem then transferring the data to tape (which is a human-intensive two-stage workflow).
- The ASC facilities often encounter performance problems when transferring data sets to or from the ASCR computing facilities. Many of the performance problems encountered are due to security restrictions—high-performance tools such as Globus are not allowed at the ASC facilities, so legacy tools (e.g., rsync over ssh) must be used for data transfer.
- Statistics from the Globus transfer service show that universities that built Science DMZs (i.e., with funding from the NSF CC-NIE program) have moved over 3PB of data using Globus over the past 4 years. Because of the degree of external collaboration between ESnet sites and other research institutions (more than 80% of data transfers across ESnet have one endpoint outside the national laboratories), the increase in data capabilities at U.S. universities are likely to result in increased data transfer volumes across ESnet to the ASCR computing facilities. Historical trends and technology projections lead to an estimate of 25% year over year growth in data transferred by Globus—it is likely that the ASCR computing facilities will experience significant growth in data transfers in the coming years.
- The Data Transfer Nodes for the Theta system, due to be deployed at ALCF in 2016, will have 40GE interfaces.
- ANL may need an additional 100G connection to ESnet in the 2016 time frame to support increased DTN load at the ALCF.
- Cosmology workflows need data transfer performance at the level of 100TB/day between the ASCR computing facilities, and between the ASCR computing facilities and project archive sites. In addition, for moving

simulation output to archival storage, a performance level of 1PB/week is required. Therefore a data transfer performance level of 15 Gbps or 1.7GB/sec between facility DTNs is needed to meet both requirements collectively.

- As the ASCR computing facilities deploy systems with burst buffer capabilities, it is expected that some experiment workflows will transfer data directly to a node attached to the burst buffer from an external source, bypassing the parallel filesystem.
- There is a need for stable, functional, hardened, production-quality workflow tools that can easily be supported in production at the ASCR computing facilities. A small number of well-understood, easily-supported workflow packages will be much easier for the facilities to support than one home-grown workflow tool per experiment type. This is an area where a modest amount of focused effort in the near term could bring about significant strategic benefits to all the ASCR computing facilities and the scientists they support.
- The LHC experiments will see very large increases in data volume in the next 5–7 years (the ALICE experiment will see a 100x increase in data volume during LHC Run 3, and the ATLAS and CMS experiments will see a 100x increase in data volume during LHC Run 4). Since the LHC experiments run data analysis jobs at the ASCR computing facilities, these data increases will have an impact on the ASCR computing facilities in addition to ESnet.
- A cadence is emerging for some projects at the Leadership Computing Facilities. During the first year, the project spends its time getting the code to scale and doing large-scale runs at the LCF. After that, a second year is spent analyzing the data from the first-year runs. In many cases, support for the second year of data analysis is *ad hoc*, and comes from discretionary allocation resources.
- As *in situ* analysis becomes more common, users will push for interactive analysis, so they can steer the analysis and perhaps the simulation. The ability to steer the analysis is important because otherwise *in situ* analysis requires *a priori* knowledge of what the user wants to discover or analyze.
- Based on current technology and projections, the traditional parallel filesystem may be nearing the end of its growth trajectory in terms of performance and scale. Other storage technologies, such as key/value stores and object stores, are likely to be deployed. The impact of these new technologies on data transfer tools and technologies (e.g. Data Transfer Nodes), and therefore on ESnet, is unknown at this time, but may be significant.
- Data transfer performance metrics (e.g. data transfer performance between Data Transfer Nodes) which PanDA collects for itself, could be useful to a variety of workflow systems, schedulers, and other tools.
- Workflow management systems need better network performance metrics to help guide decisions. This is a potential area of collaboration for ESnet.
- The plans for available storage (e.g. filesystem capacity, filesystem I/O performance) at HPC facilities in the exascale era are lower in relation to the scale and performance of the computational systems than has been the case historically. This is likely to have implications for ESnet as user behavior changes in response to scarce storage resources. One possibility is additional data movement as data sets are transferred to other facilities instead of being stored next to the exascale system. Another possibility is increased network load resulting from data transfers directly from high-performance burst buffers rather than from a traditional parallel filesystem which would be much slower in comparison to the burst buffers.
- There is an open research topic involving the determination of the correct set of interfaces, service abstractions, and metrics to allow for the productive interaction between networks, workflow systems, computing resources, and storage resources.
- There is currently no commonly-deployed and easy-to-use software toolkit/API for high-performance streaming I/O between geographically distant systems. Multiple projects could benefit from the creation of such a software toolkit.
- The ATLAS experiment would find it useful if ESnet could give PanDA access to a bandwidth allocation which PanDA could then sub-allocate to scheduled data placement tasks.



# Action Items

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

- ESnet will explore the creation of a “superfacility engineering team” composed of network engineers, software tool developers, network researchers, and HPC system experts. This team would address the collection and sharing of performance metrics to enable the next generation of workflow and co-scheduling systems.
- ESnet will explore collaboration with elements of the Exascale program which are looking at the behavior of networking inside of exascale machines.
- ESnet will coordinate with the PanDA project and the OLCF on the ATLAS simulation pilot at the OLCF.
- ESnet will collaborate more closely with the Open Science Grid on network performance issues.

# ESnet SC Requirements Review Background and Structure

Funded by the Office of Advanced Scientific Computing Research (ASCR) Facilities Division, ESnet's mission is to operate and maintain a network dedicated to accelerating science discovery. ESnet's mission covers three areas:

1. Working with the DOE SC-funded science community to identify the networking implications of instruments and supercomputers and the evolving process of how science is done.
2. Developing an approach to building a network environment to enable the distributed aspects of SC science and to continuously reassess and update the approach as new requirements become clear.
3. Continuing to anticipate future network capabilities to meet new science requirements with an active program of R&D and advanced development.

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

*(a)* A review of major stakeholders' plans and processes, including the data characteristics of scientific instruments and facilities, in order to investigate what data will be generated by instruments and supercomputers coming online over the next 5–10 years. In addition, the future process of science must be examined: How and where will the new data be analyzed and used? How will the process of doing science change over the next 5–10 years?

*(b)* Observing current and historical network traffic patterns to determine how trends in network patterns predict future network needs.

The primary mechanism to accomplish (a) is through the SC Network Requirements Reviews, which are organized by ASCR in collaboration with the SC Program Offices. SC conducts two requirements reviews per year, in a cycle that assesses requirements for each of the six program offices every three years. The review reports are published at <http://www.es.net/requirements/>.

The other role of requirements reviews is to help ensure that ESnet and ASCR have a common understanding of the issues that face ESnet and the solutions that it undertakes.

In April 2015, ESnet organized a review in collaboration with the ASCR Program Office to characterize the networking requirements for the facilities and science programs funded by ASCR.

Participants were asked to communicate and document their requirements in a case-study format that included a network-centric narrative describing the science, instruments, and facilities currently used or anticipated for future programs; the network services needed; and how the network is used. Participants considered three timescales on the topics enumerated below: the near-term (immediately and up to two years in the future); the medium-term (two to five years in the future); and the long-term (greater than five years in the future).

More specifically, the structure of a case study was as follows:

- Background—an overview description of the site, facility, or collaboration described in the case study.
- Collaborators—a list or description of key collaborators for the science or facility described in the case study (the list need not be exhaustive).

- Network and Data Architecture—description of the network and/or data architecture for the science or facility. This is meant to understand how data moves in and out of the facility or laboratory focusing on local infrastructure configuration, bandwidth speed(s), hardware, etc.
- Instruments and Facilities—a description of the network, compute, instruments, and storage resources used for the science collaboration/program/project, or a description of the resources made available to the facility users, or resources that users deploy at the facility.
- Process of Science—a description of the way the instruments and facilities are used for knowledge discovery. Examples might include workflows, data analysis, data reduction, integration of experimental data with simulation data, etc.
- Remote Science Activities—a description of any remote instruments or collaborations, and how this work does or may have an impact on your network traffic.
- Software Infrastructure—a discussion focused on the software used in daily activities of the scientific process including tools that are used to locally or remotely to manage data resources, facilitate the transfer of data sets from or to remote collaborators, or process the raw results into final and intermediate formats.
- Cloud Services—discussion around how cloud services may be used for data analysis, data storage, computing, or other purposes.

The case studies included an open-ended section asking for any unresolved issues, comments or concerns to catch all remaining requirements that may be addressed by ESnet.

# Office of Advanced Scientific Computing

## Research Overview

The mission of the Advanced Scientific Computing Research (ASCR) program office is to fund basic science in applied mathematics, networking, and computer science; deliver the most advanced computational scientific applications in partnership with disciplinary science; advance computing and networking capabilities; and develop future generations of computing hardware and tools for science. All this is done in partnership with the academic, industrial, and laboratory research communities. The strategy to accomplish this mission has two thrusts: a basic research program in applied mathematics, computer science and advanced networking; and developing and maintaining world-class computing and network facilities for science.

U.S. private- and public-sector organizations are increasingly using supercomputers to achieve breakthroughs of major scientific or economic importance. These achievements have already advanced U. S. competitiveness and were, in many cases, accomplished through access to very powerful supercomputers and High Performance Computing (HPC) experts at the DOE national laboratories using tools developed with support from ASCR. ASCR has a strong history of supporting innovative scientific computing. Researchers using ASCR facilities have:

- made discoveries in functional materials
- made fundamental studies of turbulence in chemically reacting systems
- made fundamental studies of climate change
- made fundamental studies in the understanding of the physical properties of matter, such as the quark-gluon nature of nuclear matter
- modeled 3D full-core reactor neutron transport to predict the behavior of novel nuclear fuels in fission reactors
- conducted 3D turbulent combustion simulations of hydrocarbons to increase fuel efficiency
- made U.S. airplane engines quieter, more fuel efficient, and less polluting
- made long haul trucks more energy efficient in record time
- simulated ice formation in million-molecule water droplets to reduce the wind turbine downtime in cold climates
- identified novel materials for use in extreme energy environments.

The Office of Science, through ASCR, and the National Nuclear Security Administration (NNSA), have partnered to make strategic investments in hardware, methods, and critical technologies to address the exascale technical challenges and deliver an exascale system. Such a system will help scientists harness the thousand-fold increase in capability to address critical research challenges and will maintain U.S. competitiveness in HPC. These efforts are linked with investments to advance data-intensive science and to effectively use the massive scientific data generated by DOE's unparalleled suite of scientific user facilities and large-scale collaborations. By investing in both next-generation computing and data-intensive science, the ASCR program will enable the community of HPC users to

- improve and shorten industrial design processes

- design advanced materials
- better understand dark matter and dark energy
- explore possibilities for dramatically increasing fuel efficiency while lowering emissions
- design advanced nuclear reactors that are modular, safe, and affordable
- improve accuracy of climate predictions
- predict and investigate how to control the behavior of fusion plasmas
- calculate the subatomic interactions that determine nuclear structure.

## **Mathematical, Computational, Computer Sciences, and Networking Research**

Experiments at several of SC's user facilities, such as the light and neutron sources, and experiments at the Large Hadron Collider (LHC), are migrating towards workflows that need near real-time interaction between instruments and simulations. Experiments and simulations are often deeply intertwined as simulations become necessary in the design of large-scale experiments, and data from experiments are analyzed in simulations to inform and guide further experiments. The volume and complexity of data generated have increased such that a focused effort is required to develop theories, tools, and technologies to manage data—from generation through integration, transformation, analysis, and visualization, including collaborative environments; to capture the historic record of the data; and to archive and share it. This request supports ASCR efforts in data-intensive science for collaborations with applied mathematicians and computer scientists to address end-to-end data management challenges and develop new scientific workflows.

Software, tools, and methods from core research efforts will be used by the Scientific Discovery through Advanced Computing (SciDAC) partnerships to more effectively use the current and immediate next generation high performance computing facilities.

## **High Performance Computing and Network Facilities**

The Leadership Computing Facilities (LCFs) will continue preparations for a planned 75–200 petaFLOPS upgrades at each site in the 2018–2019 timeframe. These upgrades represent technological advances in both hardware and software, and engineering efforts for the ASCR facilities that incorporate custom features to meet the Department's mission requirements. HPC and the high-performance network facilities will also expand efforts in exascale component technology research and development, system engineering, and integration, leading to the design and development of future HPC systems including prototype test beds for demonstrating the feasibility of building exascale systems, and the exascale systems themselves.

The National Energy Research Scientific Computing Center (NERSC) takes delivery of the NERSC-8 supercomputer in 2016, which will expand the capacity of the facility by 10–40 petaFLOPS to address emerging scientific needs.

Experienced computational scientists who assist a wide range of users in taking effective advantage of the advanced computing resources are critical assets at both the LCFs and NERSC. To address this DOE mission need, support continues for a post-doctoral training program for high end computational science and engineering. In addition, the two LCFs and NERSC will continue coordinating efforts to quantify scientist's computational requirements and prepare their users for future architectures.

ESnet is the DOE's high-speed science network engineered and optimized to support large-scale scientific research. ESnet interconnects and allows scientists to use DOE's unique research facilities independent of time and location with state-of-the-art performance levels by providing direct connections to more than 40 DOE sites and now offers international connections to CERN. With the 340 gigabit per second (Gbps) expansion to support SC's collaborations in Europe complete, the ESnet will continue to explore, in coordination with the International



Research and Education Network community, next generation optical networking technologies and global networking architectures for future upgrades. The outcomes of these efforts will help ESnet keep pace with the continuing growth of scientific traffic from DOE's scientific user facilities and experiments.

## **Case Studies—ASCR Facilities**

# Case Study 1

## Argonne Leadership Computing Facility

### 1.1 Background

The Argonne Leadership Computing Facility (ALCF) located on the Argonne National Laboratory campus outside of Chicago, Illinois, provides the computational science community with a world-class computing capability dedicated to breakthrough science and engineering. It began operation in 2006, with its team providing expertise and assistance to support user projects to achieve top performance of applications and to maximize benefits from the use of ALCF resources.

Awardees of compute time on the ALCF systems range from national laboratories and universities, to corporations and international collaborators. As such, data must be transferred to and from the facility, driving a range of networking requirements for the facility.

### 1.2 Network and Data Architecture

The ALCF has constructed various networks to meet the needs of high performance computing resources, evaluation and experimentation resources, and researchers.

#### Egress Connectivity

The ALCF is physically connected into the Argonne National Laboratory network using a Brocade MLXe-32 core router configured with 10x10GbE bonded interfaces to a Brocade MLXe-16 upstream router serving the data-center. This router in turn connects into the campus core at 100GbE. The core routers provide access to other campus resources and to external networks at a minimum of 10GbE.

The ESnet peering is primarily reached via a 100GbE path to transport at 710 N. Lakeshore Drive in Chicago. This peering can also be reached via two alternate 10GbE paths. One is available directly by a path to transport at 710 N. Lakeshore Drive, and the other via Metropolitan Research and Education Network (MREN).

#### Internal Architecture

A Brocade MLXe-32 core router is central to the ALCF architecture. Using a hub and spoke model for simplicity and reliability, this router is responsible for all layer-3 connectivity between resources such as the data transfer nodes (DTNs), the visualization cluster, and other experiments like the Petrel high-speed data store. Resources, such as the DTNs, that are directly connected to the router use a number of 10GbE bonded interfaces.

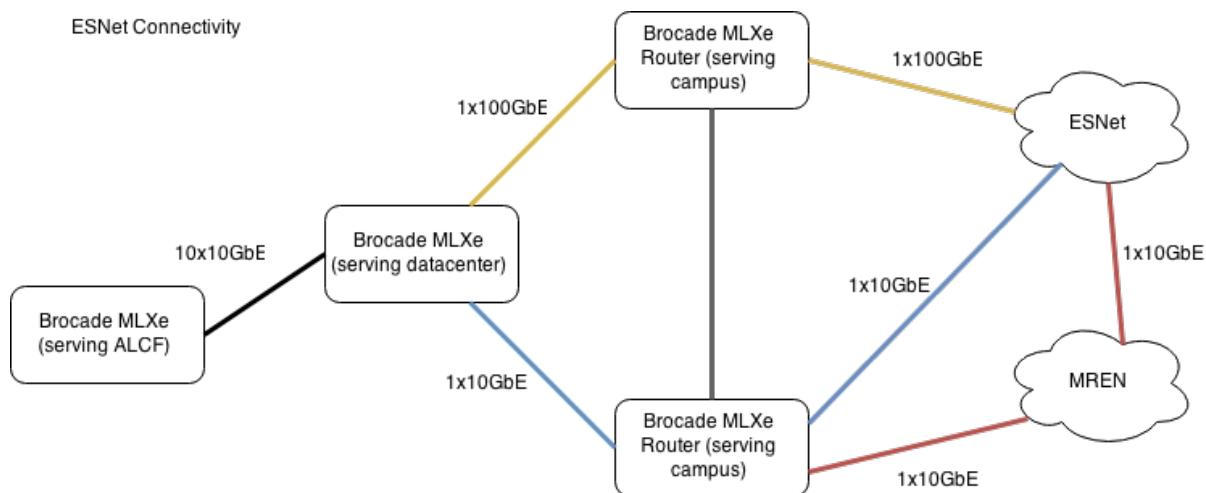


Figure 1.1: ALCF Ethernet egress topology

Supporting Brocade chassis are used for dense 10GbE aggregation and distribution using layer 2. For less bandwidth intensive access, Brocade FCX stackable switches are used. These switches will be connected into the core router using a number of 10GbE bonded interfaces.

ALCF utilizes two facilities on campus, one for housing HPC resources and the other to host experiments, such as Petrel, and for disaster-recovery diversity in tape storage locations. Dense wavelength division multiplexing (DWDM) equipment is used to create dedicated connectivity between the sites. At this time, four 10G Ethernet waves and twelve 8G Fibre Channel waves are being multiplexed between facilities.

### 1.3 Collaborators

As one of two DOE Leadership Computing Facility centers in the nation for open science, the ALCF, supported by the DOE ASCR Program, provides the computational science community with world-class computing capabilities, expertise, and assistance to ensure that every project achieves top performance on its resources.

ALCF provides support through a uniquely collaborative approach where staff are partnered with users to assist with scaling, I/O, optimization, workflow management, and domain-specific algorithm development. ALCF also partners with experts within the extreme-scale HPC community, at Argonne and throughout the world, to ensure best practices and best technologies are applied at the facility and available for the users. Last year ALCF supported 1,432 DOE-defined users and engaged in more than 342 active projects from universities, national laboratories, and industry worldwide.

### 1.4 Instruments and Facilities

#### 1.4.1 Present

Mira is an IBM Blue Gene/Q system, consisting of 48 racks 786,432 processors, and 768 terabytes of memory with a theoretical peak performance of 10 PetaFLOP/s. For message passing interface (MPI) communication, the system has a proprietary 5D torus interconnect. As this is only used for internal system communication, it is not relevant to the discussion of LAN and WAN requirements.

To reach external resources, such as the storage subsystem, Mira is equipped with 384 input/output nodes (IONs) connected to a QDR InfiniBand (IB) fabric. With all IONs operating at full capacity, Mira has an aggregate theoretical bandwidth of 1.5 TB/sec. Connectivity is provided by a fully connected network of four Mellanox IS5600




Figure 1.2: ALCF internal Ethernet topology.



QDR IB switches. Each switch provides 324 ports of edge connectivity for hosts such as IONs and file servers, and 324 ports of core connectivity between switches. Total aggregate capacity of the fabric is over 5 TB/sec.

DDN provides the storage infrastructure, using the SFA12K-20E platform. In this platform, virtual machines running onboard the disk controller couplets act as file servers for the infrastructure. The rated performance of a single couplet is 15 GB/sec. With 16 couplets, the aggregate bandwidth from Mira to disk is 240 GB/sec. This file system is intended for performant codes.

A second file system comprised of 6 couplets is available for less intensive codes. The aggregate bandwidth from Mira to this disk is 90GB/sec.

To achieve a reasonable balance between excessive data sizes and excessive storage consumption, a High Performance Storage System (HPSS) data-mover cluster is used. This system manages data migration and caching for Mira storage, providing higher throughput for users with large amounts of data to be studied over longer time periods.

For testing and debugging, two smaller Blue Gene/Q systems are available. Cetus, a four rack system of 4096 compute nodes, is connected to the same storage fabric as Mira and allows debug runs to be performed against production file systems, speeding time to resolution for users troubleshooting a failure at scale. This system, has 32 IONs connected with QDR InfiniBand, giving a ratio of 128:1 compute nodes per I/O node.

Vesta, a 2048 node two-rack system, lives in a separate, isolated FDR InfiniBand fabric. This allows experimental configurations to be tested that may involve unstable codes not fit for use on the production resource. It has been equipped with a total of 64 IONs, yielding a lower ratio of only 32:1, significantly increasing the capability of data I/O experimentation on this system.

A new cluster named Cooley will provide visualization for Mira, and later Theta. This Cray CS400-AC based system has 126 nodes, using NVIDIA Kepler K80s, with 12GB memory per compute node, achieving 223 teraFLOP/s (CPU+GPU) of double-precision performance. For inter-process communication (IPC) and storage area network (SAN) access the cluster is using a Mellanox SX6512 FDR IB switch. This switch allows for full bisectional bandwidth from any node to any other node within the system, and is uplinked to the existing SAN serving Mira. Each node is also connected at 10 Gbps into Ethernet switches with twelve 10GbE ports reserved for uplinks, providing connectivity to the ALCF WAN routers at 100Gbps aggregated theoretically.

Petrel is a pilot service for data management that allows researchers to store and share large datasets with internal and external collaborators. The pilot is joint project with the ALCF and Globus. It consists of 32 file servers and 1.7 PB of usable GPFS storage. The backing storage is four DataDirect Networks (DDN) S2A9900 Storage Systems. The file servers also serve as GridFTP DTNs, with a 1GbE WAN connection per file server. GPFS traffic utilizes a dedicated 10Gb Clos network. The GPFS was benchmarked with I/O rate at 8301.27/6059.07MB/s for read/write using all 32 file-servers as clients. Single-client performance was benchmarked with I/O rate at 1215.48/736.42 MB/sec for read/write. Networking performance has been measured with nuttcp between Mira and Petrel at 26Gbps using all 32 file servers to a test Mira DTN with a 4x10GbE aggregate.

#### **1.4.2 Next 2-5 years**

Theta will arrive in 2016 to support Argonne with the transition from Mira to Aurora. It will be based on the Cray XC series, powered by Intel's second generation Xeon processors and Knights Landing Phi coprocessors. This system should have a peak performance of greater than 8.5 PetaFLOP/s. The management infrastructure supporting this system will see standardization around 40GbE network interface cards. Anticipated workload includes large-scale compute, data analysis, and visualization.

#### **1.4.3 Beyond 5 years**

Aurora, will use Intel's HPC scalable system framework to provide a peak performance of 180 PetaFLOP/s. The system will help ensure continued U.S. leadership in high-end computing for scientific research while also cementing the nation's position as global leader in the development of next-generation exascale computing systems. Aurora

is considered a pre-exascale system. The system architecture is intended to be performant for both compute-intensive and data-centric workloads. This would require the LCF network architecture to support the migration of very large data sets.

## 1.5 Process of Science

### 1.5.1 Present

Groups awarded time at the ALCF often transfer in their data sets to the facility at the beginning of their computing time. This comes in the January time frame for INCITE awards, and the July time frame for ASCR Leadership Computing Challenge (ALCC) awards. Groups also transfer out simulation results to their home facilities or to collaborators year-round. The Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program process, which allocates 60 percent of the available time on the machine, is also open to international proposals. It is not unusual to have a trans-Atlantic awardee.

To enable data movement, a GridFTP cluster is available for the transfer of data sets. During normal operations, 12 DTNs are available for users, each with 10 GbE of connectivity. These systems read from the GPFS SAN and are built to saturate their network interface cards. To facilitate troubleshooting network connectivity there is a perfSONAR server open to the public.

The following table, Table 1.1, describes the top twenty-five ALCF DTN transfer destinations, characterized here by the Autonomous System (AS) name, ranked by the aggregate data transferred over the period of a year.

Table 1.1: Top 25 DTN transfer destinations from ALCF over the last year.

Autonomous System	Aggregate transferred in PB
NSHE-NEVADANET - Nevada System of Higher Education, US	136.682
FNAL-AS - Fermi National Accelerator Laboratory (Fermilab), US	115.416
ESNET-WEST - ESnet, US	89.159
NERSC - National Energy Research Scientific Computing Center, US	54.553
ORNL-MSRNET - Oak Ridge National Laboratory, US	20.649
BNL-AS - Brookhaven National Laboratory, US	20.578
UTAH - University of Utah, US	15.314
NCSA-AS - National Center for Supercomputing Applications, US	11.316
LANL-INET-AS - Los Alamos National Laboratory, US	9.955
TACCNET - University of Texas at Austin, US	9.870
UTEXAS - University of Texas at Austin, US	8.621
UIUC - University of Illinois, US	7.823
NCAR-AS - University Corporation for Atmospheric Research, US	3.979
BROWN - Brown University, US	1.414
CMCS - Comcast Cable Communications, Inc., US	1.254
STANFORD - Stanford University, US	1.012
UTK - The University of Tennessee Health Science Center, US	0.865
U-CHICAGO-AS - University of Chicago, US	0.680
ESNET-EAST - ESnet, US	0.599
UCDAVIS-CORE - University of California at Davis, US	0.556
FR-REMIP2000 REMIP 2000 Autonomous System, FR	0.418
USC-AS - University of Southern California, US	0.248
PPPL-AS1 - Princeton Plasma Physics Laboratory, US	0.230
CEBAF - Continuous Electron Beam Accelerator Facility, US	0.166
IASTATE-AS - Iowa State University, US	0.114

The following two charts, Figure 1.3 and 1.4, illustrate bursts of traffic over the period of a year for the six most active projects, but also all projects in aggregate. Each data point is the sum of that day's data movement for that project. This is meant to bursts in transfer activity. Notably the top bursts are generated by the cosmology projects.




Figure 1.3: Inbound traffic to ALCF DTNs.




Figure 1.4: Outbound traffic from ALCF DTNs.

### 1.5.2 Next 2–5 years

Significant new compute systems, collaboration tools, and upgrades are planned for this time period.

Currently the capabilities of the visualization cluster and the DTNs have the potential to create contention for ALCF network egress connectivity. In two to five years, there will be additional large-scale systems with supporting infrastructure standardized on 40GbE network interface cards, with their attendant DTNs. Adding to the competition for limited networking resources, several visualization walls are planned. It is also possible that users might request the ability to perform analytics in real time with a companion visualization and analysis resource located on the WAN, furthering straining the limited bandwidth. High speed data stores, like Petrel, will also be made available.

All of these resources translate into increased utilization for networking resources.

### 1.5.3 Beyond 5 years

Exascale is expected to be achieved.

## 1.6 Remote Science Activities

### 1.6.1 Cosmology Simulation Activities

Current and future astronomical surveys are pushing the limits of statistical fidelity in making cosmological measurements, such as determining the spatial clustering of galaxies. Robust cosmological inference from these and other precision cosmology datasets requires an accurate, large-scale simulation and modeling capability to make maximal use of the information contained in multiple cosmological probes. By virtue of their size, complexity, and dynamic range, cosmological simulations targeted at these problems will fully stress state of the art super-computers for generations to come.

Simulations of structure formation in the Universe are integral in survey planning, characterizing error distributions, and for calculating the observable signatures to which the data will be compared. Galaxies form in highly over-dense regions of space but trace delicate structures that span distances of hundreds of millions of light-years, demanding dynamic ranges of roughly six orders of magnitude in spatial resolution and even more in mass resolution. Much like galaxy surveys themselves, the data products from cosmological simulations are rich and can be interrogated in various ways at several levels of data reduction and compression, requiring significant resources for data archiving, data access, and data distribution.

The central engine of such a simulation workflow is a cosmological N-body code, such as the Hardware/Hybrid Accelerated Cosmology Code (HACC), that can run at scale on a variety of systems within the DOE Office of Science computing complex (ALCF, OLCF, NERSC). The outputs of a simulation can be broadly grouped into three categories based on the levels of data reduction and transformation as well as likely usage patterns. The analyses associated with the simulations can be carried out in both in-situ and offline modes, with the possibility of the latter being co-scheduled with the cosmological simulation.

One example of Level I data would be an output of all the N-body particle information at a simulation time step. For a “hero” run that is using most of a current leadership-class HPC system, the particle information for a single time step can total between 10 and 100 TB and is written into multiple individual files within a file set, with typically tens to hundreds of files per file set. The number of files and mapping between compute nodes and files is tuned to maximize the filesystem write bandwidth on each system. The purpose of “hero” runs is to study details of structure formation that can only be accessed with the highest dynamic range simulations possible, but they are so expensive in terms of computing resources that they need to be augmented by campaigns of smaller simulations to study the effects of varying cosmological parameters.

There are a few regular, predictable patterns of analysis that take Level I data for some carefully chosen time steps and produce Level II outputs that are reduced in size by roughly an order of magnitude or more. These can be run in-situ while the simulation is running or shortly after an output time step file set is produced. One example of a Level II data set is a catalog of dark matter-dominated halos and sub-halos, which mark the highly over-dense regions of the universe where galaxies can form. Scientific goals drive the time frequency requirements for the various Level I to Level II reductions, so different output time steps may be run through different sets of reductions. The extreme dynamic range available in “hero” runs generally results in more output time steps being analyzed than for a typical campaign run, with some kinds of Level II analysis being run on up to a hundred output time steps. Level II data sets are still fantastically rich and there are potentially many ways that many users would want to further process them into Level III (catalog level) data sets and share them with yet more users. Typically, Level III data sets may be interacted with in real time.

The networking requirements of the cosmological simulation workflow result from the distribution of resource allocation and availability at DOE computing facilities and user access to data. An important consideration is

the availability of data-intensive computing resources (large-scale “analysis clusters”) where offline analysis can proceed in either asynchronous or co-scheduled modes.

Level II and Level III data products may be analyzed and re-processed by many users at various facilities, but the file sets are relatively-speaking, small, and generally not part of tightly-coupled workflows, so transferring them is not a primary concern. The challenge is in moving Level I data for large simulations. Allocations of core hours at DOE computing facilities are competitively awarded, but the capabilities and policies for disk and archive access and retention vary considerably. A project may want to move a Level I file set between facilities in order to retain it on disk for a longer period of time or to archive it in a different location, or to allow another user to run a new Level I data reduction where that user has an allocation of core hours for that analysis. These transfers will not generally be part of a tightly coupled workflow, either, so some latency can be tolerated, but a file set of 100TB should be transferable on the time scale of one day.

The potential size of a Level I file set should scale roughly with the available main memory of leadership-class HPC systems, as “hero” runs are often executed in a near memory-limited configuration. However, the rate at which Level I file sets that a project needs to transfer are produced could increase more rapidly as the floating point capabilities of leadership-class HPC systems are increasing more rapidly than their memory footprints.

Over the next five years, the estimated level of simulation activity will be driven by a number of important surveys such as DESI (Dark Energy Spectroscopic Instrument), LSST (Large Synoptic Survey Telescope), WFIRST (Wide Field Infrared Survey Telescope), CMB-S4 (Cosmic Microwave Background—Stage 4). Simulations will include campaigns such as the ongoing Mira-Titan Universe simulation suite, which covers 100 large simulations run over a period of three years, to massive individual runs including hydrodynamics and a variety of subgrid models (gas cooling, star formation, astrophysical feedback mechanisms, etc.).

The large individual runs will create on the order of 10 PB of Level I data per simulation in roughly a month of wall clock time. (This is a subset of the expected output of roughly 100 PB total in Level I from all simulations.) Given current policies at supercomputing centers, moving this data set to the project’s archival storage (a mixture of disk and tape) will have to be done on a period of roughly three months, which translates to a hard requirement of moving data at the rate of 1 PB/week.

Additionally, there will be large data sets at Level II and Level III, which will be moved across supercomputer centers and storage sites as well as to sites with sufficient local analysis capabilities. The total amount of data at this level is estimated to be 50 PB. This number is also large and reflects the fact that post-processing can yield many different outputs depending on the choice of modeling parameters, even for a single base simulation. Although individual data transfers here will not be as large as the one discussed above, they too can be at the 1 PB level. Given that these datasets will be transferred essentially as part of an analysis campaign, strategies for maximizing overall throughput will be essential. These will include pre-staging of transfers, co-scheduling of analysis with the transfers, and the ability to interact with a sufficient fraction of the data set during the analysis process—a loosely coupled workflow. Such a requirement translates to several TB/hour.

## 1.6.2 Interactive and Remote Large-Data Visualization and Analysis Services

Currently, data sets in domains such as cosmology, combustion and astrophysics, range in size from  $2048^3$  to  $10,240^3$ , with a single time step of data in the range of terabytes. Using a dedicated high-end visualization resource (e.g., ALCF’s Cooley cluster) to visualize and analyze the data, single or multiple instances of a visualization application can be used to remotely visualize different variables of a particular data set. The resulting images are then streamed over the wide area over to the application scientist’s location to a display cluster that drives a multi-tiled display wall. A typical configuration for this application includes 15 LCD displays, each 1920x1200 pixels, arranged in 5 columns by 3 rows (34 MegaPixels). Streaming the full resolution of a single tile at 30 frames per second would require 2.1Gbps, with the full 15 tiles requiring 31Gbps. At the same time, there is an increasing number of high-resolution displays being deployed by science teams with a display resolution in the range of 50 to 100 MegaPixels (around 50–100Gbps) to glean insights from the higher resolution, higher fidelity, and more complex physics generated by the computational campaigns. With growing data sizes, and larger displays with more available tiles, the bandwidth requirements are anticipated to continue to grow. The data movement in this case is a parallel M-to-N data movement consisting of multiple sources and destinations. A characteristic of this data movement is reliable (lossless), low latency, low jitter, and high bandwidth.



One challenge with the current data movement is that it requires network admins at the various sites to be involved in order to set up the network path involving multiple administrative domains. This tends to be a severe impediment to science. For deterministic end-to-end performance, a multi-domain network reservation system is needed to account for scheduling policies within a site, typically involving multiple domains, together with the OSCARS reservations for the ESNNet domain. Also, it will be necessary to co-ordinate the scheduling of the job transfers with the job schedulers at each site to better facilitate the end-to-end data movement. This again necessitates a co-scheduling mechanism involving various scheduling policies along the end-to-end path

In the next 2 to 5 years, data sizes from science campaigns are expected to grow. We expect the adoption of new and cheaper 3D display walls including the CAVE2 by science teams to better understand the higher resolution and more complex datasets. This is critical to better understand molecular structures among others. This significantly increases the required frame rate (usually 60Hz) for interactivity and further increases the bandwidth needs by a factor of two. We also expect to witness the advent of 4K displays in form factors similar to current displays resulting in a 4X increase in the networking requirements for remote visualization. In general, we expect tiled displays to become more commonly accepted by remote science teams for data exploration and visualization. Low overhead network provisioning and scheduling will be critical to rapid adoption of remote visualization and analysis.

Beyond 5 years, data sizes from science campaigns are expected to continue to grow. We expect to witness advent of 8K displays, wider adoption of tiled walls (both 2D and 3D) facilitated by technologies including OLED displays. We expect tiled walls to become extremely prevalent in the science community. Seamless network provisioning and scheduling will be key to rapid adoption of remote visualization and analysis.

## **1.7 Software Infrastructure**

Please see Section 1.5.

## **1.8 Outstanding Issues**

No outstanding issues at this time.

Table 1.2: The following table summarizes data needs and networking requirements for the ALCF.

Key Science Drivers			Anticipated Network Needs	
Instruments, Software, and Facilities	Process of Science	Data Set Size	Local-Area Transfer Time	Wide-Area Transfer Time
<b>0-2 years</b>				
Mira, a compute resource.  Cooley, a new visualization cluster.  Ocular, a visualization wall.	INCITE and ALCC awards from national laboratories, universities, corporations, and international partners.	Size varies, approx. 200 TB per award, and approximately 100,000 files per award.  Visualization data sets range in size from (2048) <sup>3</sup> to (10,240) <sup>3</sup> , with a single time step of data in the range of terabytes.	Largest simulations can move 384 GB/sec peak theoretical across LAN.  The visualization wall can drive 31Gbps.	Largest users move up to 60 TB/day.  Targets are globally diverse, but typically either ESnet or Internet2.
Cosmology Simulation Activities	Requirements of the workflow result from the distribution of resource allocation and availability at DOE computing facilities and user access to data. Availability of computing resources offline analysis can proceed in either asynchronous or co-scheduled modes is also important.	Between 10 and 100 TB in files within a set, with tens to hundreds of files per set.	-	Transfers not generally tightly coupled to workflow, a file set of ~100TB should be transferable on the time scale of one day.
<b>2-5 years</b>				
Theta, a compute resource  a planned new visualization wall	Potential need for real time analysis with companion visualization and analysis resource located on the WAN.  Potential need for real time analysis with companion visualization and analysis resource located on the WAN.	Initial sizes expected to match current use, system is sized for larger appreciably overall storage for same number of projects.	Largest simulations on the new resource will experience an increase in peak theoretical speed.  The full planned visualization wall could drive up to 100Gbps.	Expectations are to reach 2x of current levels.  Targets will remain globally diverse.
Cosmology Simulation Activities	Level of simulation activity will be driven by a number of important new surveys and campaigns such as the ongoing Mira-Titan Universe simulation.	Large individual runs will create of order 10 PB of Level 1 data per simulation in roughly a month of wall clock time.	-	A hard requirement of moving data at the rate of ~1 PB/week.
<b>5+ years</b>				
Aurora, a compute resource	The egress connectivity for the ALCF must be expanded considerably to keep pace with potential demand.	Initial sizes expected to grow, with system storage being considerably larger for same number of projects.	Largest simulations on the new resource will experience a significant increase in peak theoretical speed.	Expectations are that transfers will have significantly increased.  Targets will remain globally diverse.

## Case Study 2

# National Energy Research Scientific Computing Center

### 2.1 Background

The National Energy Research Scientific Computing Center (NERSC) is the high-end scientific computing facility for the DOE's SC. With about 6,000 users from universities, national laboratories, and industry, NERSC supports the largest and most diverse research community of any computing facility within the DOE complex. NERSC provides large-scale, state-of-the-art computing for DOE's unclassified research programs in alternative energy sources, climate change, energy efficiency, environmental science and other science areas within the DOE mission.

There has been an explosive growth of observational and experimental data from DOE facilities and a number of new capabilities are expected to come online in the next 5 to 10 years including an upgrade to the Advanced Light Source (ALS) at the Lawrence Berkeley National Laboratory, the Linac Coherent Light Source (LCLS-II) experiment at the SLAC National Laboratory, increased data rates from the Planck satellite and the the Large Synoptic Sky Telescope (LSST). In genomics significant upgrades are expected in sequencing capabilities and more ubiquitous distributed sensor networks will be placed into our urban and natural environments. Some of these facilities, with large bursty workloads are expected to stress the existing networking capabilities. Large scale data analytics and simulation capabilities are required to achieve the science goals of these observational and experimental instruments and drive the need to have high data ingest rates, in some cases, directly into the supercomputer.

### 2.2 Network and Data Architecture

The network at NERSC can be roughly divided into two distinct parts. There is the internal network that serves as the high-speed interconnect between storage and computing resources, and the external-facing network that is used to access the computing and storage resources and other public-facing services (e.g., web servers and scientific portals). Each of these networks is specifically tuned for the needs that it serves.

NERSC continues to move toward a fault-tolerant architecture from the network edge through the core. Hosts are deployed with dual network connections to an Ethernet switch fabric consisting of two or more switches at the distribution layer which function as a single logical device. For the internal network, uplinks from the distribution layer to the core go to a pair of redundant routers. This architecture aims to remove single points of failure in the network, reducing the likelihood of a service interruption. It also provides two independent network links of bandwidth to the edge hosts during normal operation (Figure 2.1).

For the internal network, a significant portion of NERSC compute clusters and file system servers rely on IB as a high-bandwidth, low-latency interconnect (Figure 2.2). The clusters use a set of nodes acting as network gateways for connectivity outside of the cluster. NERSC is collaborating to deploy software routers, which will route IB over IP using standard routing protocols. This will help improve cluster load balancing and availability.




Figure 2.1: NERSC network fault tolerant architecture.




Figure 2.2: Software routers improve connectivity for Infiniband clusters.

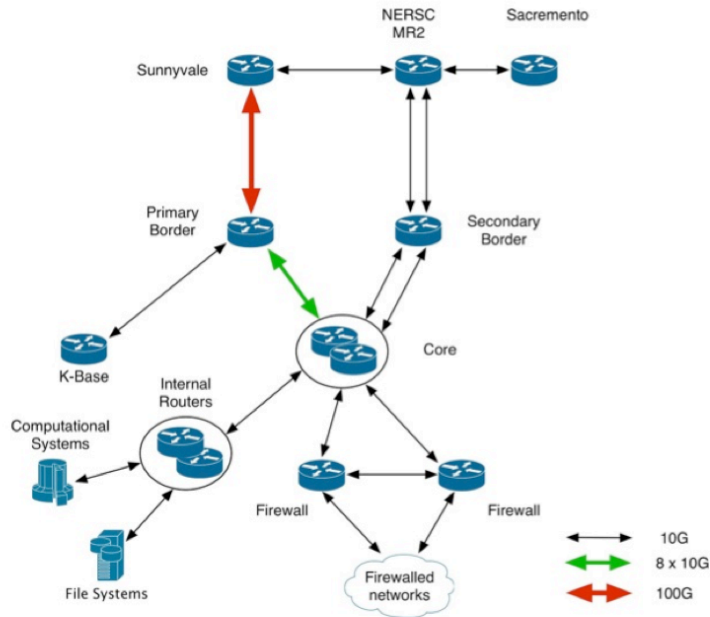


Figure 2.3: NERSC connectivity to ESnet.

For the externally facing network, NERSC’s 100Gbps border router is now fully in production and is the primary connection to ESnet (Figure 2.3). The old border router with 2x10G connections to ESnet functions as a secondary path in case of failure. The primary and backup border routers have physically separate peerings to ESnet, allowing NERSC to continue to operate during router failure or maintenance. NERSC has four DTNs which are optimized for high-bandwidth network traffic. These DTNs mount the NERSC global file system which is also mounted on the NERSC supercomputers.

## 2.3 Collaborators

NERSC has thousands of users not only across the United States, but the world. Figure 2.4 shows the number of NERSC users per country.

## 2.4 Instruments and Facilities

### 2.4.1 High Performance Computing Systems

The NERSC HPC systems are the current primary driver of network traffic at NERSC. NERSC users continue to import more data than they export, using the HPC systems to analyze data.

The largest system on the floor currently is the Edison system, which has a peak performance of 2.57 PetaFLOP/s, 133824 compute cores, 357 TB of memory and 7.4 RB of online disk storage with a peak I/O bandwidth of 164 GB/s. Two other systems, Hopper (a 1 PetaFLOP Cray XE6), and Carver (a smaller cluster) will be retired later in 2015. NERSC also operates two data-intensive systems, one for the high energy physics community, named PDSF, and one for the Joint Genome Institute, named Genepool. NERSC’s next supercomputer, Cori, will be a 30 PetaFLOP system delivered in two phases, the first phase in 2015 with Intel Haswell compute nodes and the second phase in 2016 with Intel Knights Landing compute nodes. NERSC intends make the first phase of the Cori system especially friendly for data-intensive applications, supporting real-time workflows and including an




Figure 2.4: The majority of NERSC computing hours are used by scientists at Department of Energy National Laboratories or universities.

NVRAM burst buffer, interactive nodes for advanced workflows, and improved network performance into the compute system to support streaming applications.


The graph below shows the data transferred into and out of the NERSC Center by source and destination. The graph indicates that NERSC, together with ESnet have been quite successful at encouraging users to use the highly optimized DTNs for network transfers. The DTNs support approximately 50% of the inbound traffic to NERSC, followed next by transfers directly into NERSC’s HPSS archival storage. Data from the Joint Genome Institute sequencers comprise a relatively small amount of data transferred to NERSC. The new Cori system will have 1.5 PB of a non-volatile storage as a “burst buffer,” a layer of NVRAM that will sit between the compute node memory and the file system, which will help accelerate I/O applications. NERSC is working closely with Cray, the Cori vendor, to improve networking capabilities into and out of the Cray compute nodes in order to be able to support streaming data applications. In the future NERSC could anticipate scheduling network transfers concurrently with compute capabilities.

## 2.4.2 Light Sources


Light Sources present a “bursty” use case for NERSC, in that there is a burst of network traffic when the experiment is turned on, but there is a downtime associated with the duty cycle of the experiment (i.e., sample preparation). At present, various beamlines at the ALS (at Lawrence Berkeley National Laboratory), the Linear Coherent Light Source (at SLAC National Laboratory), and the Transmission Electron Aberration-Corrected Microscope (TEAM) high-speed camera (at Lawrence Berkeley National Laboratory) are major use cases. The ALS range of beamlines (tomography, small-angle diffraction, COSMIC, infrared) are representative of major network drivers in the 0 to 5 years timeframe. At SLAC, the CXI beamline and diffractive imaging beamlines are drivers in 0-5 years. Once LCLS-II comes online, it will present a host of challenges for real-time analytics and offline computation. Finally, the proposed TEAM camera, which should come online in 2 years, will be an excellent test case for LCLS-II-like throughput rates; this specific instrument will be capable of 20,000 foot-pound-second acquisition rates.

### Process of Science


The SPOT suite is an emerging platform for handling ALS beamline workloads. In the suite, images are collected into an HDF5 suitcase, and files are transferred to NERSC using Spade and GridFTP. The LCLS instrument at SLAC utilizes the Python Script ANALysis (PSANA) workflow. The TEAM microscope will likely use HDF5 and some data transfer mechanism for efficient data movement and real-time analytics at NERSC. We do not anticipate major changes in the list of software technologies in the 2-5 year timeframe, however these tools will need to accommodate efficient multi-core execution on Cori-like architectures, and the workflows will need to take burst-buffer-like hardware into account. Both “online” real-time analytics, perhaps coupling with a simulation code, as well as of-



(a)



(b)



(c)

Figure 2.5: Graphs characterizing NERSC users by institution type, location for national laboratories, and a break down of universities.




Figure 2.6: Graph showing the total bytes (in TB) imported and exported per month at NERSC.




Figure 2.7: Graph showing NERSC's routed traffic via the wide-area network in TB/day per year.






Figure 2.8: Percentage of data in and out of NERSC by source and destination.

fine analytics, will be important use cases.

### 2.4.3 High Energy Physics

NERSC is involved in community support of a number of high energy physics (HEP) experiments. The Solenoidal Tracker at RHIC (STAR) detector at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory; the A Large Ion Collider Experiment (ALICE) and A Toroidal LHC Apparatus (ATLAS) experiments at Large Hadron Collider (CERN); the Daya Bay experiment in China, and the Large Underground Xenon (LUX) Dark Matter detector are major drivers for NERSC. The network traffic modality tends to be in a stage of “Continuous Import” and “Continuous Export.” The HEP community has a mature set of workflow tools and technologies; and the infrastructure is working well currently. The data volumes will increase by ten-fold in the next 2–5 years, which will require scaling bandwidth and compute resources. STAR and ALICE instruments present both import and export use cases, which are relatively symmetric in their bandwidth requirements. LUX and Daya Bay are primarily import-oriented use cases at this point in time.

**Process of Science** ALICE utilized the XRootD framework, which is managed with AliEn. STAR utilizes a couple of pipelines which use globus-url-copy or Globus internally. ATLAS utilizes ROOT and BeStMan for transfers. Daya Bay utilizes Gaudi/ROOT and SPADE for transfers. We do not expect major changes in the list of technologies in the 2–5 year timeframe.

### 2.4.4 Astronomy and Cosmology

The Dark Energy Spectroscopic Instrument (DESI) and Cosmic Microwave Background (CMB) telescopes present both bursty and continuous import use cases for NERSC in the 0–5 year timeframe. The DESI project will focus on targeting surveys and simulations in the 0-2 year timeframe, and focus towards serving data in the 2–5+ year period. Raw data from the DESI project will be mirrored in a “continuous export” mode in the 2–5+ year timeframe; with a modest annual data volume of 10 TB/year. Public CMB datasets (Planck and WMAP) are hosted at NERSC at present, and will continue to be available for download to collaborators in the 0–5+ year timeframe. Simulations for the LSST project will be made available at NERSC in the 2–5+ year timeframe for serving reduced data products. The major driver for the field will start with the onset of the LSST project in 2022.

Table 2.1: The following table summarizes data needs and networking requirements for NERSC's light sources use cases.

Key "Bursty" Science Drivers			Anticipated Network Needs	
Instruments, Software, and Facilities	Process of Science	Data Set Size	Local-Area Transfer Time	Wide-Area Transfer Time
<b>0-2 years</b>				
<ul style="list-style-type: none"> <li>ALS (Tomography, Small Angle Diffraction); LCLS (CXI: Coherent X-ray Imaging); TEAM (high speed cameras)</li> <li>ALS: Spade; GridFTP; HDF5. LCLS: PSANA workflow.</li> </ul>	<ul style="list-style-type: none"> <li>Highlights of current science process--ALS: collect images; prepare HDF5 suitcase; transfer to NERSC. LCLS: 1) Fast feedback analysis: check data quality while experiment is running; 2) Offline analysis: data is processed at LCLS/NERSC to perform image reconstruction. TEAM: transfer data over "LAN" to NERSC; Cori BB; real-time analytics.</li> </ul>	<ul style="list-style-type: none"> <li>Size of one data set at the ALS: 100MB-100GB; LCLS: 20 GB files; 150TB/experiment. TEAM: 1-10TB/dataset</li> <li>Aggregate dataset size (annual)--ALS: 10-300TB. LCLS: 600 TB @NERSC; 3PB/year. TEAM: not in production in this timeframe.</li> <li>What is the data set composed of? ALS: HDF5 file consisting of several images (datasets). LCLS: XTC format. TEAM: HDF5 file.</li> </ul>	<ul style="list-style-type: none"> <li>How long does it take to transfer a data set on the local network? ALS: 1-2 minutes. TEAM: target "real-time feedback."</li> <li>How frequent are the transfers? 10% duty cycle (TEAM); ALS; LCLS</li> </ul> <p>Based on 10Gbps.</p>	<ul style="list-style-type: none"> <li>How long does it take to transfer a data set offsite? LCLS: 7.5 Gbps obtained at present.</li> <li>How frequent are the transfers? LCLS: continuous (24x7) streaming 4 days/nights in a row. LCLS data generation is 2x transfer rates; night shifts allow transfer to catch up. TEAM: 8 hour shifts; low duty cycle.</li> <li>Where are the collaborating sites/destination points for the data transfers/data sets? LCLS to NERSC DTNs.</li> </ul>
<b>2-5 years</b>				
<ul style="list-style-type: none"> <li>ALS: COSMIC, IR. LCLS-II: diffractive imaging, structural biology, etc (2019). TEAM: 20,000 fps acquisition.</li> <li>ALS: HDF5, LCLS: PSANA/HDF5, TEAM: HDF5.</li> </ul>	<ul style="list-style-type: none"> <li>Cori architecture; data partition; streaming to burst buffer and real-time analytics. Porting analytics codes and workflows to Cori. Possible coupling with simulation in limited cases. LCLS: offline analysis will be critical.</li> </ul>	<ul style="list-style-type: none"> <li>Size of one data set: ALS: 1GB-100GB. LCLS: 600TB/experiment. TEAM: 10-100+TB</li> <li>Aggregate size of datasets (annual): ALS: 100TB-3PB, LCLS: 600TB/experiment 2.4PB@NERSC</li> </ul>		<ul style="list-style-type: none"> <li>LCLS: continuous (24x7) streaming 4 days/nights in a row.</li> <li>LCLS: NERSC Cray COE; streaming writes to Burst Buffer + storage. Real-time and offline analytics.</li> </ul>
<b>5+ years</b>				
<ul style="list-style-type: none"> <li>Describe any planned new data sources or software packages: LCLS: next gen XFEL.</li> </ul>	<ul style="list-style-type: none"> <li>What is the strategic direction for data flow, science process, etc.? LCLS: Resolution and sampling rates for LCLS-II: start with 10x compared to LCLS; proceeding to 100x. Robotic sample handling will improve duty cycle; Dynamical reconstruction will require processing of 10-100x data (compared to static); Simulation of experiments requiring HPC.</li> </ul>	<ul style="list-style-type: none"> <li>Size of one data set: ALS: 10GB-?; LCLS: 2-6PB data set size / experiment are possible given 2 order of magnitude improvement in technology.</li> <li>Aggregate size of datasets (annual): ALS: 100TB-5PB. LCLS: 30PB.</li> </ul>		<ul style="list-style-type: none"> <li>Reduction in duty cycle with automation might imply more frequent transfers.</li> </ul>

Table 2.2: The following table summarizes data needs and networking requirements for NERSC's HEP use cases.

Key Science Drivers: "Continuous Import and Export"			Anticipated Network Needs	
Instruments, Software, and Facilities	Process of Science	Data Set Size	Local-Area Transfer Time	Wide-Area Transfer Time
<b>0-2 years</b>				
<p>STAR@RHIC; ALICE@LHC; ATLAS@LHC; Daya Bay; LUX Dark Matter Detector. DESI: 2 telescopes.</p> <p>· What is the current/new software used in scientific process? ALICE: XRootD, managed with AliEn framework; STAR: couple of pipelines that use globus-url-copy or globus online; ATLAS: ROOT; BeStMan (SRM) for transfers; Daya Bay: Gaudi/ROOT, SPADE for transfers.</p>		<p>· What is the size of one data set? STAR/ALICE: ~100TB/dataset. ATLAS: GB-TBs. Daya Bay: ~0.6GB.</p> <p>· What is the general range of data set sizes? ALICE/STAR: ~1TB-1PB; full ATLAS: 100PB; Daya Bay: ~70TB/year raw, ~1PB derived; LUX: 100s TB/yr, file sizes vary. DESI: 1 TB/night; processed: 4TB.</p> <p>· What is the data set composed of? ALICE/STAR: ~1M ROOT files. ATLAS: 100K files; 10-100MB each. Daya Bay: 350x0.5GB files/day</p>	<p>· How long does it take to transfer a data set on the local network? ALICE/STAR: streaming from local storage to compute nodes: 10TB in 35 minutes @5GB/s</p> <p>· How frequent are the transfers? 24/7 streaming</p>	<p>· How long does it take to transfer a data set offsite? ALICE: 50MB/s - &gt; 250 MB/s. STAR transfers: 60 MB/s. ATLAS: individual file: 3 MB/s (from CERN, BNL), aggregate: 150 MB/s peak, limited by PDSF DTN bandwidth of 10Gb/s; Daya Bay: PDSF DTNs (3 MB/s aggregate from China); LUX: aggregate rates of ~20MB/s to NERSC DTNs, and 1MB/s to PDSF DTNs; goal to get PDSF DTNs to 10MB/s.</p> <p>· How frequent are the transfers? ALICE: steady state. STAR: duty cycle of 1/3. Desire to increase rates by order of magnitude: 3x via network optimizations; 2x via continuous operations. ATLAS: user initiated; between 3-30 transfers/day. Daya Bay, LUX: several times / day. DESI: 30 and 180 nights/yr.</p> <p>· Where are the collaborating sites/destination points for the data transfers/data sets? ALICE: 80+ grid sites world-wide; closest sites that dominate transfers: ORNL, UNAM, KISTI. ALICE uses ALICE grid enabled storage elements (and not NERSC DMZ dtm servers); STAR: BNL, relies on PDSF and/or NERSC DTNs. ATLAS, Daya Bay: dedicated PDSF DTNs. LUX: primary data mirror at Brown University.</p>
<b>2-5 years</b>				
<p>· What are the planned, new data sources/instruments? STAR@RHIC; ALICE@LHC; ATLAS upgrade; Daya Bay upgrade</p>	<p>· What are the foreseeable changes to data flow, science process, etc? ATLAS: 10x increase; Daya Bay: 10x increase possible. ATLAS: Increased use of on-demand WAN data access using federation technologies such as xrootd; as well as continuing pre-placement of data.</p>	<p>· Size of one data set: STAR, ALICE numbers same as 0-2 years.</p> <p>· Range of data set sizes (e.g. 500GB to 2TB depending on experiment)</p> <p>· Total ATLAS dataset size will double in 2 years; 25% increase during shutdown 2018-2019; further 2x increase by 2023.</p>	<p>· How long does it take to transfer a data set on the local network?</p> <p>· How frequent are the transfers? STAR, ALICE numbers same as 0-2 years.</p>	<p>· How long does it take to transfer a data set offsite? ALICE/STAR: target performance improvement of 2-3x as facility grows.</p> <p>ATLAS: Network usage will scale with dataset size: roughly double over next 2 years. Would like WAN networks to allow 1-10MB/s/core; around 10Gbps aggregate.</p>
<b>5+ years</b>				
<p>· Describe any planned new data sources or software packages:  Future for STAR in question; related to RHIC and plans for electron-ion collider</p>	<p>· What is the strategic direction for data flow, science process, etc? Significant changes for ALICE; 1 TB/s coming out the detector. New computing project underway to develop technology. Regional analysis facilities being developed to hold reduced datasets. PDSF could support this for ALICE-USA.</p>		<p>· How long does it take to transfer a data set on the local network? ALICE: Current target: 5GB/s based on current workflow. As facility increases (say 2x); need to scale bandwidth capacity to at least 10GB/s.</p>	<p>· How long does it take to transfer a data set offsite? ALICE: in a Tier-2 role will grow by 10x relative to present. If PDSF becomes an analysis facility; we will need to move ~1 PB-sized datasets routinely over ~3 weeks; indicating a performance target of 500MB/s from CERN-&gt;NERSC. ATLAS: need proportional scaling in bandwidth with data volume increase: 25% increase in 2018-2019, followed by 2x by 2023.</p>

Table 2.3: The following table summarizes data needs and networking requirements for NERSC's cosmology and astronomy use cases.

Key Science Drivers "Bursty stuff"			Anticipated Network Needs	
Instruments, Software, and Facilities	Process of Science	Data Set Size	Local-Area Transfer Time	Wide-Area Transfer Time
<b>0-2 years</b>				
· DESI: Targeting surveys and Simulations.	DESI activities leading up to commissioning include dedicated observing runs on other telescopes to build a list of targets to observe with the DESI instrument. The data from these runs are transferred to NERSC.	· DESI: Targeting surveys are 4-80TB each. Total transfer is ~150TB / year. DESI cosmology simulations (level-2 products) are 20TB, and we will transfer several per year.	For bursty DESI data, files are copied to HPSS and staged on NGF. Then data is copied to Lustre scratch for processing. Existing transfer times are sufficient.	For DESI targeting, data, one day of data should transfer to NERSC in less than one day. For cosmology simulations, we should be able to transfer one set of level-2 products in tens of minutes. Both of these requirements are already met.
· CMB Telescopes: Data serving	Cosmological simulation products generated at other leadership class facilities will be transferred to NERSC. Public CMB datasets (Planck and WMAP) are hosted at NERSC and available for use / download.	· CMB Telescopes: serving public data will be mostly limited to large transfers to other university and academic institutions. Roughly 50TB per year.	CMB data being served is not moved locally.	A reasonable number would be the ability to transfer ~10TB in less than a day to other US academic institutions.
<b>2-5 years</b>				
· DESI Targeting, simulations, and data serving.	· DESI targeting is ramping down during this time. DESI cosmology simulations continue. The main DESI instrument begins to collect data and this is distributed to the collaboration. Public data from CMB telescopes continues to be served, and simulated data in support of the DOE Stage-4 CMB program is generated at NERSC and distributed to the collaboration. During this time, LSST simulations will begin to be generated at other institutions and copied to NERSC.	· DESI: targeting down to 50TB / year. Cosmology sims still several at 20TB each. Serving data to public will be bursty transfers of about 100TB / year.	· For bursty DESI data, files are copied locally between NGF, HPSS, and Lustre scratch. Existing transfer times are sufficient.	· For DESI targeting, data, one day of data should transfer to NERSC in less than one day. For cosmology simulations, we should be able to transfer one set of level-2 products in tens of minutes.
· CMB Telescopes: data serving		· CMB telescopes: still serving roughly 50TB / year, ramping up to maybe 100TB per year including simulated data.	· CMB data being served is not moved locally.	· A reasonable number would be the ability to transfer ~10TB in less than a day to other US academic institutions.
· LSST: transfer of simulations to NERSC		· For LSST simulations, we expect the data to be roughly the size of one year of actual observing, which is 5.5PB.	· LSST simulations will likely be transferred directly to HPSS and then small pieces will be extracted to scratch for testing. We need to be able to transfer a day's worth of data (15TB) from HPSS to scratch in 10-20 minutes.	· The transfer of LSST simulations will be done over the WAN directly to HPSS. It would be good to be able to transfer one of these simulations in about a week.
<b>5+ years</b>				
· DESI simulations and data serving.	· We will continue bursty transfers of cosmology simulations and serving of DESI data.	· We will continue at ~100TB per year of cosmology simulation products and ~100TB / year of data serving	· For bursty DESI data, files are copied locally between NGF, HPSS, and Lustre scratch. Existing transfer times are sufficient.	· For cosmology simulations, we should be able to transfer one set of level-2 products in tens of minutes.

## **2.5 Software Infrastructure**

NERSC users use a variety of software to transfer data and manage their workflows. Most traditionally, users use scp and ftp. In the past few years more NERSC users have moved to GridFTP.

As described in the case studies many users with complex workflows use customized software for data management and data analysis.

## **2.6 Cloud Services**

NERSC uses the commercial cloud for limited business services (ticketing system) and does not anticipate using the commercial cloud for data storage, analysis, or computing.

## **2.7 Outstanding Issues**

A growing use case is the desire to be able to stream data directly into the compute partition of the NERSC supercomputers. Today, bandwidth into the supercomputing systems' compute nodes is limited by intermediate networking nodes. We are working with our supercomputing vendors, in this case Cray, to improve streaming capabilities that will allow experimental facilities to stream data directly into the burst buffer or filter data on intermediate nodes. We would appreciate collaborating with ESnet on this effort.

## Case Study 3

# The Oak Ridge Leadership Computing Facility

### 3.1 Background

The Advanced Scientific Computing Research (ASCR) program within the Office of Science at the Department of Energy (DOE) established the Leadership<sup>1</sup>Computing Facility in 2004. The ASCR supports the Oak Ridge Leadership Computing Facility (OLCF) as a highly collaborative, Office of Science national user facility dedicated to leading-edge computational capabilities that advance fundamental discovery and understanding in a broad range of scientific and engineering disciplines. The OLCF deploys HPC architectures that are 10–100× more powerful than systems typically available for open scientific research. The OLCF leverages infrastructure of massive data storage, high bandwidth network connectivity, and advanced visualization resources, resulting in the world’s leading computational science infrastructure. The facility partners with scientists, engineers, mathematicians, and computer scientists, along with the software and system development community, to continuously innovate solutions to computational science challenges pacing scientific progress across a broad spectrum of research domains.

Today, the OLCF is home to Titan, a hybrid-architecture Cray XK7 system with a theoretical peak performance exceeding 27 petaFLOPs. Titan features 18,688 compute nodes, (each with one 16-core AMD Opteron CPU and 1 NVIDIA Kepler K20X GPU), 299,008 x86 cores, a total system memory of 710 TB, and a high-performance proprietary network. The combination of these technologies allows Titan to achieve up to 10 times the speed of its predecessor, the Jaguar supercomputer—a Cray XT5 system, while consuming the same average power load and occupying the same physical footprint. The system provides decreased time to solution, increased complexity of models, and greater realism in simulations.

### 3.2 Network and Data Architecture

The OLCF network consists of various production and test networks.

#### 3.2.1 Egress Connectivity

Currently OLCF is connected to ORNL via four 10 Gigabit Ethernet connections. ORNL currently has a 100 Gigabit Ethernet connection and a 10 Gigabit Ethernet backup link to ESnet. OLCF is working with ORNL to bring the 100

---

<sup>1</sup>The term Leadership System means a high-end computing system that is among the most advanced in the world in terms of performance in solving scientific and engineering problems. [As defined in Public Law 108–423 Nov. 30, 2004 Department Of Energy High-End Computing Revitalization Act of 2004] and proposed in the Federal Plan for High-End Computing: Report of the High-End Computing Revitalization Task Force (HECRTF)—May 10, 2004.

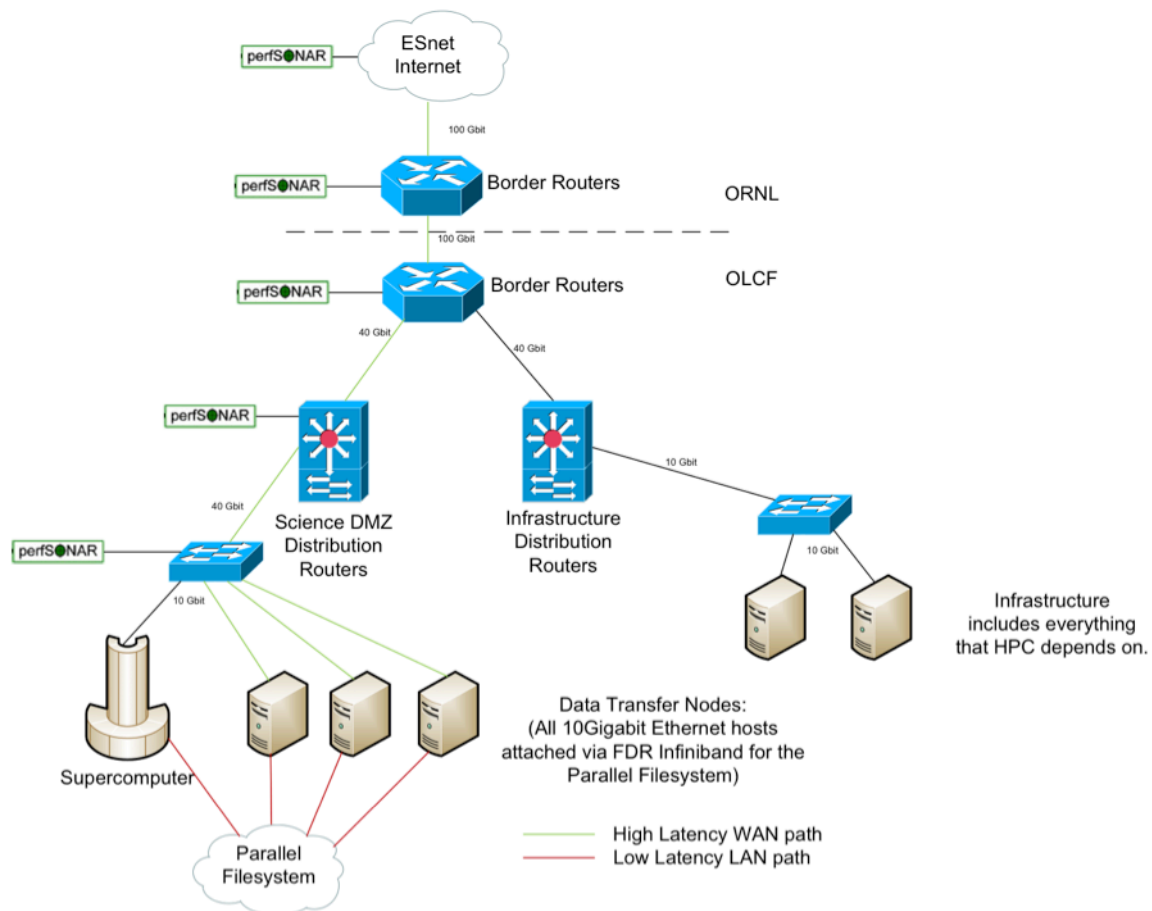


Figure 3.1: Schematic of the upgraded OLCF network architecture.

Gigabit Ethernet path into the OLCF core network, keeping several 10 Gigabit Ethernet links as backup paths. The OLCF infrastructure is deployed in a highly available architecture such that a single switch, router, interface, or server failure will not drastically impact the center. Although the HPC systems themselves are not redundant, having a redundant infrastructure in critical areas ensures that HPC resources, file systems, clusters, test systems, and critical infrastructure are able to function and are not subject to single points of failure. Figure 3.1 shows the OLCF network as it will operate once the previously described upgrades are complete.

### 3.2.2 Science DMZ

In 2009, both the OLCF and NERSC installed a number of DTNs in order to better facilitate the high-speed transfer of data between the two sites. Most users observed a significant performance gain of at least 20x when using the newly installed infrastructure. In 2015, the OLCF continued to pursue innovative operational approaches. One such approach is a ScienceDMZ model, that the OLCF worked toward in 2015.

A Science DMZ model, as defined by ESnet engineers, seeks to provide a distinct network architecture for high-performance applications that is simple, incrementally scalable/deployable, and adaptable to new and emerging hardware technologies. The ESnet model states there are three key components of a science DMZ: a friction-free network path comprised of high-performing networking devices, high-performance servers dedicated to the functions relating to the transfer of data, and the means to monitor and analyze overall performance.

The OLCF is deploying a Science DMZ, which will allow redundant 40 Gigabit Ethernet paths from the OLCF border layer down through the access layer. All HPC resources such as Titan, HPSS, and the DTNs are connected to access

switches inside of the Science DMZ. All HPC resources are connected via 10 Gigabit Ethernet on the high latency WAN path, and via FDR Infiniband on the low latency LAN path to the center wide parallel file system.

### **3.2.3 Friction Free Network Path**

The underlying objectives of a friction free network path are to utilize highly capable networking devices, logically locate them at or near the site perimeter, and devise appropriate security policies that do not hinder scientific productivity. To this end, the OLCF has focused on redesigning its network to support a tiered hierarchy. This provides enhanced redundancy, scalability, and fast implementation of new technologies with minimal to no disruption of production traffic.

In 2015, a Cisco Nexus 7710 and two Cisco Nexus 6004's were deployed within the existing network to bring in high density 40 and 100 Gigabit Ethernet within the data center. These chassis operate at wire speed and will allow the OLCF to upgrade the data path between OLCF and ORNL to a 100 Gigabit Ethernet path. This was followed by the purchase of two Fortinet firewalls that are capable of processing 160 Gigabits of throughput and handling 50 million concurrent sessions. Firewalls can limit throughput and impede transfers over wide area networks. To have a friction-free network, firewall functions are moved from the border closer to the HPC infrastructure to prevent delay and degrading transfers in the broader Science DMZ.

Selection of appropriate security policies is extremely important to adequately secure and enable high-performance data transfers. To this end, the OLCF worked with ORNL risk management to quantify the security risk and productivity gains of extending GridFTP certificate lifetimes. The OLCF tripled the maximum lifetime of grid certificates. Users can now request grid certificates with lifetimes of up to 72-hours. This extension enables weekend transfers that can occur with little to no user interaction. Network intrusion detection and monitoring systems were refreshed in CY 2015 and deployed in a clustered model. This approach allows for redundant 100Gb/s link monitoring as well as traffic aggregation, shunting, and distribution to various network security sensors that are tailored for specific applications, such as Snort and Bro IDS. OLCF deployed a local certificate authority so that users can easily obtain temporary data transfer certificates without needing to register with the Open Science Grid. This removed several steps from the data transfer workflow allowing project teams to more easily transfer data in and out of the facility. The OLCF now provides OAuth to Globus as an authentication mechanism for validating OLCF users and distributing OLCF user credentials. This is much more secure than traditional authentication methods where Globus acts as a credential proxy when a user activates an OLCF data transport endpoint. Users are now able to leverage Globus to facilitate data transfers faster and more securely than ever before.

These additional measures provide increased network and system availability, high-performance network connectivity, dedicated friction-free paths, and more accurate proactive monitoring and security capabilities, which will be of great benefit to our users.

### **3.2.4 Data Transfer Nodes**

Since the introduction of the OLCF's first data transfer nodes in 2009, the capacity and offerings have increased significantly in those areas. At present the OLCF maintains 20 data transfer nodes dedicated to interactive, scheduled, and archive specific transfer functions. In 2015, the OLCF began implementing new hardware for these DTNs, which will include an upgrade from 10 Gigabit to 40 Gigabit networking. These new nodes are managed as a diskless cluster, providing for increased speed of deployment when new DTNs or new software are needed. Additional focus was also placed on transfer-related service offerings such as Globus Online and GridFTP their interactions with extant services within the OLCF.

## **3.3 Collaborators**

LCF user projects are made without regard to funding source or affiliation (U.S. industry, academia, national laboratory, and other federal agencies). The user population tends to be highly diverse, representing a wide range of scientific disciplines (Figure 3.2). The physics category represents primarily astrophysics, plasma physics,






Figure 3.2: Wide reach of Leadership computing across many science domains is a testament to its ubiquity as a scientific tool.

high-energy/particle physics, and nuclear physics. By design, the number of Leadership Computing Facility (LCF) user projects is smaller than at many other high-performance computing facilities so that the LCF can provide the large-scale computing, storage, and networking resources, coupled with a high level of support needed to solve the most difficult scientific challenges. The INCITE program is the largest allocation program by which the scientific community gains access to the LCF. The INCITE program is highly competitive and designed to enable grand-challenge investigations and discoveries in science and engineering. Typically, OLCF will support about 30 INCITE projects per year. The ASCR Leadership Computing Challenge (ALCC) program allocates large resources to projects of interest to the DOE with an emphasis on high-risk, high-payoff simulations in areas directly related to the DOE mission and for broadening the community of researchers capable of using Leadership computing resources. Recently, ASCR has been selecting approximately 25 project per program (1 July through 30 June) year for OLCF. An additional allocation program, Director's Discretionary (DD) allocations, provides awards of resources for development work, outreach to new users, and strategic laboratory projects. In order to meet the goals for the DD program, OLCF typically awards 170 to 200 DD projects, 3 months to 12 months in duration, over the course of one year. During 2015, the OLCF supported 1,176 users on over 316 projects active at any given point during the year.

## 3.4 Instruments and Facilities

### 3.4.1 Present

#### Cray XK7 (Titan) Resource Summary

The OLCF upgraded the existing Cray Jaguar from a model XT5 to a model XK7, releasing it to production on May 31, 2013. The resulting system contains 18,688 NVIDIA K20X (Kepler) accelerators, in which each existing AMD Opteron connects to an NVIDIA Kepler to form a CPU-GPU pair. The completed XK7 system, with more than 27 petaflops of peak computational capacity, is named Titan.

### **Cray XC30 (EOS) Resource Summary**

Eos is a four-cabinet Cray XC30. The system, with 736 Intel Xeon E5-2670 compute nodes and 47.6 TB of memory, provides the OLCF user community with a substantive large-memory-per-node computing platform. The Eos nodes are connected by Cray's Aries interconnect in a network topology called "Dragonfly." All INCITE users are automatically granted access to the XC30.

### **Lustre File Systems (Spider II) Resource Summary**

In October 2013, the OLCF released Spider II, its next-generation Lustre parallel file system, to production. Spider II contains two instantiations of the /atlas file system, with an aggregate capacity of more than 30 petabytes (PB) and block-level performance of more than 1.3 TB/second. The Spider II file system is the default high-performance file system for all compute systems. The previous generation Lustre file system, Spider I (collectively the four /widow file systems) was decommissioned during the 2013 reporting period.

### **Data Analysis and Visualization Cluster (Rhea) Resource Summary**

Rhea is a 512-node large memory data analytics Linux cluster. The primary purpose of Rhea is to provide a conduit for large-scale scientific discovery via pre- and post-processing of simulation data generated on Titan. Users with accounts on INCITE- or ALCC-supported projects are automatically given accounts on Rhea. DD projects may request access to Rhea. Each of Rhea's nodes contain two 8-core 2.0 GHz Intel Xeon processors with hyper-threading and 128 GB of main memory (upgraded in 2015 from 64 GB). New in 2015, Rhea offers nine additional nodes, each of which boast 1 TB of main memory and 2 NVIDIA Tesla K80 (Kepler GK210) GPUs. Rhea is connected to the OLCF's 30+ PB high-performance Lustre file system, Spider II.

### **High Performance Storage System (HPSS) Resource Summary**

The OLCF provides a long-term storage archive system based on the High Performance Storage System (HPSS) software product co-developed by IBM, Los Alamos National Laboratory, Sandia National Laboratories, Lawrence Livermore National Laboratory, Lawrence Berkeley National Laboratory, and ORNL. The ORNL HPSS instance is currently over 50 PB in size and provides up to 200 Gbps of read and write performance. The archive has taken in over 225 TB in a single day several times in the last year; the previous daily maximum was just over 150 TB.

The archive is built from hardware from Dell, Hewlett Packard, Brocade, NetApp, DataDirect Networks, and Oracle. An 18 PB disk cache allows burst rates into the archive at up to 200 GB/s; there is 26 GB/s of read/write bandwidth to the archive via 154 Oracle T10K series tape drives. There are 6 Oracle SL8500 tape libraries for tape archival storage that each contain 10,100 slots; the archive's maximum capacity is over 500 PB, using these libraries.

### **Visualization Resource Summary**

The Exploratory Visualization Environment for Research in Science and Technology (EVEREST) has three computing systems and two separate state-of-the-art visualization display walls. The primary display wall spans 30.5×8.5 feet and consists of eighteen 1920×1080 stereoscopic Barco projection displays arranged in a 6×3 configuration. The secondary display wall contains sixteen 1920×1080 planar displays arranged in a 4×4 configuration, providing a standard 16:9 aspect ratio. The stereoscopic capabilities allow the user to experience binocular depth perception. An array of sequentially pulsed infrared LED cameras record the physical position and orientation of the user, and the resolution density provides an optimal solution for human visual acuity. These combined technologies, along with OLCF staff expertise, allow scientist to analyze complex scientific datasets in an immersive environment and communicate abstract concepts in an intuitive visual format.

Table 3.1: OLCF production computer systems.

System	Type	CPU	GPU	Computational description			Interconnect
				Nodes	Node configuration	Memory configuration	
Titan	Cray XK7	2.2 GHz AMD Opteron 6274 (16-core)	732 MHz NVIDIA K20X (Kepler)	18,688	16-core SMP + 14 streaming multiprocessor (SM) GPU (hosted)	32 GB DDR3-1600 and 6 GB GDDR5 per node; 598,016 GB DDR3 and 112,128 GB GDDR5 aggregate	Gemini (Torus)
Eos	Cray XC30	2.6 GHz Intel E5-2670 (8-core)	None	736	2 × 8-core SMP	64 GB DDR3-1600 per node; 47,104 GB DDR3 aggregate	Aries (Dragonfly)

## OLCF Computational Resource Summary

### 3.4.2 Next 2–5 Years

Summit will arrive for OLCF users beginning in 2018 with full user operations beginning in 2019 for INCITE and ALCC programs. Summit, the 4th major refreshment of OLCF resources (OLCF-4), will be a supercomputer emerging from the OpenPOWER vendor consortium, including IBM, NVIDIA, and Mellanox. Summit will provide a compute and data capability for applications that is 5× to 10× greater than Titan, and is considered a pre-exascale system. It will advance a hybrid, accelerated architecture with new memory technologies such as high bandwidth memory and NVRAM introducing new levels into this hierarchy. A new GPFS-based center-wide file system will be introduced into the OLCF data center during this time frame.

The OLCF anticipates that user requirements for advanced data analytics and workflows will drive the development of new capabilities and services, e.g., workflow-management solutions, data portals, and data-analytics capabilities as these have recently emerged from active user engagement. The OLCF anticipates that these new services will be provided in integration with a “science cloud” infrastructure being built at ORNL called the Compute and Data Environment for Science (CADES).

### 3.4.3 Beyond 5 Years

Beyond 5 years, the OLCF anticipates that the exascale era for OLCF users will occur in 2023 with the culmination of the Exascale Computing Project, and will represent a computing capability 50× to 100× greater than what is available today within Titan.

## 3.5 Process of Science

### 3.5.1 Present

Groups awarded time at the OLCF often transfer in their data sets to the facility at the beginning of their computing time. This comes in the January time frame for INCITE awards, and the July time frame for ALCC awards. Groups also transfer out simulation results to their respective home facilities or to collaborators year-round. The INCITE program process, which allocates 60 percent of the available time on the machine, is also open to international proposals. It is not unusual to have a trans-Atlantic awardee, with the request to facilitate international data transfers.

The OLCF also has a robust industrial partnership program. Data transfers with industrial partners are, at times, challenging because of the relatively low bandwidth into and out of the industrial user’s data facility. Often, this results in their data remaining within the OLCF longer than would otherwise be desired.

To enable data movement, a GridFTP cluster (DTN nodes) is available for the transfer of data sets. During normal operations, 20 DTNs are currently available for users, each with 10 GbE of connectivity. The OLCF is currently deploying a refreshed and enhanced cluster of DTNs that will provide a greater number of higher-performing nodes. This project will be completed in 2016.

### **3.5.2 Next 2–5 Years**

The OLCF anticipates strong growth in distributed services, e.g., remote visualization, distributed workflow management and execution, and new data portals to service integrated compute, analysis, visualization, and, curation of large and/or significant datasets generated at the OLCF facility.

### **3.5.3 Beyond 5 Years**

The generation of (i) exascale datasets from within the OLCF and, (ii) the management of distributed data workflows with large experimental facilities will drive increased demand for network services.

## **3.6 Remote Science Activities**

Increasingly, the OLCF is coupling its unique computational and data resources with experiment and observation data (EOD) across a broad range of scientific domains. This coupling is driven by a number of factors, including the need for large-scale simulation-based analysis, near real-time analysis requiring massive ensemble runs, and large-scale data storage resources. The OLCF is often partnering with OLCF users to implement and evaluate distributed workflow technologies and science use cases that are beneficial to users. Although the workflow of each pilot project had unique components, common requirements are emerging; many of them are being met by building upon existing scalable computing and data technologies and practices in operation at the OLCF.

Workflows and workflow systems enhance developer and scientist productivity. As an initial step toward understanding the science of workflows, OLCF conducted studies via several pilot projects by collaborating with its users in 2015 and continues this thread into 2016. The aim of this exercise was to derive a practical understanding of the current state of the theory and practice of workflow systems. Workflow requirements and expectations are documented based on discussions with several INCITE (e.g., CyberShake, Hardware/Hybrid Accelerated Cosmology Code, and Accelerated Climate Model for Energy [ACME]) and Director's Discretionary program (BEAM, BigPanDA, and Center for Nanophase Materials Science [CNMS]) projects. These discussions concluded with the observation that the current proliferation of workflow systems in response to perceived domain-specific needs of scientific workflows makes it difficult to choose a site-wide operational workflow manager, particularly for the leadership-class machines. However, there are opportunities where facilities can centralize workflow technology offerings to reduce anticipated fragmentation. This is especially true if a facility attempts to develop, deploy, and operate each and every workflow solution requested by the user community. Through these evaluations, the OLCF seeks to identify interesting intersections that are of the most value to OLCF stakeholders. As a result of their dependence on ESnet services and solutions, the following two examples are highlighted.

### **3.6.1 High-Energy and Nuclear Physics Workflows on Titan**

The largest scientific instrument in the world—the Large Hadron Collider (LHC)—operates at the CERN Laboratory in Geneva, Switzerland. The ATLAS and ALICE experiments at the LHC explore the fundamental nature of matter and the basic forces that shape our universe. The BigPanDA project<sup>2</sup> has provided the first important demonstration of the capabilities that a workload management system (WMS) can provide for improving the uptake and utilization of leadership computing facilities from both the application and systems points of view. Support from DOE ASCR and DOE HEP has led to the successful deployment of the BigPanDA workflow management

<sup>2</sup>BigPanDA: DOE ASCR and HEP -funded project (2012–2015) to extend the ATLAS workload management system (a.k.a. PanDA) beyond the Grid, in particular to clouds and supercomputers.




Figure 3.3: Titan core-hour usage per month by ATLAS, 1 May 2015 through 30 April 2016.

tools on Titan. Today, Titan is used by the ATLAS collaboration for Monte-Carlo Geant4 simulations. Working with user-collaborators from Brookhaven National Laboratory (BNL) and the University of Texas–Arlington, OLCF adapted PanDA for Titan and the OLCF environment, reusing much of the existing PanDA components and workflow.<sup>3</sup>

The project team developed and implemented a new capability in PanDA to collect information about unused worker nodes on Titan and, based on that information, adjust workload parameters to fill free, and otherwise unused, resources through intelligent backfill. Proof-of-concept tests of this mechanism, executed over a few days, achieved increased system utilization levels and provided short wait times to ATLAS and ALICE for jobs submitted to Titan via PanDA. All of this was accomplished with no negative impact on OLCF’s ability to schedule large, leadership-class jobs. Perhaps most important, Titan was fully integrated with the ATLAS PanDA-based Production and Analysis system, and today the ATLAS experiment routinely runs Monte-Carlo simulation tasks there. All operations, including data transfers to and from Titan, are transparent to the ATLAS Computing Operations team and physicists.

Titan can contribute a significant fraction of computing resources for ATLAS simulations, and Titan is regularly appearing near the top of the contributor list for wall clock consumption for ATLAS simulation jobs worldwide. (Note, simulation is the only task currently run on Titan.) Over the period 1 September 2015 through 30 April 2016, ATLAS utilization of Titan has averaged 3.5 million Titan core hours to run detector simulation jobs, using only opportunistic, backfill resources (Figure 3.3). As a result, the PanDA WMS has off-loaded an average of 7 TB of data per month to the ATLAS Tier 1 site at BNL over the ESnet. We expect this volume of data transfer to grow over the next two years as the PanDA-Titan integration continues to mature and new use-cases and performance objectives are explored.

### 3.6.2 Near Real Time Analysis of Light Source Experiment Data

Working with users Alexander Hexemer, staff scientist at Lawrence Berkeley National Laboratory (LBNL) and Craig Tull, group leader of the Science Software Systems Group at LBNL, OLCF has demonstrated the use of Titan to facilitate near real-time analysis of organic photovoltaics (OPV) using x-ray scattering at the Advanced Light Source (ALS). As data were collected at the ALS, data movement to and subsequent analysis on Titan were triggered using more than 8,000 compute nodes running HipGISAXS, a massively parallel high performance x-ray-scattering data analysis code. This analysis was used to solve an inverse problem, allowing scientists to understand the OPV

<sup>3</sup>A. Klimentov et al. “Next generation workload management system for big data on heterogeneous distributed computing,” J. Phys. Conf. Ser. 608 (1), 012040 (2015).




Figure 3.4: Fraction of largest 100 data transfers at OLCF by transfer software.

material structures from scattering data in the context of theoretical models and then drive the next stage of the experiment while the data was being collected. Moving the data to ORNL made sense because only Titan has the computational capability to run HipGISAXS in real time with the data streaming from ALS experiments. This demonstration required co-scheduling of computational resources with the ALS experiment, remote triggering of analysis running on Titan, high-performance data transfer over ESnet from the experiment end station, and near-real-time feedback of analysis results through a web portal interface.

### 3.7 Software Infrastructure

Data transfer methods must be easy to use and widely and uniformly available since the given method must be functional and installed at both ends of the data transfer. OLCF supports GridFTP and single TCP stream tools like scp.

Figure 3.4 shows the breakdown of the 100 largest data transfers at OLCF grouped by transfer tools (scp, bbcp, and Globus) for 2015 and 2016. OLCF users still perform a large fraction of transfers with the single tcp stream, scp, due to the fact it is available and functional at almost any transfer destination.

However, for the largest transfers, users are shifting to GridFTP clients such as Globus. Globus is a hosted GridFTP service that allows the use of a browser to transfer files between trusted sites called endpoints or control transfers with a command line interface. Currently eight DTNs serve the OLCF globus endpoint. As mentioned in section 1.2.4, OLCF now generates temporary X.509 certificates for GridFTP transfers with the user's existing OLCF credentials. The temporary credentials are delegated to Globus such that users do not need to manage their own personal X.509 certificates. This greatly streamlined the authentication process for GridFtp. OLCF also recently increased the time that Globus credentials and endpoints can stay active from 12 hours to 72 hours. This allows users to operate Globus-based workflows with less frequent manual re-authentication. This easier method of authentication and more workflow friendly approach has led to the dramatic increase in usage of Globus at OLCF in 2016 that can be seen in Figure 3.4.

bbcp is a multi-streaming point-to-point network file copy application created at SLAC as a tool for the BaBar collaboration. We provide bbcp as one more alternative to GridFTP. Only about 10% of users use bbcp.

## **3.8 Outstanding Issues**

### **3.8.1 The Last-Mile Problem**

OLCF does have an outstanding issue for a number of our user projects. These are user projects that are not able to transfer to their home institutions large datasets generated at OLCF because of their home institution's relatively poor connectivity to the broad-area network. OLCF does not view this as an ESnet issue necessarily, but is an issue of loss of performance and service closer to the user's institution.

## **Case Studies—ASCR Research**



## Case Study 4

# Advanced Network Services

### 4.1 Background

The DOE domain science disciplines are continuously evolving, innovating, and are reaching unprecedented complexity and scales as progress is made towards extreme-scale data and computing. ESnet traffic has increased by a factor of 10 every 48 months, and this trend has remarkably been continuing over the past 25-year history of ESnet. The result is increasingly sophisticated and large-scale science workflows that subsequently require unprecedented capabilities from the network infrastructures. We are in a particularly intense phase of this innovation cycle now with the combined emergence of big data, extreme-scale computing, cloud-based infrastructures, and sophisticated large science instruments. There is concern that current Research and Education (R&E) network infrastructures are not designed with the resource management flexibility and advanced services that will be needed by the future domain science environments.

A corresponding application evolution is occurring in the commercial space, largely driven by data center use cases. Empowered by host/compute/storage virtualization technologies, large-scale data centers have reached unprecedented levels of flexibility, scale and automation in their deployment and operation. Their network infrastructures are now recognized as the major bottleneck with respect to provisioning agility and resource management flexibility.

The R&E and the commercial sectors are both looking toward the emerging Software-Defined Networking (SDN) technologies to enable new Advanced Network Services (ANS). The objective is not just for networks to keep up with their changing application environments. The goal is for networks to provide advanced services that enable increased application level innovations. As a result, network architectures, designs, and feature sets are currently on the edge of a paradigm shift which is more significant than anything that has happened in networking since the wide-spread deployment of Dense Wavelength Division Multiplexing (DWDM) in the late 1990s. While the DWDM systems have successfully offered orders of magnitude increases in communication capacity to date, we are already exhausting the maximum capacity attainable from DWDM technologies.

The focus of this case study is: i) describing how the innovation trends in DOE domain science applications, host/compute/storage systems, and SDN/ANS will interact, and ii) identifying what the related impacts may be to ESnet and other R&E network infrastructures.

### 4.2 Network and Data Architecture

A key focus here is on the end-to-end ecosystem of resources utilized by DOE domain science communities. Most of these distributed scientific workflows rely on the DOE network resources such as the ESnet wide-area network and the individual laboratory networks. In addition, many of the science flows on ESnet include a remote side that is located at an academic or other external research organization. As a result, these collaborative science efforts also rely on the higher education network infrastructure consisting of Internet2, the regional networks, and the

academic campus infrastructures. The end-system resources located at the DOE and external facilities include science instruments, compute facilities, storage systems, DTNs, and individual researcher computers.

In order to establish the proper context for future network services discussion, a short summary of the current state of the art for advanced R&E network infrastructure and services is provided. This current state of the art revolves around the following deployed capabilities and features sets:

- 100 Gigabit/second (Gbps) Core Links: Most of the wide area links and some of the regional network links are now operating at 100 Gbps.
- Layer-3 IP Routing: Best effort routed services is still the most common method for science applications to move data. This is especially true when looking at all of the networks in the end-to-end path.
- Layer-2 Path Provisioning across Core Networks: DOE has been a pioneer in the use of Layer-2 provisioned paths for data movement across the ESnet infrastructure using OSCARS (On-Demand Secure Circuits and Advance Reservation System). This provides mechanisms for a science application to obtain an isolated Layer-2 path across the wide-area network in an automated fashion. Extending this Layer-2 path across the regional, laboratory, or campus network typically requires manual configurations, and usually the end-network does not support mechanisms to provide the Quality of Service (QoS).
- Science DMZ: At the edge of the laboratories and campuses networks, a Science DMZ is often deployed to enable high throughput flows from/to the wide-area networks.
- 10 and 40 Gbps End-system Interfaces: The current standard for network interface speed for end systems is 10 Gbps. End systems, especially purpose-designed DTNs, with 40 Gbps interfaces are becoming more common, but use of parallel data movement to multiple 10Gbps connected end systems is still the most common method for moving large amounts of data.
- End System Software and Protocols: To this network infrastructure, the domain science communities connect end systems configured with various middleware, data movement protocols, storage and compute systems, and domain science specific applications and workflows. The data movement protocols are typically based on TCP and UDP, although increased experimentation is ongoing using protocols such as iSCSI and RDMA over Ethernet.

### 4.3 Collaborators

The majority of the DOE domain science communities include collaborations with partners at academic or other external research organizations. As a result, these collaborative science efforts are multi-domain endeavors that utilize networks and other resources from DOE Laboratories, ESnet, and R&E networks consisting of Internet2, regional networks, and academic campuses. An example end-to-end flow is depicted in Figure 4.1. This diagram shows a flow that crosses DOE laboratories, university campuses, Science DMZs, multi-layer wide-area networks, and exchange points. There are also typically regional networks in these paths.

### 4.4 Network Infrastructure and Facilities

A review of the current state-of-the-art network services highlights the fact that network architectures and services have been relatively stagnant compared to the innovation that has occurred in the compute and storage system space. As a result, network infrastructures are far behind host and storage system technology with respect to dynamic resources instantiation and provisioning agility. At the same time, DOE extreme-scale science workflows are becoming increasingly distributed and complex, thereby requiring flexible, adaptive and optimizing high-performance networks. The next-generation science domain applications will require flexible and seamless integration across multiple resources, namely compute, storage, instruments, and networks. This need is motivated by several paradigm shifts in the science domain application spaces. The first shift is that science domain applications are becoming big-data driven, wherein the data considerations are greatly increasing with respect to the location, volume, mobility, and persistence requirements. Another important shift is the increasingly distributed nature of the resources (storage, compute, and instrument) needed by science workflows. While these

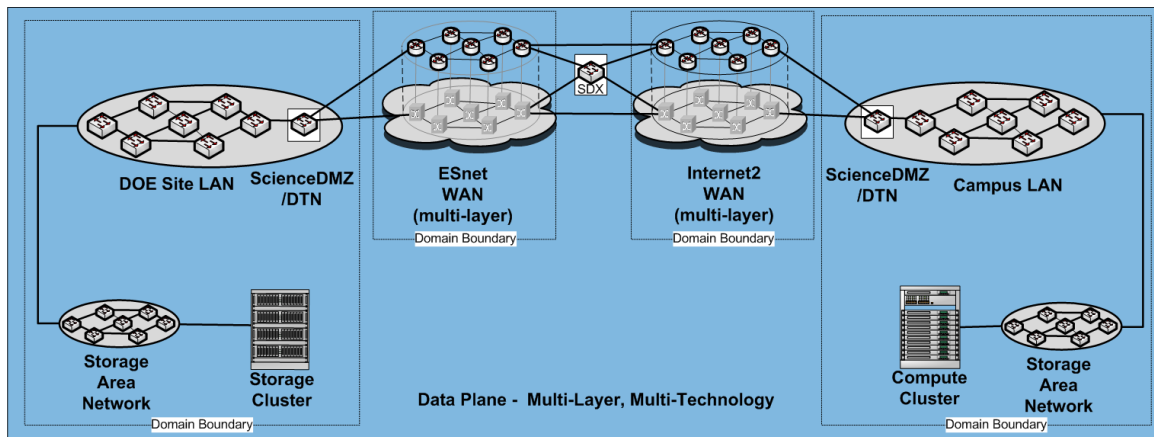


Figure 4.1: Example end-to-end flow showing collaborative organizations.

resources have always been physically distributed, the science workflow use cases are rapidly evolving to require real-time adaptations to adjust the specific resource set they are using. This will require specific capabilities from the network with regard to bandwidth, latency, and rapid re-provisioning.

In this case study we identify four key emerging technology areas we think will drive this evolution of R&E networks moving forward:

- Software Defined Networking (SDN)
- Powerful, Programmable End Systems
- Emerging Data Plane with Software Defined Networking Technologies
- Advanced Network Services (based on the integration of the above three items)

Additional discussion and information for each of these areas is presented below.

#### 4.4.1 Software Defined Networking

The R&E community is not alone in its observation that networks now represent a major bottleneck with respect to provisioning agility and resource management flexibility. Empowered by virtualization technologies, large-scale data centers have reached unprecedented levels of flexibility, scale and automation in their deployment and operation. Indeed, it is now practical to install and operate large complexes of servers and storage systems in flexible and agile configurations, using powerful automated software. Having realized great benefits from the virtualization innovation in the end-system, compute, and storage spaces, the commercial industry is now turning its attention to the network infrastructure. The emerging SDN technologies are part of a network infrastructure innovation cycle that holds an enormous potential to close this gap. SDN is expected to greatly change the way networks are constructed and operated in the future. The high-level objective is to apply virtualization concepts to networks with hopes to realize innovations similar to what has been seen in the host and storage space where these technologies resulted in new paradigms and use models.

SDN is a broad term from which many individual technologies are emerging. Below are four main concepts we believe capture the SDN core features. A brief overview of these is provided below.

- Control Plane/Data Plane Separation: The fundamental concept behind SDN is the decoupling of the network control plane from the data plane. This enables network programmability thru a controller that interacts with the data plane element forwarding engines via a southbound Application Programming Interfaces (API). OpenFlow is one example of such a southbound API. However SDN is much broader than OpenFlow, and in many cases OpenFlow may not even be part of SDN implementations.
- Northbound API: This is the API that higher level agents, such as workflow engines or multi-domain orchestration systems, would use to interact with an SDN enabled network. This API was not an early SDN

architecture definition focus area. However, it is increasingly becoming an important area of research and development.

- **Network Virtualization (NV):** These technologies are focused on two modes. The first is one where the network is truly sliced at a network dataplane technology level such as ports, DWDM, QoS protected Labels (VLANs, MPLS, others), or some other flowspace. This allows multiple “virtual” isolated networks to exist on one physical infrastructure. The other NV mode is creation of an environment where Virtual Machines (VMs) are interconnected in an agile and dynamic fashion. This type of NV utilizes tunnel mechanisms, or overlays, across an existing network infrastructure to create unique VM interconnection topologies. This mode of NV typically relies on technologies such as Virtual Extensible LAN (VXLAN) or NVGRE (Network Virtualization using Generic Routing Encapsulation) to construct the overlay networks.
- **Network Functions Virtualization (NFV):** These technologies build on the NV paradigm to add network functions such as firewalls, intrusion detection, or load balancing features in software, typically running in an elastic scalable virtual machine environment.

SDN is a nascent and promising communication network paradigm. Its foundational underpinnings are not fully understood, validated, secured, and tested. During this time of active SDN development by the commercial sector is an ideal time for DOE to evaluate how these technologies can be utilized or adapted to the DOE uses. The opportunity to leverage the commercial sector expertise and to influence the designs and standards are both available.

There are a large number of SDN concepts and individual solutions being generated in the SDN intellectual market place. ESnet has a good foundation of SDN implementation as a result of their OSCARS development and use. Some key SDN related areas that DOE may want to focus upon are:

- Definition and implementation of SDN-enabled advanced network services. Higher level agents may utilize these services to improve network resource utilization, performance, or access other advanced capabilities. This should include the development of northbound APIs to enable access to these SDN features sets.
- Evaluation and testing of emerging vendor capabilities in the SDN space to determine how they meet the needs for the R&E community, and provide feed back to the commercial sector.

#### **4.4.2 Powerful, Programmable End Systems**

There is a clear trend in the R&E and commercial infrastructures that includes placement of increasingly powerful end systems at end sites or on the edge of regional and wide-area networks. In the R&E space, this is in reaction to the observation that the limiting factors with regard to domain science researcher application end-to-end performance now reside in the end system host, storage, and application codes. The large regional and wide-area networks are generally excellent performers when tested with performance verification tools such as perfSONAR. The ScienceDMZ and DTN concepts, originated by ESnet, are part of this trend. This approach is rapidly gaining momentum as the default mechanism for end-sites to facilitate and maximize end-to-end performance. A series of National Science Foundation (NSF) programs, starting with the Campus Cyberinfrastructure—Network Infrastructure and Engineering Program (CC-NIE) in 2012 has greatly accelerated this effort for the academic regional and campus network infrastructures. In addition, many campuses are deploying local on-premise cloud systems, based on OpenStack or vendor systems, as part of a hybrid model where services move back and forth between on-premise and off-premise cloud-based infrastructures.

Looking forward, we believe that the results of this trend will be:

- Well engineered resources such as high-performance file systems and multi-tenant virtual-machine-based application hosting facilities will increasingly be attached to the Science DMZs in an effort to increase application end-to-end performance.
- Researchers will discover that an order of magnitude more application throughput is readily available as a result of connecting to these well engineered edge locations.
- These two factors will build on each other to a sufficient degree that traffic profiles across the R&E infrastructure will change from bandwidth utilization being dominated by a few large science projects to one

where hundreds of smaller researchers and projects using significant bandwidth will be more of a driving requirement.

A large university campus could easily see dozens of researchers move from occasionally using hundreds of megabits per second (Mbps) of data to frequently being able to initiate near 10 gigabit-per-second (Gbps) flows. This would likely be sufficient to require upgrades to campus and regional network infrastructures.

Several research projects are focused on automated engineering of the end systems to maximize performance without requiring network and end-system expert participation on a per-flow basis. One example are efforts to combine programmatic control and configuration non-uniform memory access (NUMA) based multi-core end systems with SDN capabilities. This would allow the mapping and correlation of network flows to end system processor cores and internal data paths. Additional information on this research area is provided below.

#### *Multicore and SDN for high-speed data movement*

Due to the fact that networks are getting faster and CPU cores are not, it is increasingly difficult for a single core to keep up with the high-speed link rates. To date, numerous efforts have been made to allow host systems to keep up with high-speed networks, through a combination of parallelism, network acceleration, and server platform improvements:

- At the application level. Various data movement tools or technologies have been developed, such as TCP-based GridFTP and BBCP, and UDP-based UDT. Parallel data transfer technologies are now widely used in large-scale data movement, providing significant improvement in aggregated data transfer throughput. These data transfer tools typically employ a multi-threaded architecture. For a data transfer, multiple threads can be spawned and assigned to different cores, with each thread handling one or multiple flows.
- At the operating system (OS) level. Major OSs (e.g., Windows and Linux) have been redesigned and parallelized to better utilize additional CPU cores. Modern network stacks can exploit cores to allow either message- or connection-based parallelism to enhance both performance and processor efficiency.
- At the system platform hardware level. Server platform performance keeps on improving. The use of NUMA systems is on the rise, due to the scalability advantage of NUMA architecture over tradition UMA (uniform memory access) architecture. New I/O technologies such as Intel QuickPath Interconnect (QPI), AMD HyperTransport (HT), and PCI Express Gen3 significantly advance server I/O bandwidth.

These efforts have been effective. A high-end host can now saturate multiple 10GE network interface cards (NICs). However, we are rapidly moving towards 40GE-connected and (eventually) 100GE-connected systems. As with previous transition to 10GE, the initial transition to 40GE and 100GE will create a fast-network, slow-host situation. Initial experiments on 40G NICs indicate that serious packet drops would occur if a faster host sends TCP data to a slower host. This is a strong signal of the “fast-network, slow-host” phenomena. We suspect that some of the aforementioned mechanisms and techniques should still be effective in the forthcoming 40/100GE host realm, while others will likely suffer from scalability limitations. To allow host systems to keep up with 40/100GE networks, we have identified a list of challenges and open questions that need to be addressed and answered, among which is the parallelism vs. I/O locality challenge on NUMA systems.

Massive parallelism is needed to handle the widening the speed mismatch between CPU cores and high-speed networks. In a multicore system, it is necessary to distribute network I/O accesses to a 40GE or 100GE NIC across many cores to maximize processing parallelism. In the case of a NUMA system, cores from different NUMA nodes may be involved. However, on NUMA systems, I/O devices (e.g., NIC and storage) are connected to processor sockets in a NUMA manner. This results in NUMA effects for transfers between I/O devices and memory banks, as well as CPU I/O accesses to I/O devices. Consequently, remote I/O accesses require more system resources than local I/O accesses on a NUMA system (Figure 4.2). Investigations show that I/O throughputs on a NUMA system can be significantly improved if applications can be placed on cores near the I/O devices they use (i.e., I/O locality) while excessive remote I/O accesses tend to degrade overall system performance. Therefore, it is also a significant challenge to optimize the tradeoff between these competing requirements

A possible solution to address the above challenge is to install multiple NICs in a NUMA system, with each NUMA node configured with at least one local NIC. Further, these installed NICs can be logically bonded as a “virtual” NIC, sharing a single IP address. In this way, each core in the NUMA system can access its local NIC(s) to send/receive packets. Therefore, remote I/O can be totally avoided, resulting in improved system performance (Figure 4.3).

Existing link bundling technologies (e.g., LACP) allows to bundle several physical ports together to form a single logical channel. However, these link bundling technologies typically cannot support I/O locality in multicore systems, due to disconnect between network applications and underlying networks. As shown in Figure 4.4, a local host transfers bulk data to a remote host. The local host is a NUMA system with two NUMA nodes, and each NUMA node is configured with a local NIC. These two NICs are bundled together to form a single logical channel. In the forward direction, the application can access the local NIC (NIC1) to send traffic, only involving local I/Os. However, in the reverse direction, the incoming traffic may be steered to the remote NIC (NIC2). The application would incur remote I/Os when accessing incoming traffic, leading to degraded performance.

In an SDN network, traffic can be steered on a per-flow basis. Therefore, we can use SDN to improve network I/Os on multicore systems. As shown in Figure 4.5, the traffic in the reverse direction can be steered to NIC1 by using SDN technology.




Figure 4.2: The parallelism vs. I/O locality on NUMA systems.




Figure 4.3: A NUMA with I/O locality.




Figure 4.4: Existing technologies cannot ensure I/O locality.

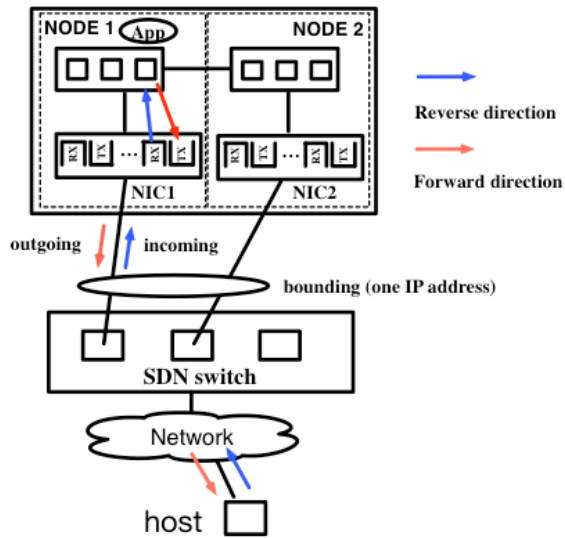


Figure 4.5: Existing technologies cannot ensure I/O locality.

### 4.4.3 Emerging Data Plane with Software Defined Networking Technologies

SDN also provides an opportunity to better manage the heterogeneous nature of the underlying network data plane. The current data plane consists a variety of technologies that includes fiber-optic wavelength-based transport, layer-2 Ethernet services, Optical Transport Network (OTN) based services, and layer-3 IP routed based services. These data plane elements are deployed in many complex environments that often include multi-layer, multi-technology, and multi-vendor configurations. Many of the core feature sets are often locked within their layer/technology/vendor regions. SDN's programmability can be leveraged to manage this complexity and facilitate the design and operation of agile networks suitable for distributed science. There are many areas that require further research and development as it relates to SDN control for multi-layer, multi-technology, and multi-vendor environments. Options based on single high-availability SDN controller versus a hierarchical or collaborative SDN controller system are active areas of research and development.

This control plane evolution also provides an opportunity to develop and integrate new data plane technologies into future network infrastructures. One of the promising areas of research and development is based on a re-evaluation of how the spectrum is allocated and utilized in optical networks, know as Elastic Optical Networks (EON).

#### *Elastic Optical Networking*

In telecommunication networks, initially, the increases in capacity demands were successfully met by deployment of wavelength division multiplexing (WDM) technology in the late 1990s. WDM systems have offered orders of magnitude increases in remarkable communication capacity from 10 Gbps to multiple terabits-per-second (Tbps). More recently, rapidly emerging new services and data centers are driving the peak link capacity demands beyond 10 Tbps. In practice, the usable bandwidth of single mode fiber communications is limited by the bandwidth of amplification technologies, which is approximately 5–10 THz for commonly used erbium-doped fiber amplifiers (EDFAs). Hence, recent optical networking advances have addressed optical communications with high-spectral efficiency beyond 1–10 b/s-Hz employing advanced modulation formats. However, it is extremely difficult to support such high-spectral efficiency and high-capacity under dynamically changing traffic conditions especially due to their sensitivity to physical layer impairments such as fiber amplifier noise, chromatic and polarization dispersion, and optical nonlinearity. Hence, the commercial telecommunication networks and the system vendors are ramping up their development and trials of EON technologies for deployment in the near future.

As Figure 4.6 illustrates, EONs utilize flexible (or elastic) spectral bandwidths for each data link without using fixed wavelength grids. For this reason, EON is also often called FlexiGrid Networks. The flexibility in spectrum allocation brings many appealing features to network operations. Current networks are designed for the worst case impairments in transmission performance and the assigned spectrum is over-provisioned. In contrast, the flexible-bandwidth networks can operate with the highest spectral efficiency and minimum bandwidth for the given traffic demand. In the case of a link failure in the network, flexible-bandwidth networks are more adaptive and likely to have spare spectrum to allocate to the re-routed signal ensuring a high-survivable restoration compared to conventional optical networks. EONs employ coherent optical orthogonal frequency division multiplexing (CO-OFDM), coherent optical WDM (CO-WDM), or Nyquist WDM technologies, and adopt various modulation formats depending on the reach. EON promises (a) to provide a large superchannel bandwidth upon demand, (b) to achieve high-spectral efficiency by eliminating stranded spectrum between the fixed grid bandwidths, (c) to support both subchannel and superchannel traffic, (d) to provide multiple data rate and modulation formats optimized for each link.

Due to the very strong interest from Internet Service Providers (ISPs) such as Verizon, Deutsche Telecom, NTT, and Google and from systems vendors such as NEC, Ericsson, Cisco, Ciena, Fujitsu, and Infinera, some of the underlying subsystems and systems are already commercially available (e.g., Flex-grid Wavelength Selective Switches (WSS), Optical Coherent transponder, coherent optical transmission systems). Initially, EON faced numerous challenges owing to lack of architectures and technologies to support bursty traffic on flexible spectrum. Under DOE and NSF support, a team at the University of California Davis (UC Davis) has recently developed many technologies, subsystems, algorithms, and testbed demonstrations for EON. In a recent UC Davis Software Defined Elastic Optical Network testbed demonstration, self-adaptive and impairment-responsive networking with observe-analyze-act cycle has been demonstrated. These studies so far aimed at achieving the following key results:

- 1 Adaptive and impairment-responsive networking optimized for each flow and for each link condition (adaptive to distance and impairment),



- 2 Interoperability with legacy WDM, CO-OFDM, CO-WDM, or Nyquist-WDM networks,
- 3 Automated, QoS-aware, and impairment-responsive network control and management,
- 4 Resilient and adaptive network operation, and
- 5 Spectrally-efficient terabit-per-second networking with superchannel and subchannel support with providing high level of availability and high throughput.

Further continuing studies are important for practical demonstrations of Software-Defined EON in the context of future ESnet with the following attributes:

- 1 Self-optimizing and automated QoS-aware, and impairment-responsive network control and management in support of big data transport upon demand,
- 2 Rapid and dynamic assignment of superchannel flows (400 Gbps and beyond),
- 3 Universal Network Access System (UNAS) edge client interface development for interoperability with legacy IP and big-data applications,
- 4 Network operating system in support of multi-domain SDN, and
- 5 Protection and restoration of EON in a single domain and a multi-domain scenarios




Figure 4.6: Comparison between (a) a standard WDM network with fixed wavelength grid and (b) a flexible bandwidth network with flexible spectrum assignment. (Flex BW TRX : Flexible Bandwidth Transmitter and Receiver.)

#### 4.4.4 Advanced Network Services

Each of the items discussed in this section represent to some degree a point technology area or solution that can contribute to improved network performance and services for users. Providing value-added services and/or performance enhancements for domain science application users and workflows should be the main objective of these research and development activities. For this reason, an overarching objective of these efforts should be to integrate the multiple technology components into systems that provide value-added functionality, i.e. advanced services.

These advanced services will sometimes be a set of tools, perhaps accessed via SDN APIs, which the network operators utilize to better manage traffic. In this manner the advanced feature set will not be transparent to the domain science user, but the improved performance will be the result. In other situations, domain science application and workflow agents will be more directly involved in the per-flow management and directly accessing the advanced network services. The interactive run-analyze-adjust-run method discussed in Section 4.5 makes an example of this.

The increasing scale and complexity of science workflows is driving a need to reevaluate the concept of end-to-end, which traditionally focused on network resources. For science workflows, the end-to-end includes all the systems between the data source and sink: SAN, LAN, Science DMZ, regional network, wide-area networks, and end-systems. This end-to-end view should be a focus for the network community as part of the development of new architectures to support big-data driven science. While this need for flexible resource integration is not

new, the technology and capability advances in each of these resource realms represents a paradigm shift where this lack of integration is now a limiting factor. While this broad consensus is indeed becoming more in focus, there are still many unknowns about what it really means to seamlessly integrate data, compute, and networking in a manner which provides the flexibility and simplicity that domain science applications require. Advanced networking infrastructures and capabilities are the cornerstone technology to enable this integration. Network attachment is a common and unifying feature around which subsequent resource integration and coordination activities can be organized. Future network infrastructures for DOE science all point to the need for networks to evolve into a flexible, agile, and programmable infrastructure, scalable to extreme-scale. Networks to be able to participate in science application workflows operations as a first class resource on the same level as compute, storage, and instrument resources.

It has been observed by some domain science researchers that the soon-to-be-routine high-performing end-systems that will appear across the R&E infrastructure can quickly be detrimental to the use of shared infrastructure. Once the knowledge and equipment for obtaining good end-system performance is more widely deployed, better coordination between networks, end systems, and end-to-end flow management will be likely be a required capability.

One goal may be to provide applications and workflows with a “deterministic performance” environment. That is, while applications will not always be able to have all the resources or end-to-end performance they would like, it should be possible for critical applications to determine what level of performance they can expect on an end-to-end basis. This will allow applications to optimize their workflows for the operational environment.

*While it is not possible, or desirable, to manage all flows in the network, it should be possible to manage “any” flow in the network.*

The other key observation about advanced network services is that they really need to be end-to-end, which means that multi-domain federated technologies are needed for these advanced network services. Figure 4.7 updates the earlier diagram to reflect distributed multi-domain SDN-based orchestration.

Resource description and service advertisement is another important capability that will be needed to enable service planning and navigation through this federated, multi-domain, multi-resource ecosystem. The OSCARS systems uses a standard known as Network Markup Language (NML) to describe ESnet network resources as part of the Network Service Interface (NSI) based provisioning. The DOE ASCR-funded Resource Aware Intelligent Network Services (RAINS) project is researching methods to extend the NSI/NML technologies to describe other types of resources which are connected to networks, such as end systems, compute, storage, and science instruments. This will allow topology computations and advanced services provisioning to consider of all the elements which constitute the end-to-end topology.

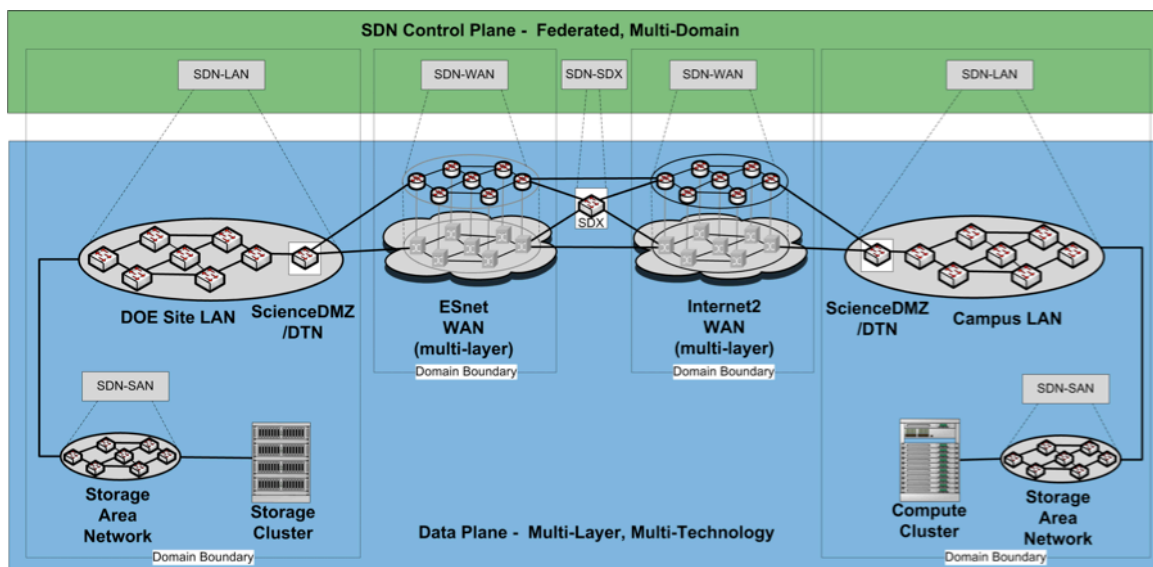


Figure 4.7: End-to-end flow with multi-domain SDN orchestration.

#### 4.4.5 Timeframes

##### Present (0-2 years):

The near term enhancements will likely revolve around efforts by the end-sites to deploy more powerful end-systems based on the Science DMZ, DTNs, and hybrid cloud solutions. The R&E networks will see an increasingly distributed and dynamic environment for high-throughput data flows. The emergence of Exchange Points powered by SDN technologies will likely emerge during this time period. These interconnect facilities are referred to as Software Defined Exchanges (SDX), and are likely to play important role in future network architectures.

##### Next 2-5 years:

Subsequent efforts will need to focus on the development of SDN-based technologies and services to better accommodate the increased numbers of high-performance end-systems and the more dynamic use of these systems as application workflows adjust to leverage these capabilities.

##### Beyond 5 years:

A longer term focus will likely be focused on a radical shift in network architectures and services based on SDN and the new data plane technologies that allow the network agility and dynamism to a degree which is compatible with the compute/storage/instrument resources and workflow operations.

### 4.5 Process of Science

The process of science is expected to change significantly as a result of the technologies and issues described in this case study. The overall theme of these changes is domain science applications and workflows driving the need for a network that is more agile from a resource management perspective, and more intelligent from an application workflow interaction perspective. We identify several generic capabilities and feature sets that we believe network infrastructures will have to play a role in providing as part of more holistic integration between network, computing systems, storage, and instrument resources.

##### Interactive and Adaptive Network Enabled Scientific Workflows:

As science workflows become more sophisticated, an ability to interact and adapt in near-real-time with the key resources is becoming increasingly important. In this context these resources are typically a workflow specific combination of compute, storage, and instrument resources. The timescales for this interaction and adaptation are typically minutes to hours and may involve some preliminary analysis of data, in order to adjust a compute process, instrument setting, or data access/storage action. This type of run-analyze-adjust-run method has always been part of the science process. However, currently this series of steps often includes long delays associated with offline data movement (i.e., FedEx) methods or over the network data transfer rates that effectively reduce the workflow to a non-real-time process. The next generations of science workflows need to transition to a true near-real-time interactive environment. This will require network infrastructures to be more flexible, adaptive, and intelligent as part of its role in providing these capabilities.

##### Intelligent Data Movements:

Each of the capability sets described require fast data movement in support of their specific focus areas. Whether the need is based on near-real-time interaction and adaptation, support of a distributed data infrastructure, or an attempt to get compute and data resources together in a timely fashion, there is a need to maximize the throughput for data movement. These types of operations are a limiting factor and key bottleneck for today's workflow. This will be an increasingly limiting factor as the volume of data, and degree of resource and scientist distribution are expected to greatly increase. Much progress has been made in this area with well-engineered edge resources such as the Science DMZ and DTNs as two example technologies. However, these capabilities must continue to improve to enable the next generation of science workflows. In particular a true end-to-end data movement paradigm needs to be developed which includes not only the wide-area network and site DMZ resources, but also extends to the local area network, storage area networks, compute, storage, instrument and data storage systems. A tighter integration and coordination between networking, storage area networks, storage systems, and project unique resources will likely be needed.

##### Smart Services for Distributed Science Big Data:

Many of the domain science applications and workflows depend on project specific data distribution, replication,

archiving, and access. This typically requires a data storage and distribution infrastructure that allows for discovery and access of project specific data. This data is typically replicated and stored in multiple data depots or repositories that are geographically distributed. Advanced network infrastructures and services that allow for increased flexibility and end-to-end capability awareness are needed to optimize these infrastructures. Intelligent cyberinfrastructure services which can facilitate decisions regarding where to store and how to access data in a workflow context are needed to enable enhancements in these areas.

#### High Performance Computing and Big Data Integration:

The timely combination of compute and data resources is a persistent problem within the domain science application community. In this context, computation may be performed on leadership class supercomputers at DOE Laboratories, distributed compute environments such as Open Science Grid (OSG), or local compute resources. Future science applications and workflows need greatly improved capabilities for a flexible integration of data and compute resources. Solutions will likely involve multiple approaches including improved mechanisms to: moving data to the compute; moving the compute to the data; improving remote access of data; feature extraction and data reduction at the sources to reduce the volume of data to be transferred. These new capabilities will leverage the other capabilities described previously in this section. These services will have to be developed in the context of the compute job execution environment and access mechanisms.

#### Real-time Interaction and Adaptation:

Real-time interaction across distance for computer-to-computer or human-driven remote control applications are expected to be of increasing interest. These interaction timescales could be an order of magnitude smaller than the near-real-time Interactive and Adaptive Workflow Support scenario discussed earlier. The types of applications and workflows for which this will be important will be based on computer-to-computer interactions or ones where there is realtime human interaction such as remote steering operations. Even within these two categories, there are orders of magnitude differences in latency and responsiveness requirements. Both of these types of realtime interactions scenarios are considered longer term requirements and goals.

#### Present (0-2 years):

The near term enhancements will likely involve efforts by the domain science applications to integrate the more powerful end-system platforms now being deployed into their workflows.

#### Next 2-5 years:

Subsequent efforts will focus on resource discovery and capability planning across multi-domain federated infrastructures.

#### Beyond 5 years:

Longer term focus will likely be focused on multi-domain resource reservation and scheduling services to move to a more real-time dynamic and interactive workflow-based operational model.

## **4.6 Remote Science Activities**

This case study focus is on the distributed domain science application workflows use of the multi-domain R&E network infrastructure. As a result, the majority of the issues and topics discussed revolve around both local and remote resources. For this reason, additional information is not provided in this section regarding remote resource utilization, as this is the default assumption in the overall case study.

## **4.7 Software Infrastructure**

From a network infrastructure and data movement perspective some of the key software in use today includes the following:

- OSCARS: Provides the ability to schedule and dynamically provision network paths. This system is run on ESnet, Internet2 AL2S, and multiple other regional and international networks.

- Internet2 OESS (Open Exchange Software Suite), FSFW (FlowSpace Firewall): Provides the ability to dynamically provision point-to-point and multi-point network topologies across the Internet2 Advanced Layer2 Service (AL2S) network.

There is other software that runs on top of these systems to provide additional value added services.

From a data movement perspective, the main software in use today by domain science applications is the GridFTP, extensions based on GridFTP, or science project unique systems.

Present (0-2 years):

The near-term enhancements will likely involve efforts by the domain applications to integrate the more powerful end-system platforms now being deployed into their workflows. This may include incorporation of current data movement protocols such as GridFTP in to these systems as well as other technologies, perhaps based on RDMA over Ethernet or iSCSI technologies.

Next 2-5 years:

Subsequent efforts will likely focus on utilization of the new SDN features sets which are expected to emerge during this time frame. It is not clear exactly what these will be, but features sets to allow scheduled, dynamic resource allocation and access to new data plane technologies are expected.

Beyond 5 years:

Longer term focus will likely be focused on multi-domain resource reservation and scheduling services to move to a more real-time dynamic and interactive workflow based operational model.

## 4.8 Cloud Services

As discussed in section 4.4, hybrid cloud environments will likely be a standard feature at most university campuses in the future. In addition, domain science applications will want to utilize public and private cloud environments as a part of an increasingly heterogeneous and distributed compute model. It is expected that there will be domain science focused cloud infrastructures deployed for this purpose, in addition to the use of commercial services such as those available from Amazon Web Services (AWS) and others.

ESnet, DOE Laboratories, and other R&E networks, will likely need to include high-performance flexible network connections to these cloud infrastructures to enable domain science workflows to easily incorporate these resources into their operations.

Some R&E networks are already providing this capability for cloud systems such as Amazon Web Services (AWS) via the direct connect service.

## 4.9 Summary and Outstanding Issues

This case study focused on the evolving end-system and network technologies, and the impact expected for the R&E network infrastructure in general, and DOE ESnet in particular.

As discussed, this is a period of rapid change in the network technologies and services. These changes are evolving in multiple dimensions including the wide spread deployment of powerful end-systems, new science instruments, and the emergence of cloud based models. The commercial and R&E communities are both looking to SDN as a new network paradigm to provide new capabilities and services. It is still early in this process and SDN means different things to different people. As a result it is difficult at this time to say definitively where SDN and Advanced Network Services are going, and exactly how ESnet should respond. However, the following broad themes and directions can be observed:

- SDN Technologies for R&E Environments: SDN is a nascent and promising communication network paradigm. Its foundational underpinnings are not fully understood, validated, secured, and tested. During this time of active SDN development by the commercial sector is an ideal time for DOE to evaluate how these technologies can be utilized or adapted to the DOE uses. The opportunity to leverage the commercial sector expertise and influence the designs and standards are both available.

- **New Data Plane and End System Technologies:** New data plane technologies such as EON and programmatic control of end-system resources such multi-core data flow management will be important parts of next generation infrastructures. The SDN paradigm provides mechanisms to amplify the utility and value of these new capabilities.
- **Advanced Network Services:** In many ways, this network architecture paradigm shift is unique, because it is happening in parallel with a similarly momentous change in application and workflow designs being driven by big data. There is a synergistic and iterative relationship between the emerging SDN network infrastructures and the big-data-driven applications. The requirements of these next-generation big science applications will drive the next-generation network infrastructures and services. The SDN-based next-generation networks and services will drive what new and innovative workflows operations domain science applications can develop. The result is that a new and important service boundary layer can be identified which sits in between the next-generation network infrastructure and the next-generation big-data-driven domain science applications. The requirements and designs for this service boundary and the associated features sets required by both the networks and the applications/workflows are currently undefined. As a result, a group of researchers will be needed who can work collaboratively across this boundary to maximize the benefit for the network operators and the domain science application and workflow developers.
- **Federated, Distributed, Multi-Domain Services:** DOE science applications workflows are generally distributed and multi-domain. A typical workflow includes resources across DOE Laboratories, wide-area networks, regional networks, and university campuses. As a result, federated and multi-domain SDN technologies will be needed. In this environment, autonomous SDN domains will need mechanisms to interact with each other, or with higher-level workflow agents in order to coordinate operations that cross multiple domains. Past experience indicates that commercial development efforts may not focus on these issues due to the business considerations associated with the multi-provider and multi-vendor Internet topology. The R&E community is well positioned to address these issues associated and it is necessary that solutions be developed in these areas.
- **SDX:** This is a concept that has been discussed as a mechanism to facilitate multi-domain services. SDXs are well-defined points of peering which may offer opportunities to realize a rich policy-based automation of network peering and services exchange. SDXs are also envisioned as a mechanism to facilitate the transition to multi-domain SDN infrastructures where non-enabled SDN networks may need to interconnect with SDN enabled networks.

Table 4.1: The following table summarizes data needs and networking requirements for ASCR’s Advanced Networking Services research area.

Key Science Drivers			Open Research Areas	
Network Infrastructure and Facilities	Process of Science	Advanced Network Services Needed	Network Research Areas	Application Research Areas
<b>0-2 years</b>				
<ul style="list-style-type: none"> <li>End-site deployments of powerful end-systems as part of ScienceDMZ, DTN, and hybrid cloud technologies become more common. Will drive new traffic profiles on R&amp;E networks.</li> <li>Prototype deployment and testing of SDN enabled R&amp;E network feature sets. New SDN enabled dataplanes, new APIs for existing network functions.</li> <li>Emergence of SD Exchanges (SDX) with basic functionality for automated exchange operations</li> </ul>	<ul style="list-style-type: none"> <li>Domain science applications to integrate the more powerful end-system platforms now being deployed as key parts of their workflows</li> <li>Domain science applications explore how to utilize SDN APIs and services as part of workflows</li> </ul>	<ul style="list-style-type: none"> <li>Well defined SDN service definitions and APIs</li> <li>Ability to interact with application agents for the purpose of flow identification and management</li> </ul>	<ul style="list-style-type: none"> <li>Development and testing of SDN features sets. Which services should be available for programmatic interaction? Which should be strictly for internal network management and optimization?</li> <li>How to accommodate many more powerful end-systems becoming available on the edges of networks? From resource planning and access management perspective.</li> <li>What are the peering and feature sets needed in future SDXs?</li> </ul>	<ul style="list-style-type: none"> <li>How to incorporate powerful end-systems connected to network edges and ScienceDMZs in to their workflows?</li> </ul>
<b>2-5 years</b>				
<ul style="list-style-type: none"> <li>Many powerful programmable end-systems, highly tuned for maximum throughput. Great increase in the number of researchers who can routinely initiate large data flows</li> <li>New SDN enabled network production deployments</li> <li>SDN networks begin to incorporate other resources embedded in their core, or at SDX facilities. This may for NFV or for application focused middlebox type of functions.</li> <li>Establishment of SD Exchanges (SDX) as key component of R&amp;E infrastructure. Rich set of SDX peering options, and services to facilitate multi-domain SDN service coordination</li> </ul>	<ul style="list-style-type: none"> <li>SDN techniques are integrated with control of other resources in application workflows such as hosts, compute, storage, instruments</li> <li>Prototypes of resource discovery and capability planning across multi-domain federated infrastructures</li> </ul>	<ul style="list-style-type: none"> <li>Well defined SDN service definitions and APIs</li> <li>Ability to interact with application agents for the purpose of flow identification and management</li> <li>Resource and service discovery mechanisms for SDN, SDX, and embedded services</li> </ul>	<ul style="list-style-type: none"> <li>How do SDN enabled networks interact with higher level agents engaged in application workflow and multi-resource orchestration?</li> <li>What are the value added functions that may be embedded in SDN networks or SDXs?</li> <li>What are the peering and feature sets needed in future SDXs?</li> </ul>	<ul style="list-style-type: none"> <li>How can SDN APIs and features be utilized to better plan, schedule, and troubleshoot end-to-end operations?</li> <li>How can SDN APIs and features be utilized to support real-time run-analyze-adjust-run methods?</li> </ul>
<b>5+ years</b>				
<ul style="list-style-type: none"> <li>SDN capabilities enabled across many infrastructures within the R&amp;E ecosystem</li> <li>Domain science and end-system integration with SDN enabled network resources is common place.</li> <li>New network architectures and feature sets based on new SDN enabled dataplane technologies</li> </ul>	<ul style="list-style-type: none"> <li>Multi-domain resource and capability discovery available across all the elements in an end to end application workflow</li> <li>Multi-domain resource planning as part of movement to more real-time and interactive workflow based operational model</li> </ul>	<ul style="list-style-type: none"> <li>Network features and services which provide agile resource management to a degree comparable with other resources in an application end-to-end path</li> </ul>	<ul style="list-style-type: none"> <li>What is the proper level of resource management needed in advanced R&amp;E networks? Is isolated individual network optimization sufficient? Is flow and resource management as part of a larger multi-domain, multi-resource topology an important objective?</li> </ul>	<ul style="list-style-type: none"> <li>What are the specific features sets that end-systems and application agents need and want from the network resources?</li> </ul>

## Case Study 5

# Bulk Data Transfer

### 5.1 Data Transfer Tools

Over the years, various data movement tools or technologies have been developed, many based on the transmission control protocol (TCP) such as Globus GridFTP [2], BBCP [6], the Secure Copy Protocol (SCP), the File Transfer Protocol (FTP); or the Universal Datagram Protocol (UDP), such as the UDP-based Data Transfer (UDT) protocol [21],<sup>1</sup> and RDMA-based File Transfer Protocol (RFTP). TCP-based tools are widely used in shared network environments, however, the TCP-based tools typically experience performance constraints on high-speed networks because the standard TCP congestion control algorithm (i.e., TCP Reno) limits the efficiency of network resource utilization. There have been numerous efforts to scale TCP over high-bandwidth networks, such as FAST TCP, High-Speed TCP (HS-TCP) [19], BIC-TCP, CUBIC-TCP, Hamilton TCP (H-TCP) and Scalable TCP (STCP) [22]. In addition, to overcome TCP's inefficiency in high-speed networks, UDP-based tools have been proposed as TCP replacements. These tools include Reliable Blast UDP and UDT. Applications can benefit from selecting among various available tools or technologies and adapting them to different networking environments. For example, in certain cases, exclusive access to the entire connection bandwidth could obviate the need for complex TCP mechanisms. Alternative transmission protocols, such as NACK-based UDT, that can make more efficient use of dedicated channels may provide a simpler, more efficient approach to data transport. Lately, RFTP utilizes the RDMA-based technology that was developed for low-latency, high-performance interconnect and extends its capability into data transfer over wide-area networking (iWARP) and Software-Defined Networks (RoCE). RFTP gains significant performance improvement due to its employed off-loading and kernel bypass technologies.

### 5.2 Data Transfer Services

In reality, bulk data transfer may encounter many abnormal conditions, including server failures, transient network failures (fiber cut, line card malfunctions, etc.), data corruption, and other errors. Therefore, bulk data transfer by hand is a human-intensive process. Researchers have developed various data transfer services on top of data transfer tools (e.g., GridFTP) to automate bulk data transfer. The HEP communities have developed several high-throughput data-transfer management systems, for example the PhEDEx [17] and ATLAS Distributed Data Management (DDM), to manage data movement for LHC experiments. The Laser Interferometer Gravitational Wave Observatory (LIGO) project developed the LIGO data Replicator. Argonne National Laboratory and University of Chicago have developed the Globus transfer service, a hosted service to which users can direct requests to transfer or synchronize files and directories between two locations. Under the covers, Globus orchestrates GridFTP transfers and handles security, monitors transfer, and restarts upon failure. With more than 10,000 active endpoints as of April 2015, Globus is an important element of the research networking ecosystem.

---

<sup>1</sup>Globus GridFTP can be configured to run UDT.



### 5.3 Data Transfer Nodes

Network engineers from ESnet have observed that the computer systems being used for wide-area data transfers perform far better if they are purposefully built, dedicated, and tuned to the function of wide-area data transfer. These systems are called DTNs. Dedicated DTNs have been deployed in DOE computing facilities and universities. See, for example:

- OLCF, [https://www.olcf.ornl.gov/kb\\_articles/employing-data-transfer-nodes/](https://www.olcf.ornl.gov/kb_articles/employing-data-transfer-nodes/)
- ALCF, <https://www.alcf.anl.gov/user-guides/data-transfer>
- NERSC, <https://www.nersc.gov/systems/data-transfer-nodes/>

### 5.4 The Science DMZ Approach

Science DMZ refers to a special DTN subnet that is typically close to a site’s network perimeter. The hardware devices, software, configuration, and policies in the Science DMZ are structured and optimized for high-performance data transfer.

The primary components of a Science DMZ are:

- Dedicated network paths for science data leveraging Access Control Lists (ACLs) for security instead of firewalls,
- High-performance DTNs with parallel data transfer tools such as GridFTP or BSCP,
- A network performance measurement system, such as perfSONAR, and
- Routers/switches with deep buffers to avoid packet drops.

DOE computing facilities are adopting or have already adopted the Science DMZ architecture and deployed multiple DTNs, improving throughput at the sites.

The NSF Office of Cyberinfrastructure has funded significant universities across the United States through its Campus Cyber-infrastructure—Network Infrastructure and Engineering Program (CC-NIE) to accelerate the deployment of Science DMZ architectures.

### 5.5 Bulk Data Transfer Related Research

In summary, these research projects can be categorized into three major areas:

Research Areas	Research Projects
Data transfer tool performance optimization	<ul style="list-style-type: none"><li>• MDTM</li><li>• GridFTP</li><li>• RFTP</li></ul>
Intelligent network service	<ul style="list-style-type: none"><li>• RAINS</li><li>• Virtual network control</li></ul>
End-to-End data transfer optimization	<ul style="list-style-type: none"><li>• Concerted Flows</li><li>• PROPER</li><li>• Synthesis of Source-to-Sink High-performance flows</li><li>• An adaptive end-to-end approach for terabit data movement optimization</li><li>• RAMSES: Robust Analytic Modeling for Science at Extreme Scales</li></ul>

## 5.6 Bulk Data Transfer Case Study

Network usage for bulk data transfer is driven by the following factors:

1. Data volumes generated by various science domains and the remote analysis requirements,
2. End-system infrastructure including the DTNs and the local network capabilities,
3. High-speed data movement tools,
4. Ease of using the data movement tools and the network, and
5. Service capabilities of transit (WAN) service providers

In recent ESnet requirements reviews for the DOE SC programs, data volumes and the need for distributed and remote analysis are increasing in almost all science domains. There is a lot of activity in terms of end-system infrastructure upgrades, thanks to the formalization the Science DMZ design pattern [12] by ESnet, and the NSF CC-NIE and follow-on grants. High-speed data movement tools such as GridFTP and BSCP have matured and there are a number of research activities focused on enhancing and building new tools and techniques to optimize transfer performance. Hosted services such as Globus [3] make data transfers much easier for users and increase performance via automated optimization, particularly when used in conjunction with DTNs and Science DMZs. The Globus Connect software makes deploying Globus endpoints straightforward.

Still, many researchers either do not use the network or do not use the network efficiently (because they use inefficient tools). A surprising number of people still use SCP rather than Globus, for example. If there is sustained funding for the above mentioned activities, more researchers can be converted to use the network (efficiently) for moving data and the usage of the network will grow tremendously. One example is to efficiently utilize the RDMA technology such as InfiniBand, iWARP, and RoCE. The RFTP tool offers significant performance improvement due to the hardware off-loading of protocol processing and software kernel zero-copy techniques.

Since the first item above has been the subject of ESnet's requirements reviews with other Office of Science program offices, here we focus on the later factors, review their current state of the art, historical trends, perceived future trends in technology, and discuss future projections.

### 5.6.1 0-2 years

#### Key science drivers

##### *Instruments, Software, and Facilities*

ALCF, OLCF, and NERSC each has a distinct set of DTNs for bulk data transfer. All data sets moved in or out of these computing facilities are transferred using the dedicated DTNs. A high-performance DTN typically features:

- One or multiple high-speed multicore processors
- High-speed storage (e.g., RAID, SSD, a parallel distributed file system)
- One or multiple high-speed NICs (e.g., 10GE, 40GE)
- High-performance motherboard that supports multiple PCIe 3 slots

Starting from 2012, NSF has been funding 15-20 universities every year to upgrade their network, and build a Science DMZ including the associated infrastructure such as dedicated DTNs for bulk data transfer. These dedicated DTNs typically run parallel data transfer tools such as GridFTP and BSCP and are configured as endpoints on the Globus Transfer service to enable "fire and forget" data transfer jobs for users. Globus transfer service logs show that the NSF CC-NIE funded universities have moved ~3PB using Globus. (A subset of the ~85 PB moved via the Globus transfer service and the >1 EB moved by Globus GridFTP servers over the past four years.) Some subset of this data (precisely how much is not known) traversed ESnet. As the deployments mature at these universities and as more universities are funded through this program, we can anticipate more traffic from these endpoints on ESnet in the subsequent years.




Figure 5.1: Usage data for Globus GridFTP servers and the Globus Transfer service.

Researchers from Fermi National Laboratory and Brookhaven National Laboratory are working on the multicore-aware data transfer middleware project, which aims to harness multicore parallelism to scale data movement toolkits at multicore systems. The project also develops several performance optimization techniques during user request preprocessing, in order to maximize the runtime utilization of system resources such as network capacity and disk bandwidth performance. If this project is successful, the MDTM middleware will be also deployed in DTNs. Other research projects mentioned in Section 5 will also develop methods to enhance performance of data transfers.

With these technologies, it is expected that bulk data transfer throughput in ESnet will significantly increase.

The left chart in Figure 5.1 shows the yearly total bytes transferred by Globus GridFTP servers (for those that report usage statistics to Globus usage collector) for the past 8 years. The percentage increase in the total bytes transferred from 2012 to 2013 is 18.6% and from 2013 to 2014 is 17.4% and the projected increase from 2014 to 2015 (based on the data transferred in the first quarter of 2015) is ~18%. The right chart in Figure 1 shows the yearly total bytes transferred by the Globus transfer service since 2011. This usage has been increasing significantly and the projected increase for 2015 (based on the data transferred in the first quarter of 2015) is ~20%.

### Process of Science

Groups awarded computing resources at Leadership Computing Facilities (e.g., OLCF, ALCF, NERSC) often transfer their data sets to the facilities at the beginning of their time slots. Also, simulation results needs to be transferred to home institution or collaboration sites.

Each large-scale science collaboration (e.g., ATLAS CMS, ESGF) typically has its own data movement plan and schedule. Large sets of data are often transferred among participating sites, and supercomputer centers. Give the ubiquity of big data recent years, the clear demarcation between HPC and HTC (physics event processing) start to interfuse. On the one hand, those traditional HPC simulations generate a large amount of data that need to be processed for steering subsequent simulations; On the other hand, typical data-intensive science programs, for example, LHC event discovery, and high-confidence data reconstruction of X-ray experiment, are fundamentally computing-intensive, and may benefit from the exascale supercomputers' energy-efficient computing cycles. ASCR-funded Next Generation Workload Management and Analysis System for Big Data to address these challenges and extend this successful workflow system from physics domain to more DOE's science programs:

#### Run ATLAS PanDA workload Management System for Big Data on Supercomputers and Cloud Computing

ATLAS LHC acquired about 160 PB so far with another 40 PB expected in 2015 and use a workload distribution system known as PanDA to coherently aggregate that data and make it available to thousands of scientists via a globally distributed computing network at 140 heterogeneous facilities around the world. The system works similar to the web, where end users can access the needed files, stored on a server in the cloud, by making service requests. The distributed resources are seamlessly integrated; there is automation and error handling that improves the user experience, and all users have access to the same resources worldwide through a single submission system. Lately, the tools of PanDA and the handling of big data migrate into the realm of supercomputers. So far, Panda makes opportunistic use of OLCF's capacity. OLCF allocates all Titan's small job slots to PanDA that can not accommodate large simulation tasks. Due to PanDA's asynchronous design, all PanDA jobs run concurrently and efficiently with Titan's programmatic large, long-duration simulation jobs. Through test, we estimates

that PanDA can utilize a total of 300M hours per year, more than 10% of the host's total capacity, which otherwise would be wasted. There are 300,000 cores in OLCF, and 30,000 cores might be utilized by PanDA every hour. They will activate a commensurate number of data transfer jobs to stream data from its host.

Currently ORNL runs PanDA services as a general batch queue and export its computing capacity to HEP, BER, and other science domain.

### **Anticipated Network Needs**

Based on the historical trends and future technology growth, we anticipate an average of ~20% increase in GridFTP traffic each year.

## **5.6.2 2-5 years**

### **Key science drivers**

#### *Instruments, Software, and Facilities*

- Science DMZ and dedicated DTNs will be widely deployed.
- High-performance DTNs with 40G NICs will be standard configuration.
- It is expected that the software-defined network (SDN) technologies will be adopted and deployed in DOE Labs in next 2-5 years. By doing so, DTNs can be better integrated with network resources. Better bulk data transfer performance can be achieved.
- Parallel data transfer tools such as GridFTP and BSCP run on DTNs.
- The MDTM middleware will also be deployed in DTNs.
- ESnet support for the Globus transfer service, as an essential element of the science networking ecosystem, will enable its continued operation and further optimization, for example to incorporate results from research projects.

#### *Process of Science*

Groups awarded computing resources at ASCR computing facilities (e.g., OLCF, ALCF, NERSC) often transfer their data sets to the facilities at the beginning of their time slots. Also, simulation results need to be transferred to home institutions or collaboration sites.

Each large-scale science collaboration typically has its own data management plan and schedule. Large sets of data are often transferred among participating sites, and supercomputer centers.

### **Anticipated Network Needs**

Based on the historical trends, perceived future adoption of Globus transfer service, research progress and technology growth, we anticipate an average ~25% increase in GridFTP traffic every year.

## **5.6.3 5+ years**

### **Key science drivers**

#### *Instruments, Software, and Facilities*

- Wider deployments of Science DMZ and dedicated DTNs.
- Advanced motherboard with PCIe-4 slots and 100GE NICs will be available in the market.

- High-performance DTNs with PCIe-4 and 100G NICs will be standard configuration.
- Parallel data transfer tools such as GridFTP and BSCP run on DTNs.
- The MDTM middleware will also be deployed in DTNs.
- The Globus transfer service will provide research institutions, network providers, and individual researchers with powerful monitoring and management capabilities, enabling far more efficient usage of high-speed networks.

#### *Process of Science*

Groups awarded computing resources at ASCR computing facilities (e.g., OLCF, ALCF, NERSC) often transfer their data sets to the facilities at the beginning of their time slots. Also, simulation results needs to be transferred to home institution or collaboration sites.

Each large-scale science collaboration typically has its own data management plan and schedule. Large sets of data are often transferred among participating sites, and supercomputer centers.

#### **Anticipated Network Needs**

Based on the historical trends, perceived future adoption of Globus Transfer service, advancements in transfer protocols and technology growth, we anticipate an average of ~30% increase in GridFTP traffic every year.

For all associated timeframes (0-2 years, 2-5 years, and 5+ years), data set sizes and network requirements for the local-area and wide-area transfer times will continue to vary based on the domain science and the applications.

## Case Study 6

# Remote Analysis and Visualization Services

### 6.1 Background

Analyzing and/or visualizing data is the key to developing an understanding of the science being conducted. In this case, data is being generated by computationally intensive applications running on ASCR compute facilities, but the user, who wishes to conduct visual data analysis and exploration, is located elsewhere. This case study focuses on the process by which remote users perform remote visualization and analysis, with an eye towards the role of the network in this process.

### 6.2 Network and Data Architecture

The term *remote and distributed analysis and visualization* (RDAV) refers to a mapping of visualization pipeline components onto distributed resources. Historically, the development of RDAV was motivated by the user's need to perform analysis on data too large to move to their local workstation or cluster, or that exceeded the processing capacity of their local resources.

From a high-level view, there are three fundamental types of bulk payload data that move between components of the visualization pipeline: "scientific data," visualization results (geometry and renderable objects), and image data. In some instances and applications, the portion of the pipeline that moves data between components is further resolved to distinguish between raw, or unprocessed, and filtered data, which could include the results of analysis processing. For simplicity, these three partitioning strategies are referred to as *send images*, *send data*, and *send geometry*, as shown in Figure 6.1.

#### 6.2.1 Send Images Partitioning

In a send-images partitioning, all processing needed to compute a final image is performed on a server, then the resulting image data is transmitted to a client. Over the years, there have been several different approaches to implement this partitioning strategy.

These include use of well established protocols, like X11 forwarding (e.g., OpenGL), and custom layer-7 protocols built into libraries and applications for the purposes of moving image data (e.g., VNC, OpenGL Vizserver, VirtualGL, Chromium Renderserver, VisIt, ParaView). These are all TCP-based approaches for moving image data.

The primary advantage of the send-images partitioning is that there is an upper bound on the amount of data that moves across the network. That upper bound is a function of image size,  $I_s$ , rather than the size of the




Figure 6.1: The partitionings of the remote and distributed visualization pipeline: send images (top), send data (middle), send geometry (bottom). Image source: High Performance Visualization, Bethel et al., CRC Press, 2012.

data set,  $D_s$ , being visualized. Typically, when  $D_s \gg I_s$ , the send-images partitioning has favorable performance characteristics, when compared to the other partitioning schemes.

Its primary disadvantage is related to its primary advantage: there is a minimum amount of data, per frame, that must move across the network. The combination of latency to produce the frame and the time required to move it over the network may be an impediment on interactive levels of performance. For example, if the user desires to achieve a 30 frame-per-second throughput rate, and each frame is 4MB in size,<sup>1</sup> then, assuming zero latency and zero time required to render the image, the network must provide a minimum of 120MB/s of bandwidth. Some systems, such as VNC, implement optimizations—compression and sending only the portion(s) of the screen that changes—to reduce the size of per-frame pixel payload.

The other disadvantage of send-images is the potential impact of network latency on interactivity, which will impose an upper bound on absolute frame rate. This upper bound may be sufficiently high on local-area networks to support interactive visualization when using the send-images approach, but may be too low on wide-area networks. For example, achieving 10 frames per second is possible only on networks having less than 100 ms of round-trip latency: for  $1000/2L \geq 10$  fps, then  $L \leq 50$  ms.

## 6.2.2 Send-Data Partitioning

The send-data partitioning aims to move scientific data from server to client for visualization and analysis processing and rendering. The scientific data may be the “source data,” prior to any processing, or it may be source data that has undergone some sort of processing, such as noise-reduction filtering, a computation of a derived field, statistical analysis, feature detection and analysis, and so forth.

In practice, the send-data approach may prove optimal when two conditions hold true: (1) the size of the source data is relatively small and will fit on the client machine, and (2) interactivity is a priority. However, as the size of scientific data grows, it is increasingly impractical to move full-resolution source data to the user’s machine for processing, since the size of data may exceed the capacity of the user’s local machine, and moving large amounts of data over the network may be cost-prohibitive.

<sup>1</sup>1024<sup>2</sup> pixels, each of which consists of RGB $\alpha$  tuples, one byte per color component, and no compression.

### 6.2.3 Send-Geometry Partitioning

In the send-geometry partitioning, the payload moving between the server and client is “drawable” content. For visualization, a class of visualization algorithms, often referred to as “mappers,” will transform scientific data, be it mesh-based or unstructured, and produce renderable geometry as an output. In a send-geometry partitioning, the server component runs data I/O and visualization algorithms, producing renderable geometry, then transmits this payload to the client for rendering.

One disadvantage of the send-geometry approach is the potential size of the renderable geometry payload. In some circumstances, the size of this payload may exceed the size of the original data set, or it may be so large as to exceed the capacity of the client to hold in memory all at once for rendering. Streaming approaches are one mechanism for accommodating rendering data sets too large for a client’s memory, yet the relatively slower network connection may be a more significant barrier.

The primary advantage of the send-geometry approach is that once the geometry content is resident in the client’s memory, the client may be capable of very high rendering frame rates. This approach may be the best when: (1) the geometry payload fits entirely within the client memory, and (2) interactive client-side rendering rates are a priority.

### 6.2.4 Application Use Models and Scenarios

#### Single Application, *Post Hoc* Use

Historically, RDAV applications have consisted of finished applications, like VisIt and ParaView, or research prototypes aimed at demonstrating some particular type of capability. Applications like VisIt and ParaView use a client-server model, where the client executes on the user’s remote machine, and the server executes, often in parallel, on the central HPC platform. Connections between the client and server are routed over a custom TCP socket that is often brokered through an SSH tunnel to accommodate site-specific authentication procedures.

Applications like VisIt and ParaView have the ability to switch between send-images and send-geometry, depending on circumstances. They both implement a strategy for minimizing the impact of network latency during interactive transformations by allowing the user to transform (e.g., rotate, scale, translate) a “wireframe model,” perhaps a bounding box, on the client machine at interactive rates, then request a full-resolution rendering from the server.

Historically, during the 1990s and 2000s, it was often the case that network requirements were couched in terms of bandwidth needed to support transmission of image data at a given rate to support interactive visualization. During those days, when 100 Mbps “fast Ethernet” was common and insufficient for those needs, custom network solutions were often the only answer. Such solutions included things like dedicated fiber-based local-area network solutions targeting high-throughput image movement.

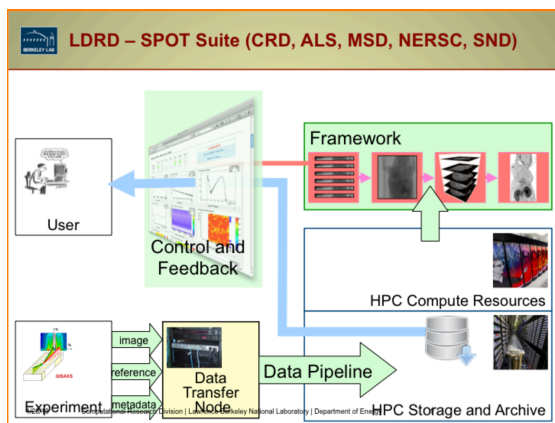
These days (in 2015), given the growth of the size of scientific data combined with the rapid growth and complexity of the underlying computational platform), there is more concern on designing and architecting algorithms and implementations to effectively tackle data of scale on large HPC platforms. The time cost of moving images tends to be much less than the time cost of analysis and visualization processing.

One typical use pattern that has emerged over the years is where a user generates a dataset on an HPC platform, saves the resulting data onto persistent storage, then invokes an application in a *post hoc* fashion to perform visual data analysis or exploration. This use pattern is still common today.

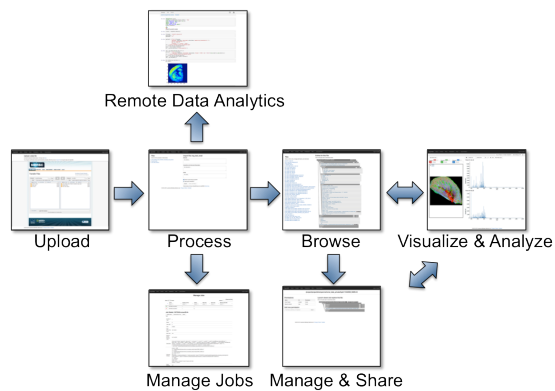
Applications like VisIt and ParaView both support some form of collaborative visualization, where multiple clients can connect to one server, and all clients can “see” the same visualization or analysis. In this case, there is only one visualization or analysis pipeline being run; within the set of clients, one is the “master” and the rest are “slaves,” though the role of master and slave can be migrated, with user coordination, from one client to another.

In the case of the current VisIt and ParaView architecture/design, it would be likely impractical to have a situation where there are multiple clients connecting to a single server, where each client is running a completely





(a) SPOT Suite. Image courtesy of Craig Tull(LBNL).



(b) OpenMSI. Image courtesy of Oliver Rübél (LBNL).

unique visualization or analysis pipeline. Part of the limitation is due to the fact that each such pipeline generates intermediate data, which will rapidly consume scarce memory resources on the HPC platform.

In more recent times, VisIt and ParaView both offer the ability to display visualization results in a web browser. Generally speaking, this type of operation is a send-images partitioning, where the server performs visualization and analysis processing, and sends a finished image to a remote client, which is a browser, over an HTTP connection.

### As part of workflows

It is increasingly the case that analysis and visualization tools/applications are part of a more elaborate processing chain that is orchestrated by a *workflow* system. While these systems tend to be highly focused on a particular science domain or problem, it is useful to mention them here.

Figure 6.2a shows the components and data flow paths of one such system, SPOT Suite. SPOT manages the collection of data acquired at ALS beamlines, sends the data to NERSC for storage, analysis and visualization processing, and returns data products and analysis/visualization results to the user, who is located at an ALS beamline.

The OpenMSI system, shown in Figure 6.2b, manages the collection of data from a mass spectrometry instrument, moves the data to NERSC for additional processing, analysis, and visualization, and then dissemination to a user or community of users.

In both of these examples, as is the case with many other similar examples, data moves through the workflow, where it undergoes many types of operations. Some operations are analysis and visualization, others are data “processing” (e.g., data reorganization, format conversion, and so forth). The partitioning of tools and data tends to be either send-data or send-images. Often, it is the case that the placement of components in the workflow is done so as to optimize for some performance characteristic, be it minimizing data movement or minimizing the “response time” from data acquisition to presentation of results to the user. Depending upon the needs of a particular workflow, some data paths may be entirely local, while others involve use of wide-area networks. These issues are largely outside the scope of this case study; they are the purview of a specific science-focused workflow case study.

### In Situ Methods

In response to the widening gap between our ability to compute data and our ability to store data for *post hoc* analysis or exploration, a new approach known as *in situ* methods has become increasingly promising and the subject of active R&D in ASCR. These methods, while potentially applicable to data from experimental and observational sources, are primarily focused on data produced by simulations run on ASCR large-scale HPC platforms.

The basic idea with *in situ* methods is that since saving full spatiotemporal resolution data to persistent storage is increasingly prohibitively expensive, then it makes sense to perform as much visualization and analysis processing as possible while simulation data is still resident in memory.

Though there exist several *in situ* framework implementations (e.g., ADIOS, Glean, ParaView/Catalyst, VisIt/libsim), the fundamental use model is more or less the same: the simulation code is coupled to some *in situ* infrastructure, which will perform analysis and visualization processing while the simulation runs. Some *in situ* frameworks (VisIt/libsim, ParaView/catalyst) are “tightly coupled” with the simulation, meaning they have the potential to alter how the simulation executes; they can be used for computational steering, for interactive debugging, and so forth. To varying degrees, *in situ* frameworks support the notion of executing potentially elaborate (and even distributed) workflows as part of their *in situ* processing. Some such configurations may involve interactions with a remotely located user. It is likely the case that such interactions would make use of primarily a send-images style partitioning, though send-data may make sense if the result is the result of some type of analysis processing.

### 6.3 Collaborators

The majority of RDAV-focused R&D, couched within the context described in the previous sections, involves use of ASCR (and Advanced Simulation and Computing, ASC) computational facilities, and science stakeholders located around the country. It is difficult to estimate the number of remote users: NERSC has  $O(3000)$  users, all of whom are remote, and of which a not insignificant number (hundreds) make use of tools like VisIt and ParaView.<sup>2</sup> Furthermore, tools like VisIt and ParaView have achieved a broad market penetration: they both are in use at NSF HPC centers like TACC and NCSA, as well as at many HPC centers abroad (e.g., CSCS, the Swiss National Supercomputing Centre).

It is reasonably safe to say that ASCR and ASC programs that fund RDAV work target deployment of technologies at HPC centers for use by science stakeholders, many of whom are remotely located. The potential impact of this work is quite broad.

### 6.4 Instruments and Facilities

Given the focus on RDAV as a service for simulation-based sciences at HPC centers, or where RDAV tools at HPC centers are brought to bear on EOD stored at those centers, then the growth of those facilities in terms of capacity is best drawn from reports directly from those facilities.

For the sake of discussion here, we can safely assume:

- Present:  $O(PF)$  class platforms.
- Next 2-5 years:  $O(10-100PF)$  class platforms.
- Beyond 5 years:  $O(100-1000PF)$  class platforms.

Given the information in the background sections, it is likely the case that most RDAV methods will continue to employ a send-data methodology for the foreseeable future. Therefore, while a significant ongoing investment in RDAV R&D focuses on scaling methods to run on larger platforms, ultimately those results will need to appear on a display somewhere, and so it may be growth in display resolution that drives an increased demand for network capacity.

Currently, typical displays use a 1080 pixel format, which is  $1920 \times 1080$  pixels. “Specialty” displays use the 4K ultra-high definition (UHD) format, which is  $3840 \times 2160$  pixels. Presumably, the 4K UHD format, which is “specialty” in the present will become commonplace in the not too distant future. An 8K UHD format is comprised of  $7680 \times 4320$  pixels, and a future standard, as yet unnamed (perhaps “16K UHD”) would use  $15,360 \times 8640$

---

<sup>2</sup>Please note, we are unable to estimate the number of users at ALCF and OLCF, or the number of such users who make use of RDAV tools like VisIt and ParaView.

pixels. Table 6.1 shows these various formats, the pixel dimensions associated with each, the size of an uncompressed image at 4 bytes/pixel, and the bandwidth required to achieve a sustained 30fps display rate. Here, we are assuming zero latency and no image compression in space or time.<sup>3</sup>

Format	Image width	Image height	Bytes/pixel	Image Size (MB)	BW in MB/s for 30fps
1k <sup>2</sup>	1024	1024	4	4.0	120.0
1080p	1920	1080	4	7.9	237.3
4K UHD	3840	2160	4	31.6	949.2
8K UHD	7680	4320	4	126.5	3796.8
16K UHD	15360	8640	4	506.3	15187.5

Table 6.1: Various display formats, image dimensions, sizes, and amount of data (in MB) moved in one second to achieve a 30fps display rate.

Another potential driver of network capacity may be an increase in the number of remote users. For example, only O(100s) of O(3000) NERSC users presently use VisIt and ParaView for RDAV. What about the remaining users? More investigation is needed to better understand their needs: are they engaging in moving data to their local machine for analysis? Do they need RDAV tools with different needs? These questions may be best explored by the centers themselves as part of their ongoing operations.

## 6.5 Process of Science

As discussed above, there appear to be three primary use modalities for RDAV technologies: strictly *post hoc*, as part of workflows (that may be distributed across resources and centers), and *in situ* methods.

- Present: primarily *post hoc*, a few workflow examples, a few *in situ* examples.
- Next 2-5 years: likely no decrease in *post hoc*, increasing numbers of workflow and *in situ* examples.
- Beyond 5 years: likely no decrease in *post hoc*, regular, production use of workflow and *in situ* methods.

## 6.6 Remote Science Activities

One significant change we anticipate in the next 5-10 years is the increasing “coupling” of instruments and experiments to HPC centers. This change will likely be motivated by the need for extreme-scale computational capacity to perform analysis and visualization on ever-larger datasets, combined with some projects’ need for specific throughput requirements: the need to use analysis results of live data to modify or alter a live-running experiment.

This change will not displace traditional center-centric use patterns of the past (*post hoc*) or present/future (workflow and *in situ* methods). It will complement those approaches, where the visualization and application tools themselves are used in new ways.

The demands of wide-area bulk data movement will likely grow significantly as part of this trend. Defining the nature of such growth is outside the scope of this case study.

## 6.7 Software Infrastructure

Given the three primary use modalities for RDAV technologies—strictly *post hoc*, as part of workflows (that may be distributed across resources and centers), and *in situ* methods—growth and evolution of software infrastructure will occur to continue to provide for those three use modalities into the future. Much of the growth and evolution

<sup>3</sup>Typical compression rates vary, depending on compressor and image characteristics, from between about 4× to about 20×.

will be driven by changes in the underlying computational architecture, as well as science-specific needs for new types of analysis and visualization methods.

## 6.8 Cloud Services

While not necessarily an RDAV-centric issue, several collaborators have discussed the idea of using cloud-based computing services to implement portions of their scientific workflow. The reasons for doing so vary, but a common theme is the need to optimize for performance or throughput in some way. For example, a given science experiment may have real-time processing and throughput requirements that, for whatever reason, may not be satisfied by a DOE HPC center, but could be using a third-party cloud-based resource. Solving problems such as this require an accurate cost estimation model (dollar, time, etc.) for all aspects of the workflow, which includes the cost of moving data, the cost of the computations (analysis, visualization), and latency associated with each workflow stage. These issues are largely beyond the scope of this case study, though future RDAV components may need the ability to provide a reasonably accurate cost estimate for their runtimes to a coordinating workflow infrastructure.

From an RDAV perspective, we anticipate there may be growth in cloud-based deployments, and as such, will need to ensure that RDAV technologies are functional in those environments. The network requirements, in terms of bandwidth and latency, are unclear. It is likely the case that a more complete understanding of those requirements would emerge from case studies that focus on specific scientific workflows.

## Case Study 7

# Remote Analysis and Visualization at Sandia National Laboratories

### 7.1 Background

Data analysis and visualization is a vital part of science through computation by providing the necessary mechanisms and tools to reason about and make conclusions from scientific data. Remote computing has long been a cornerstone of analysis and visualization in HPC facilities; the user and the machine are simply not in the same location.

As one of the three national laboratories funded by the National Nuclear Security Administration (NNSA), Sandia National Laboratories (SNL) has computing equipment specially dedicated to projects under this administration. ASCR projects at SNL can leverage some of this infrastructure, but it is sometimes more appropriate to request time on facilities at, for example, OLCF, ALCF, and NERSC.

### 7.2 Network and Data Architecture

Under the ASC program, SNL has multiple private networks at various security levels with dedicated high-speed links to Los Alamos National Laboratory (LANL) and Lawrence Livermore National Laboratory (LLNL) for speed and security.

Transferring data to/from locations outside of these secure networks (e.g., ASCR supercomputing facilities) can be difficult. Tools like Globus work well when transferring between ASCR LCFs (e.g., from ALCF to OLCF), but security restrictions prevent these tools from working. Typically we have to resort to something like rsync and wait.

Network connections for interactive tools like ParaView and VisIt can be established in a straightforward manner by using port forwarding or simply using remote client desktops. Setting up these connections, however, often requires some configuration by users. These types of applications can also be very sensitive to latency in the network.

### 7.3 Collaborators

SNL participates in many ASCR projects. Two predominant application science areas are combustion and climate. These projects are often large and involve many institutions collaborating on math, computing, and modeling.

SNL also participates in several ASCR projects involving co-design, fundamental math, and computing. The size of these collaborations varies significantly.

SNL also has a good deal of ASC work, but the computation for this work rarely leaves the confines of the three nuclear laboratories.

## 7.4 Instruments and Facilities

- Present: Transitioning from Cielo (traditional x86) to Trinity (x86 + Xeon Phi). Trinity is scheduled to have at least 80 PB storage with theoretical bandwidth of 1.45 TB/s. There will be a 3.7 PB solid-state drive (SSD) burst buffer.
- Next 2-5 years: Trinity's operational lifespan is to 2020. At the end of this time period Trinity's successor will just be rolling out.
- Beyond 5 years: It is unclear what Trinity's successor will look like, but it will likely rely heavily on "accelerator" processors to achieve the desired computational bandwidth. NVRAM will likely play a bigger role in the operation. Network and storage will improve, but not commensurately with computation.

## 7.5 Process of Science

- Present: Many computational scientific workflows still offload data to permanent disk storage and perform analysis and visualization offline later. However, the analysis and visualization community is building practical tools to work online to reduce the storage requirements. Early adopters are beginning to use these tools.
- Next 2-5 years: The usefulness of *in situ* and other online visualization tools will grow. The visualization community is working toward providing better exploratory visualization in *in situ* workflows.
- Beyond 5 years: The disparity between computation and storage will force many users to shift their workflows. Although probably not a complete replacement, *in situ* visualization will become ubiquitous in science simulation.

## 7.6 Remote Science Activities

The DOE HPC visualization tools (ParaView and VisIt) have supported remote usage since their inception in the early 2000s. These visualization servers have bursty behavior to local storage and interconnects, but the demands on the external connections are extremely low. Although other network traffic could effect the performance of the remote visualization application, it is highly unlikely for the remote visualization application to have an impact on other network traffic.

In the past it has been common for an LCF to build specialized visualization equipment, but it has been the philosophy at SNL for many years to instead leverage the computing nodes. We can run our visualization servers on the same nodes as the simulation (although it is generally necessary to have an interactive queue). This makes remote visualization feasible even when there is no special visualization or rendering hardware.

That said, it is still the case that some users opt to transfer the data to local facilities. This often simplifies configuration and removes latency problems. And since the visualization tools are designed to "run anywhere," the HPC software at the local end does not have to have specialized visualization equipment.

## 7.7 Software Infrastructure

- Present: Large scale visualization tools (i.e., ParaView and VisIt) are widely available and deployed across DOE HPC systems. These tools are deficient in leveraging accelerator architectures, this technology is rapidly becoming available. In situ visualization libraries are being integrated with more simulations and early adopters are beginning to use these tools.
- Next 2-5 years: ParaView and VisIt will be updated to leverage accelerator-type processors for the most widely used analysis functions. The usefulness of *in situ* and other inline visualization tools will grow.
- Beyond 5 years: Although probably not a complete replacement, *in situ* visualization will become ubiquitous in science simulation.

## 7.8 Cloud Services

Little if any scientific analysis for SNL is done using cloud services. “Traditional” HPC remain the most effective platform for scientific computation, analysis, and visualization.

## 7.9 Outstanding Issues

Those that manage the SNL networks understand any potential issues, however, their priority and concerns are more focused on cybersecurity than data movement. Although this helps maintain trust in protecting data, the network infrastructure team does not appreciate the difficulty in moving petabytes of data.

The respective development teams for VisIt and ParaView work hard to deploy visualization tools on the DOE supercomputing facilities and make them accessible. Often these configurations work “out of the box,” but it is difficult to cover every possible client-server. We might pursue easier and broader deployment by using remote desktop services. (TACC has reported success with this approach.)

One major issue we should address is our response to the recent OSTP memo to increase access to research. How do we make petabytes of data available to, for example, university students?

## Acknowledgment

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-AC04-94AL85000.

SAND 2015-2659 O

## Case Study 8

# Workload Management Systems

### 8.1 Background

In this case study, we explore the networking requirements placed on ASCR Leadership Computing Facilities from non-traditional distributed workloads of large SC experiments.

Workload management systems (WMS) are used by large-scale HEP, nuclear physics (NP) astro-particle and fusion energy sciences (FES) experiments to distribute and execute scientific applications on a wide class of resources. Examples of such systems include Production and Distributed Analysis System (PanDA), ALICE Environment (ALieN), DIRAC, CMS (Compact Muon Solenoid) Remote Analysis Builder (CRAB), and various experiment specific set of scripts which manage submissions to batch systems. PanDA operates at a scale of one million jobs processed daily at hundreds of grid, supercomputing and cloud sites. A variety of scientific workloads are processed, including advanced event simulations, fundamental models of particle interactions, data processing and reprocessing, and statistical analysis. WMS like PanDA typically support scientific user communities up to a few thousand people. In Figure 8.1, we show the number of jobs completed per month by PanDA at hundreds of sites, between 2011 and 2014, for the ATLAS experiment at the LHC.

Traditionally, leadership class HPC machines have not been accessible by WMS in HEP, NP, astro-particle and FES experiments. This has changed recently, with PanDA being used to process event simulations at ALCF, OLCF and NERSC (as well as at European supercomputing centers in Nordic countries, Germany and Switzerland). A limited number of HEP and NP scientific workloads have been ported to the ASCR computing facilities, and users have successfully processed event simulations at these facilities for scientific publications. Usage through these prototype and test studies has already crossed tens of millions of CPU hours. Serious efforts are underway to diversify the usage to many other workloads, to unify the usage through WMS like PanDA, to enable access to computing facilities by thousands of scientists, and to increase the usage to a hundred million hours per year. We explore the requirements on networking arising from these new patterns of HPC usage at the leadership class facilities.

### 8.2 Network and Data Architecture

We use the example of PanDA WMS in ATLAS to illustrate the typical network architecture for large-scale HEP and NP experiments. LHC computing resources were originally configured according to a strict hierarchical model. CERN was considered a Tier-0 site, since it is the primary location of experimental data. For ATLAS, 11 Tier-1 centers are deployed worldwide, connected via the LHCOPN (LHC Optical Private Network). Each Tier-1 center is connected to many Tier-2 centers (on the order of 5–10 Tier-2 sites each). The Tier-1 and Tier-2 centers are mostly connected through a Science DMZ called LHCONE (LHC Open Network Environment). In the original Models of Networked Analysis at Regional Centres (MONARC) computing model the flow of data strictly followed this hierarchy of tiered centers.






Figure 8.1: Number of jobs completed per month by PanDA for ATLAS.

During the early stages of Run 1 at the LHC, the MONARC model was relaxed in favor of a mesh model. Data transfers to/from Tier-2 centers are no longer restricted to a specific Tier-1 center, or associated Tier-2 centers. The evolution of the LHC networking architecture is shown in Figure 8.2. All LHC experiments have now migrated from a strictly hierarchical model to the mesh model. This has introduced new performance requirements for both the underlying networking infrastructure, as well as a WMS like PanDA.

Figure 8.3 shows the integrated ATLAS data transfer volume between grid sites over a period of three days. This example shows transfers after data taking had ended at the LHC. Simulation and statistical analysis are the primary activities during this period. The volume is quite high, requiring typically 100 Gbps connectivity between Tier-1 sites, and 10–40 Gbps connectivity at Tier-2 sites.

### 8.3 Collaborators

For ATLAS with 1 Tier-0 center, 10 Tier-1 centers, 100 Tier-2 centers, there are around 3000 physicists. Just in the United States, the ATLAS experiment has 1 Tier-1 center (BNL), 5 Tier-2 centers, and around 800 physicists. The ALICE experiment has around 1500 physicist and engineer collaborators.

### 8.4 Instruments and Facilities

The ATLAS experiment at the LHC is designed to explore the fundamental properties of matter and particles for the next few decades. In operation since 2009, the experiment has distributed hundreds of petabytes of data worldwide. Thousands of physicists analyze tens of millions of collisions daily, leading to more than 400 publications of new results in peer-reviewed journals.

The scale and scope of the computing challenges in ATLAS are unparalleled in the scientific community. Through the Worldwide LHC Computing Grid (WLCG), which provides access to the Open Science Grid (OSG), European




Figure 8.2: Evolution of the LHC computing model for HEP and NP experiments.




Figure 8.3: ATLAS data transfers integrated over three days.

Grid (EGI) and NorduGrid, ATLAS has seamlessly combined the resources of almost a hundred computing centers around the globe. On average, a hundred and fifty thousand jobs run simultaneously, accessing hundreds of petabytes of deployed storage worldwide and utilizing hundreds of gigabits of network bandwidth. Scientists access these computing resources transparently through the PanDA distributed computing system developed for ATLAS. The ATLAS distributed software system is highly flexible and is continuously evolving to meet the needs of thousands of physics users.

Another example of a future facility is the Belle-II experiment, which reported that all centers in the United States should be connected to ESnet. The organization of Belle-II centers will be similar to LHCONE (pairing Tier-1 and Tier-2 centers) and the data volume will be comparable to the ATLAS data volume during Run 1 (2011–2013).

Requests for bandwidth will grow, according to estimates from the two leading experiments in HEP and NP, ATLAS and ALICE by 2020. ALICE (and ATLAS) is(are) expected to have a 100-fold increase in storage and network traffic in comparison to 2014 which averaged 80 GB/s to storage, and 50 GB/s to storage respectively.

## **8.5 Process of Science**

The HEP experiments at the LHC are probing the fundamental laws of nature at the highest energies available. Until 2012, the LHC operated at a maximum 8 TeV (tera-electron volts) collision energy, which is four times the energy previously available at particle accelerators. This year, the maximum energy will increase to 13 TeV. Solving fundamental mysteries of mass and dark matter are just two examples of the exciting physics potential of the LHC experiments. A rich and varied menu of physics studies in ATLAS have led to over four hundred peer reviewed publications. Dozens of new topics will be explored in the next decade of explorations at the LHC. The NP experiments at the LHC are similarly exploring a rich menu of physics topics at the highest energy densities available at nuclei colliders.

### **8.5.1 Higgs Boson Discovery**

The 2013 Nobel Prize in Physics credited the discovery of the Higgs particle to two experiments at the LHC: ATLAS and CMS. The discovery of the Higgs particle, proposed almost 50 years ago to account for the mass of elementary particles, is a major triumph for high-throughput, network-enabled big data science.

### **8.5.2 Dark Matter Searches**

Astronomical observations 70 years ago hinted at the existence of dark matter in the universe. Subsequent observations have confirmed that the vast majority of matter in the universe is dark. ATLAS is actively searching for the fundamental particles of dark matter. Supersymmetry is a theory proposed more than 40 years ago, which may hold the clue to dark matter. Experimentally, supersymmetry has never been observed. It is a top priority for the LHC. Discovery of supersymmetry will require carefully searching through billions of events distributed worldwide, requiring high-bandwidth networking capabilities.

## **8.6 Remote Science Activities**

The large-scale HEP, NP, astro-particle and FES experiments, which form the focus of this study, have complex distributed computing and data analysis infrastructure. By necessity, they have remote facilities and users worldwide. Most of these experiments have distributed tiered computing centers, requiring high-bandwidth connections. We address the connection requirements in other sections of this document.

## 8.7 Software Infrastructure

The ATLAS experiment at the LHC uses PanDA to manage the execution of all workloads at distributed computing facilities. PanDA delivers transparency of data and processing in a distributed computing environment to ATLAS physicists. It provides execution environments for a wide range of experimental applications, automates centralized data production and processing, enables analysis activity of physics groups, supports custom workflow of individual physicists, provides a unified view of distributed worldwide resources, presents status and history of workflow through an integrated monitoring system, archives and curates all workflow, manages distribution of data as needed for processing or physicist access, integrates with the underlying networking infrastructure, and provides other features. This rich menu of features, coupled with support for heterogeneous computing environments, makes PanDA ideally suited for data-intensive science.

Through PanDA, ATLAS physicists see a single computing facility that is used to run all data processing for the experiment, even though the data centers are physically scattered all over the world. Central computing tasks (Monte Carlo simulations, or MC simulations, processing and reprocessing of LHC data, reprocessing of MC simulations, mixing and merging of data, and other tasks) are automatically scheduled and executed. Physics groups production tasks, carried out by groups of physicists of varying sizes, are also processed by PanDA. User analysis tasks, providing the majority of activities by individual physicists leading to scientific publications, are seamlessly managed.

File Transfer Service version 3 (FTS3), is the service responsible for globally distributing the majority of the LHC data across the WLCG infrastructure. FTS3 offers features and functionality that were requested by the LHC experiments and computing sites following their usage of FTS2 for Run 1. The main FTS3 features include:

- Transfer auto-tuning/adaptive optimization,
- Endpoint-centric virtual organization (VO) configuration,
- Transfer multi-hop,
- VO activity shares
- Multiple file replica support,
- Bulk deletions,
- Staging files from tapes,
- Transfer and access protocols support on top of GFAL2 plug-in mechanism (SRM, GridFTP, HTTP, xroot), and
- Session / connection reuse (GridFTP, SSL, etc), which is ideal for many small file transfers.

Recently the FTS3 team has addressed the question of HPC integration, in response to the LHC experiments showing increasing interest in using available HPC resources in addition to the usual grid ones. The difference here is that HPCs usually have unique architectures that do not match those at normal grid sites from WLCG. Work is in progress, as shown in Figure 8.4, to integrate FTS3 with PanDA and also making FTS3 capable of managing file transfers between non-grid HPC resources and standard grid storage endpoints.

FTS and PanDA will remain two main pillars of LHC data processing and data transfer for the next 5–10 years.

## 8.8 Cloud Services

ATLAS cloud R&D was started in 2009, led by a team at BNL. Currently ATLAS Distributed Computing routinely uses academic, national and commercial clouds. Amazon EC2 is one of the ATLAS “PanDA sites.” We were also the first experiment to conduct a 3-month common project with Google to demonstrate that Google cloud computing facilities can be integrated and used by ATLAS at large scale. A Tier-2 virtual center was set up in GCE and operated for 2 months. Figure 8.5 shows the 2014 ATLAS production running in clouds in the United States, Europe, Canada, and Australia.

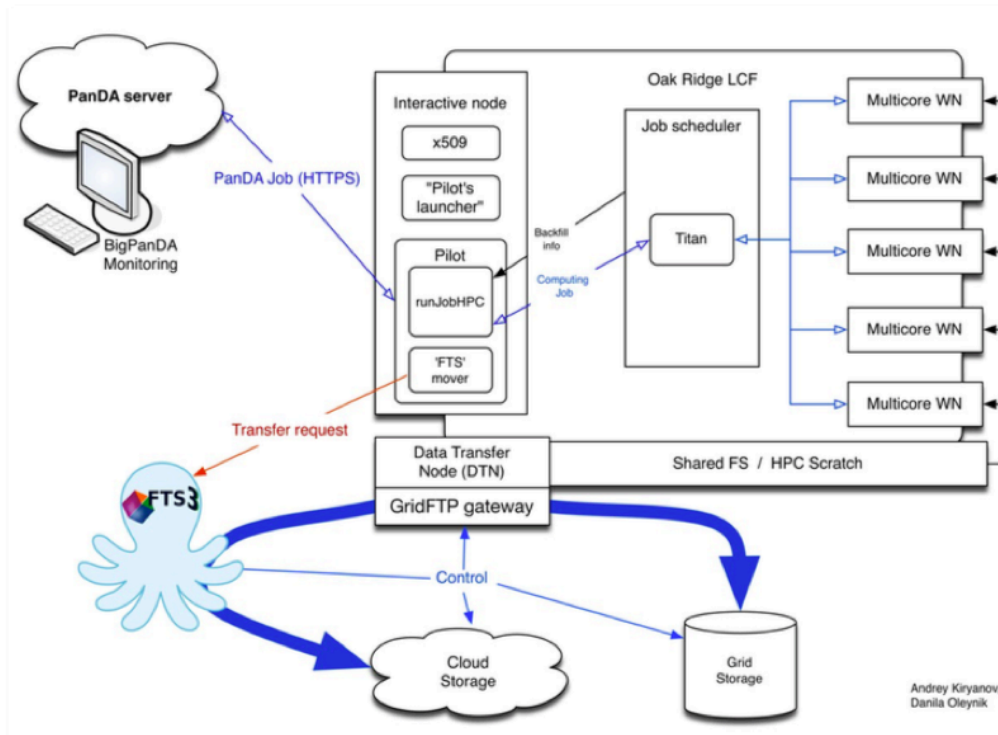


Figure 8.4: FTS3 integration with PanDA used for Titan at OLCF.




Figure 8.5: CPU consumption in seconds for ATLAS jobs running in academic and commercial clouds.

BNL received a \$200k grant from Amazon to run ATLAS workloads at large scale. The output from jobs is transferred to BNL Tier-1 WLCG center. We now use cloud infrastructure routinely (academic clouds in Canada and Australia) and usage is increasing. Currently there are no plans to use clouds as data storage, but at Amazon (and other clouds) for ATLAS production (and LHC in general) simulation and analysis will be continued. Remote data access from cloud to the grid will be an interesting and important option.

## 8.9 Conclusion

The use of ASCR Computing Facilities by large-scale experiments in HEP, NP, astro-particle and FES represents both an opportunity and a challenge. The potential for ground-breaking discoveries in the fundamental sciences is high. Rapid progress is being made in adapting WMS technology to the challenges of ASCR computing facilities. Networking improvements are needed in parallel. The scale of usage at each computing facility is equivalent to a Tier-2 center, although the workflow is different. We expect a high level of modeling and simulations to be carried out at the ASCR computing facilities. Overall, we expect the throughput to be similar to a Tier-2 facility. Therefore, 40–100 Gbps network connectivity is required between the computing facilities and the primary Tier-1 and Tier-2 facilities in the United States. For example, 100 Gbps connectivity between OLCF and BNL, NERSC and BNL, and NERSC and FNAL will be required. In the short-term 10 Gbps is required between computing facilities and Tier-2 sites, rising to 40 Gbps as computing facility usage scales up.

Typically, one hour of the GEANT4 simulation on a Titan node (16 cores) produces 200 MB output. Therefore, 10 million CPU hours will require 2 PB to be transferred at a minimum, for output. Other workloads will require additional bandwidth. As we scale up HPC facility usage by a factor of 5–10, connectivity of 100 Gbps will become a necessity. Expectations for network performance need to be raised significantly, so that collaborations do not design workflows around a historical impression of what is possible. Networking needs to be included into the resource planning process, in addition to CPU and storage, to determine how much/what is needed based on a comprehensive cost/benefit analysis (as it was stated in findings from the Snowmass Community Planning Process).

We want to add that networking information (metrics) should be taken into account by WMS and data transfer applications, network awareness should be added to workflow engines and data placement—the subject is addressed within the ASCR-, HEP- and NSF-funded Big PanDA and PanDA ANSE projects. Named-Data Networking is a new way of accessing content that is promising, without worrying about where the data is located. It will be a very important feature if it will be implemented.

## Case Study 9

# Streaming Workflows: Fusion Experimental Data Processing Workflow

### 9.1 Background

Fusion experiments provide critical information to validate and refine simulations that model complex physical processes in the fusion reactor as well as to test and postulate hypotheses. Monitoring, predicting, and mitigating instabilities are critical components of Fusion experiments. Unstable high-energy plasmas can cause serious damage to the reactor chamber, costing hundreds of millions of dollars to repair or substantial loss in productivity. Support of near real-time remote analysis workflow executions and collaboration is necessary. For the last several years, we have been researching and developing systems to support such challenging workflow scenarios through the Adaptable I/O System (ADIOS) framework. We extended ADIOS to support remote analysis workflows with WAN staging [11, 29].

### 9.2 Network and Data Architecture

Local and wide-area networks. Datasets are stored in files on file systems. Streaming data through wide-area networks; streaming experimental data in near-real-time in order to support remote analysis.

### 9.3 Collaborators

Korea Superconducting Tokamak Advanced Research (KSTAR), a fusion experiment facility located in Korea, Joint European Torus (JET) in the United Kingdom, Princeton Plasma Physics Lab (PPPL), LBNL, and ORNL.

### 9.4 Instruments and Facilities

#### 9.4.1 Present

JET and KSTAR are the current fusion experiment facilities in UK and Korea, respectively. Currently, JET, the world's largest magnetic confinement plasma physics experiment in the UK, is collecting 60 GB of diagnostic data per pulse [18]. An imaging system, called Electron Cyclotron Emission Imaging (ECEI), in KSTAR alone generates 10–100 GB of images per pulse [32]. Mostly post and batch-based data/image analysis is performed locally.




Figure 9.1: Fusion instability monitoring and mitigation workflow.

## 9.4.2 Next 2–5 years

Due to the continued advancement in sensor technologies, we expect 2–5x increases in data volume in the next 5 years. The rapid imaging system development will contribute on the data explosion. We expect the rate and spatial coverage will be 2–4x faster and wider in the next 5 years, leading 10–100x increased data volumes. Researchers need to perform near real-time analysis without restrictions on data locality. Stream-based analysis and workflows through wide area networks need to be supported.

## 9.4.3 Beyond 5 years

ITER, the next generation fusion facility being built in France, is going to start its initial plasma experiments in 2020. We expect 300–3,000 second pulses, which is 10–100 times longer than current ones produced in JET and KSTAR. Not only near real-time local/remote analysis, but also on-line feedback workflows over wide area networks will take an important role in ITER.

## 9.5 Process of Science

### 9.5.1 Present

Fusion experiments provide critical information to validate and refine simulations that model complex physical processes in the fusion reactor as well as to test and postulate hypotheses. Recent advances in sensors and imaging systems, such as sub-microsecond data acquisition capabilities and extremely fast 2D/3D imaging, allow researchers to capture very large volumes of data at high rates for monitoring and diagnostic purposes as well as post-experiment analyses. However, currently most data and image analysis is performed locally after experiments.

### 9.5.2 Next 2–5 years

The volume, velocity, and variety (data elements from thousands of sensors) of data will make it extremely challenging for researchers to analyze the data only using computational resources at experiment facilities. Researchers need ability to compose and execute workflows spanning local resources and remote large-scale high performance computing facilities. Moreover, near-real-time (NRT) analysis and decision-making is of paramount importance in fusion experiments. Monitoring, predicting, and mitigating instabilities during an experiment need strong NRT analysis capabilities. Unstable high-energy plasmas can cause serious damage to the reactor chamber, costing hundreds of millions of dollars to repair or substantial loss in productivity. A workflow to monitor, predict, and mitigate instabilities is being considered (Figure 9.1). This workflow is a multi-level workflow in that each box consists of one or more sub-workflows. Figure 9.2 shows an example workflow for analyzing 2D imaging data as part of analysis workflows for instability prediction during experiment run using a previously trained model.






Figure 9.2: Workflow for analysis of 2D image data as one of ensemble of workflows for instability prediction during experiment run.

To facilitate more efficient experimental work in fusion science, analysis workflows and underlying middleware infrastructure to execute them on local and remote resources should be able to handle thousands of streams of multi-dimensional sensor data within near-real time analysis constraints.

We have been researching and developing systems to support various data challenges in fusion science for the next 2-5 years, which involves the development of ICEE framework to support science workflows execution over the wide area network (WAN). ICEE is developed to support near-real-time streaming of experiment data to and from an experiment site and remote computing resource facilities. We focus on how we execute remote workflows over WAN with NRT requirement.

### 9.5.3 Beyond 5 years

We anticipate that fusion researches will have more remote workflows scenarios and require strong NRT supports in order to collaborate with remote scientists and exchange live feedbacks. Streaming data thorough WAN will be an important technical element in managing and executing remote workflows.

## 9.6 Remote Science Activities

Remote science activities in fusion experiments can be divided into a few categories. During the run of an experiment, collaborators at multiple sites will want to monitor the experiments and apply analyses to mitigate problems that may arise from instabilities. Between experiments, collaborators may analyze experimental results to evaluate hypotheses as well as design new experiments.

## 9.7 Software Infrastructure

### 9.7.1 Present

A variety of software systems and methods, mostly developed and maintained by research groups in house, are used locally. There is strong need to develop software and tools for stream data processing and large scale data management.

### 9.7.2 Next 2–5 years

We expect variety of stream-based signal processing and data mining methods need to be integrated in the fusion data processing workflows. Strong NRT support is also necessary. In order to keep up with high-speed data generations, intensive researches will need to be performed on data management technology for the next-generation infrastructure, such as indexing, compression, and feature detection. Hardware and network development needs to be aligned with software development for NRT support.

### 9.7.3 Beyond 5 years

We expect the complexity of software and workflow system will be highly increased. Efficient software and network infrastructures need to be developed.

## **9.8 Cloud Services**

No use of cloud services yet.

## **9.9 Outstanding Issues**

The volume, velocity, and variety (data elements from thousands of sensors) of data make it extremely challenging for researchers to analyze the data only using computational resources at experiment facilities. Researchers need ability to compose and execute workflows spanning local resources and remote large-scale high performance computing facilities. Moreover, near-real-time (NRT) analysis and decision-making is of paramount importance in fusion experiments.

## Case Study 10

# Scientific Workflows

### 10.1 Background

The scope of “scientific workflows” in ASCR is enormous, with the potential to encompass much of the science performed on DOE facilities. We focus here on a small subset of the scientific workflow space, namely that relating to analysis of data from experimental facilities and yet more specifically, light sources—with a particular emphasis on applications at the Advanced Photon Source.

#### 10.1.1 Experimental Science

The modern research environment encompasses many participants and a large and complex collection of experimental facilities, computer systems, laboratories, data stores, publications, software repositories, and other resources. Within this environment, researchers search for data; design and conduct experiments; consult with colleagues; share and publish results; develop, test, and run software; and much more. Meanwhile, dramatic (often exponential) increases in the scale and complexity of the scientific environment (e.g., amount of data, complexity of computers, size of collaborations) place extreme stresses on the abilities of individuals, institutions, and organizations to maintain effectiveness. These challenges occur across all sciences but are particularly intense within DOE due to the scale and complexity of its facilities and science.

In such settings, automation is the key to change. We become more efficient by removing time-consuming tasks from our work processes and indeed our consciousness. To that end, we must identify repeatable patterns of activities (“workflows”) and then automate those patterns in a fashion that is so intuitive, reliable, and efficient that researchers no longer need to think about them. Such workflows must be easily tailored to meet the needs of a specific facility, experiment, or project, without losing the economics of scale inherent in automation.

Workflow tools may be used to coordinate computations both within individual computer systems (*in situ* workflows) and across multiple facilities (*distributed* workflows). We focus here on the latter case, as that is where ESnet becomes important. We focus yet more specifically on the linking of experiment and computation, a particularly important (although certainly not the only) driver for distributed workflows within DOE.

DOE operates dozens of experimental facilities, of widely varying different types and scales. In order to provide further focus for this discussion, we focus in particular on the requirements of light sources (e.g., Argonne National Laboratory’s Advanced Photon Source, LBNL’s Advanced Light Source, and BNL’s National Synchrotron Light Source-II), user facilities that support thousands of users per year. Until recently, light sources assumed that data generated at the facility would be e-mailed or be transferred to portable media, and then taken back to the user’s home institution for analysis. Larger data volumes and new experimental modalities are changing this assumption. Increasingly, data is being moved over networks to local or remote facilities during experiments. Immediate computation is often required for reduction or analysis to ensure the experiment is functioning properly. The requirements placed on both workflow technologies and networks in such settings can be extreme, as a single experimental session can operate at multiple time scales; engage both distributed systems and tightly

coupled parallel computers; and require interactions with data archives and human collaborators. Timely and reliable workflow execution can enable important new experimental modalities (e.g., real-time studies of battery charge-discharge cycles). They can also enable a more efficient use of expensive facilities: without online feedback, a vastly expensive experimental session, often scheduled months in advance, can be entirely wasted: see Section 10.5.

### 10.1.2 Data-driven Workflows

Science today often requires the processing and analysis of vast amounts of data in search of postulated phenomena (e.g., climate sciences, material sciences, and bioinformatics), and the validation of core principles through the simulation of complex system behaviors and interactions (e.g., Earth System Modeling simulations). In order to support the computational and data needs of today's science, new knowledge must be gained on how to enable scientists to leverage the distributed computing infrastructure from their desktop in an accessible, reliable, and scalable way.

Even though scientists are now using workflows to express complex computations, there is still a lack of understanding of the expected workflow behavior in heterogeneous environments. It is difficult to correlate what is observed by the scientist or the workflow management system with what is happening in the infrastructure (network, storage, and compute resources). Additionally, it is difficult to predict the expected behavior of a workflow given the use of shared resources and their variable behavior.

The DOE Panorama [28] project aims to develop models of workflows to enable performance prediction, fault detection, and fault diagnosis. The project uses the Pegasus [14] workflow management system (WMS) and the ASPEN performance modeling system [30] to analyze the workflow and to develop models of expected behavior given a particular computing environment, such as an HPC system, clusters distributed over wide-area networks, or clouds. From a coupled model of the application and execution environment, decisions can then be made about resource provisioning, application task scheduling, data management within the application, etc. Panorama has identified three important application use cases involving advanced workflows that are the initial focus of our modeling efforts: parameter refinement workflows for the Spallation Neutron Source (SNS) [26], climate simulation automation for the Accelerated Climate Modeling for Energy (ACME) project, and the MG-RAST metagenome analysis [27].

## 10.2 Network and Data Architecture

The resources used in DOE experimental science span a wide range, as described in Section 10.4, as do the underlying network architectures.

Taking the APS as an example, this light source facility has more than 60 beamlines, each with its own experimental setup(s) and specialized data formats and requirements. A small compute cluster at the APS facility is used by some beamlines for processing; roughly 1km away is the Laboratory Computing Resource Center (LCRC), a modest-sized cluster; the Petrel high-speed data store (1.5 PB); the Magellan cloud resource; and the large Argonne Leadership Computing Facility (ALCF) supercomputer. Various experiment-computation pilot workflows have been developed that link APS beamlines to each of these resources, and also to remote facilities at NERSC and elsewhere. (For example, in one recent pilot involved K. Kleese Van Dam and colleagues, data was collected at APS, shipped to PNNL for reconstruction, and then viewed back at the APS beamline.) These pilots are far from production services, but they provide insights into opportunities, challenges, and future requirements.

APS networking has improved progressively over several years. Figure 10.1 shows the status as May 2015. The network between the APS and other facilities at ANL, and the APS and the outside world is far better than it was, but is still capped at 20 Gbps and, in practice, has other bottleneck links: e.g., 1 Gbps links at many beamlines. A planned initiative will connect APS and ALCF at far higher speeds via the deployment of new fiber and then the acquisition of faster hardware. The ultimate goal is 1 Tbps. Individual beamlines, meanwhile, are connected at either 1 or 10 Gbps.

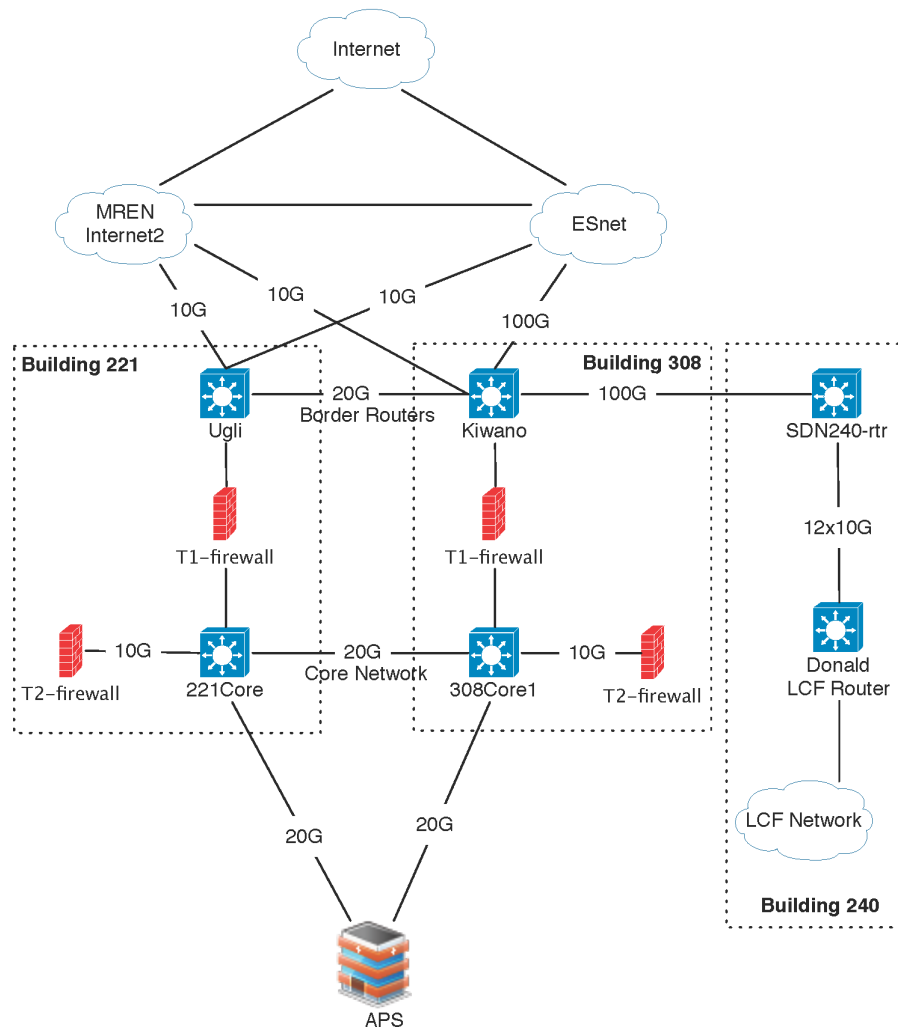



Figure 10.1: ANL network architecture, showing connectivity between the APS, the ALCF (in building 240), and the outside world.

The ALCF, like most high-end DOE computing facilities, operates a powerful array of DTNs configured with Globus GridFTP and other software to enable high-speed data transfer. The APS operates GridFTP servers, but not DTNs. An experimental deployment of a DTN at APS is currently underway.

A recent ESnet pilot project conducted by Jason Zurawski and colleagues configured a direct connection from an ANL border router to a specific beamline to demonstrate how a Science DMZ near a beamline could enable multiple gigabit-per-second rates for external transport.

Figures 10.3 and 10.4 provide some additional perspectives on network usage at DOE light sources. These figures summarize data transfers sent via the Globus transfer service from two endpoints, at APS (`aps#clutch`) and ALS (`alsuser#b1832data`), over a roughly two-year period. Figure 10.3 shows transfer destinations on United States and world maps, to provide some perspective on the geographic diversity of transfers. Figures 10.4 plot transfer rate as a function of transfer size and transfer distance. The APS and ALS servers had a 1Gbps and 10Gbps network interface, respectively, resulting in different peak bandwidths. We see considerable diversity in destinations and transfer sizes, and overall rather modest transfer performance.



9 Oct 2014  
DWL, MLW

Figure 10.2: APS network architecture showing 20 Gbps connectivity between the outside world and high-end servers such as Orthos, but with only 1 or 10 Gbps to beamlines.

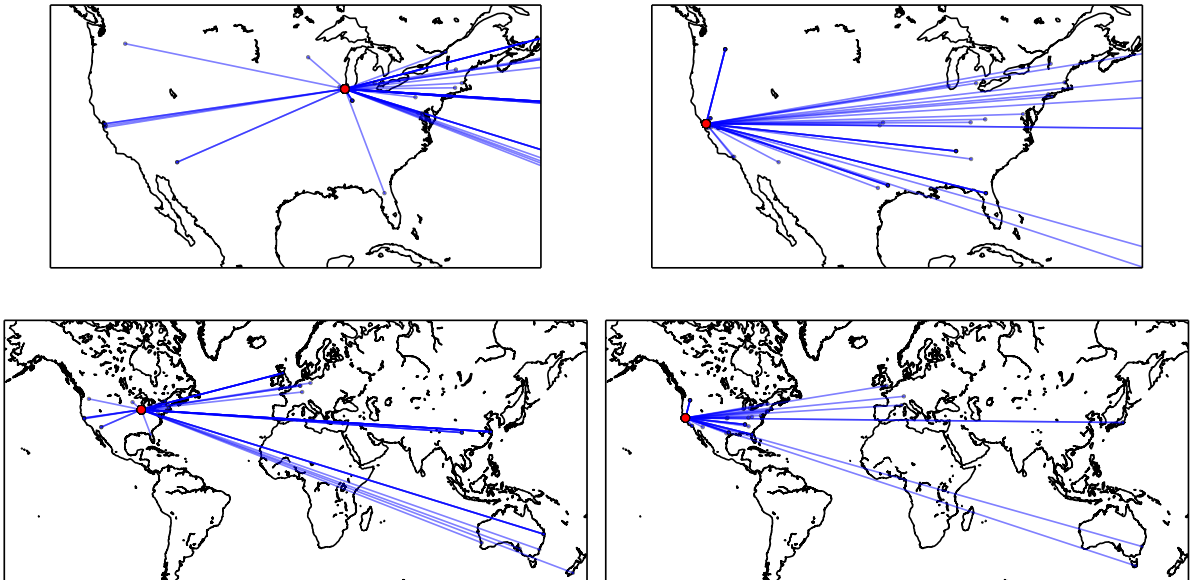


Figure 10.3: Destinations for 2,275 transfers to 119 destinations from an APS beamline (left), and 5,841 transfers to 102 destinations from an ALS beamline (right), both over a roughly two-year period.




Figure 10.4: Details of the APS transfers (on the left) and ALS transfers (on the right) from Figure 10.3. Top: Transfer rate as a function of transfer size. Bottom: Transfer rate as a function of great-circle distance from source to destination.




Figure 10.5: Self-reported data rates in 2012 from selected APS beamlines. Top: Total data per day at present, in TB—total is 168 TB/day. Bottom: Current and estimated future burst rates, in Mbps. (The light vs. dark shading can be ignored.) Source: F. De Carlo [13].

### 10.3 Collaborators

DOE experimental facilities typically support large numbers of users. The APS, for example, supports more than 5,000 users each year. An individual experiment may involve on-site beamline scientists, visiting researchers, and remote collaborators, and may make use of networks only to transfer data to remote storage at the end of an experiment or, alternatively, to transfer data during an experiment for online reconstruction and/or collaborative analysis.

### 10.4 Instruments and Facilities

As discussed in Section 10.2, the scientific workflows discussed in this case study span a wide range of resources, including experimental facilities, computational facilities, and data stores.

#### 10.4.1 Present

Different experimental facilities can generate data at vastly different rates. For example, Figure 10.5 shows estimates of current and future data rates at a number of APS beamlines. Depending on the nature of the experiment, the data rate estimates range from megabits to gigabits per second. Note that the APS can already generate data at a rate greater than that of the Large Hadron Collider at CERN.

Other facilities can achieve yet higher data rates. For example, tomographic experiments can already generate data at ~10 GB/s [25] (e.g., Dectris Eiger X, 9.6 GB/s), with the limiting factor being the camera bus: detectors can collect data internally at far greater rates. (Figure 10.5 does not reflect that data rate, as such equipment is not yet installed at APS.)

Brookhaven National Laboratory’s new National Synchrotron Light Source II (NSLS II, the sixth Office of Science light source) is expected to generate about 15 PB per year later this decade. The LCLS at SLAC can support time-resolved experiments with high spatial and temporal resolution; detectors running at maximum output can generate 1TB/hour today.



## 10.4.2 Next 2–5 years

Future developments are expected to increase data rates further. As new detectors are purchased, we can expect an order of magnitude increase in the data production rates for many APS beamlines within a few years. The performance of detectors is limited by their ability to internally pipeline photon events with data rates of circa 10 GB/s. This already challenges both storage and networks. There is every reason to believe that faster data rate detectors will be available within a few years. Further, it is also possible for a beamline to increase efficiency through the use of several detectors in parallel. We should anticipate that data rates of 100 GB/s will soon be possible. At LCLS, there are estimates that upgrades will permit a single detector to generate 1PB/hour.

## 10.4.3 Beyond 5 years

The impending APS upgrade will increase brilliance substantially allowing two and three orders of magnitude improvement in sensitivity and throughput. This will fuel development in detector technology, and in the early 2020s, upgrades are expected to increase at least another order of magnitude larger in data rates.

## 10.5 Process of Science

Many experimental facilities still practice collecting during an experiment and then analyzing the data only after the experiment has completed, either on-site or at a remote computing facility. However, a growing number of pilot projects are demonstrating the value of on-line analysis. We list three representative examples here.




Figure 10.6: Computational activities involved in a diffuse scattering experiment, including data collection and reconstruction (e.g., CCTW), identification of Bragg peaks, and the use of DIFFEV. Figure credit: J. Wozniak.

- **Single-crystal diffuse scattering** (Wozniak, Osborn, Wilde, et al.) The goal of this work is to understand defect structure in disordered materials. Wozniak et al. have developed a range of workflows, illustrated in Figure 10.6, including rapid reconstruction during individual experiment (hundreds of cores), analysis of data for peaks (thousands of cores), and evolutionary optimization for inverse modeling, using DIFFEV (100K+ BG/Q cores; Swift+OpenMP).

- **X-ray nano/microtomography** (Bicer, Gursoy, Kettimuthu, De Carlo, et al.). Rapid image reconstruction enables new applications in biological, geographical, and material science imaging, but requires large-scale on-demand computing. In one recent pilot, on-slice parallelization permitted reconstruction of a  $360 \times 2048 \times 1024$  dataset in  $\sim 1$  minute, using 32000 Blue Gene/Q cores, vs. many days on a typical cluster, enabling quasi-instant response.
- **Near-field high-energy X-ray diffraction microscopy** (Almer, Sharma, et al.). This method is used to characterize microstructures in bulk materials. Reconstruction on 10000+ Blue Gene/Q cores (Swift + MPI-IO) gives results in  $\sim 10$  minutes, vs.  $>5$  hours on an  $O(100)$  core cluster or months if data taken home. This workflow was recently used to detect errors in experiment configuration that would have otherwise resulted in a total waste of beamtime. Figure credit: J. Wozniak.

In each case, data from an experiment is shipped to a computer for analysis and results are returned to the beamline to guide further experiments. Experimental data, reconstructed data, and simulation output may also be shipped to a data store (e.g., a user's home institution, the Petrel data store) for long-term storage. The different examples vary greatly in terms of the data volumes, amount of computation, and time constraints involved.

### 10.5.1 Present

Current practices, as represented by the examples just listed, are necessarily constrained by the capabilities of existing networks; availability of computational and storage resources; and the maturity of the workflow, online reconstruction, and analysis software. Nevertheless, it already suggests or alludes to new areas of research and possibilities.

### 10.5.2 Next 2–5 years

We expect that behavioral changes (delivery of data via networks rather than physical media), policy changes (requirements for archival of all data generated at facilities), methodological changes (routine integration of large-scale computing into experimental procedures), process changes (the increasing amount of computing performed at facilities), and technological changes (new detectors) will together result in far more data movement and computing being associated with facilities than at present.

It is easy to imagine aggregate data output from all APS beamlines—currently 168 TB/day reaching 1 PB/day in this timescale. That is 12 GB/s on average, but the traffic will be bursty, and if experimentalists start wanting to link experiments with computation in near-real-time, the required burst capacity may be much larger.

The nature of the networking challenge depends on where the computing power used in experiments will be located: at APS, at ALCF, or elsewhere (e.g., NERSC).

### 10.5.3 Beyond 5 years

Computation is addictive. As better results are obtained, increased attention is placed on better algorithms and workflows. Not only can we expect to see a further order of magnitude increase in the facility's raw data rate, we can also anticipate that scientists will expect to extend the nature of their numerical processing, further increasing the demand to move raw processed datasets to large-scale computational facilities.

## 10.6 Remote Science Activities

At present, macromolecular crystallography measurements are conducted at the APS and most other light sources using controls specifically designed for remote presence. These do not demand high data transfer rates, but are very sensitive to data latency and network interruptions. We can expect that remote data collection utilization of light and neutron source facilities will grow to address the radiological, security and training demands needed

to bring users on-site, as well as decrease travel costs. One important challenge for future networks is that they must handle the previously discussed massive and bursty data transfers while at no time increasing data latency for remote experimental access.

The scientific workflows described above can involve remote facilities and collaborations in a variety of ways. For example:

- As already noted, the ability to access compute facilities at extremely high speeds will be important for new experimental modalities. In many cases, these computers will be located close to facilities and so ESnet may not be directly involved in their use. However, there will surely always be situations in which local compute power is either not present or is inadequate.
- Collaboration: share results with team members. These scenarios will become more common as substantial storage and computational resources are used more frequently in experiments. For example, R. Osborn and J. Wozniak have developed methods that enable remote analysis and visualization of diffuse scattering data via a modification to the popular NexPy data analysis software that permits partitioning of server-side data analysis from client-side interactions. This form of interaction permits interactive access to large data sets over modest bandwidth networks.
- Transfer data to remote storage.
- Coupling of experiments at multiple facilities. It is commonplace for researchers to want to characterize the same sample with multiple experimental modalities. Researchers talk about performing two (or more) experiments at the same time, e.g., at SNS and APS, so that one can guide the other.

## 10.7 Software Infrastructure

We outline in the following a few examples of software that is used in our workflow.

### 10.7.1 Present

Beamlines are controlled by EPIC software. However, this has limited applicability outside the narrow purview of controlling beamlines.

Data is collected by control computers that typically run instrument-specific data collection software and have limited local storage.

At the APS, data is transferred from beamline computers to other locations via a variety of means. Many beamlines run GridFTP servers or are configured as Globus transfer endpoints.

The following example, an expanded description of the diffuse scattering example described above, illustrates some of the software components that may be used when analyzing light source data. This workflow provides visual data analysis results to beam users while using the beam. This processing pipeline provides the user with visual experimental results in reciprocal space and real space, and results from inverse simulation and Bragg peak analysis.

As shown in Figure 10.7, the pipeline begins with the creation of raw image data on the detector computer ①. This data is transferred to ALCF resources for stable storage ② and processing. The raw data is tagged in the Globus catalog ③, along with pipeline outputs as they are produced. Then, multiple components operate on the data. If necessary, the detector background signal is subtracted from the data ④. The raw image files are merged into large NeXus files, which are visualizable in NeXpy ⑤. Then, the maximal peak and other peaks are discovered in the data ⑥. The data is transformed into real space via the Crystal Coordinate Transformation Workflow (CCTW) ⑦, which runs as a subcomputation. This subcomputation produces the visualizable real space NeXus file and produces inputs for further processing—inverse simulation-based modeling ⑧ and Bragg peak modeling ⑨. Implemented as a Swift script, it runs automatically on a parallel cluster as data is ingested, and is capable of using the whole 100-node cluster, concurrently transforming one data set per node.

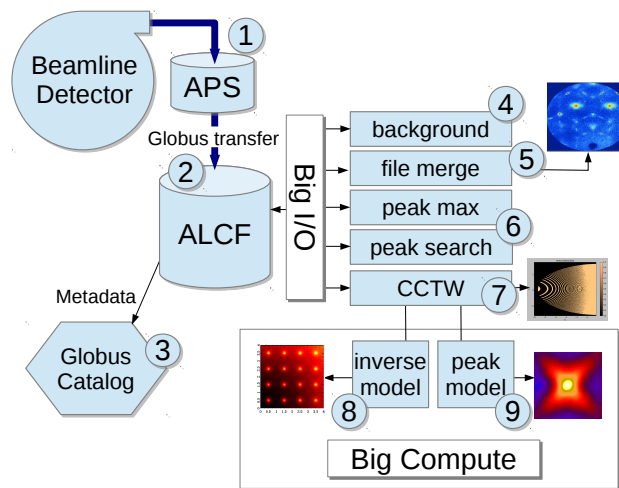


Figure 10.7: Experiment-time data analysis

CCTW, the new transformation code developed for this project, is a nearly-all-new C++ code that operates on NeXus or other HDF data sets. CCTW may be called in an automated manner as part of the pipeline. Additionally, the C++ interfaces are exposed to Swift, allowing the parallelization of CCTW itself—a feature that is critical for real-time experiment calibration, etc., as the first visualization in a run must be done quickly (in less than 10 minutes). As of early 2015, we have collected and processed about 50TB of data.

### 10.7.2 Next 2–5 years

We expect to see:

- Data collection, transfer, and management infrastructure at experimental facilities enables rapid collection of data, routing of data to on-site or remote storage, and return of analysis results to beamlines.
- A wide variety of high-performance reconstruction, analysis, and modeling and simulation codes, adapted for execution on modern high-performance computing platforms.
- Sophisticated data stores capable of storing, tracking, and enabling analysis of data produced by a wide range of experiments.

### 10.7.3 Beyond 5 years

Bigger and better.

## 10.8 Cloud Services

We see much potential for the use of cloud computing in the scientific workflows that we have just described. We identify two major use cases, quite different in their nature and implications for ESnet.

**On-demand computing and storage:** The workflows considered here have frequent needs for substantial on-demand computing and storage, and we expect such needs to increase greatly over time. The question of what sort of facility will best meet these needs has yet to be determined. However, cloud platforms, either private or public, are a potential target.

At ANL, early experiments with Magellan have been conducted and shown promise in terms of their ability to support dynamic deployment of required software and on-demand allocation of required computing and storage.

However, it remains to be seen whether DOE will invest in cloud systems at the scale required to support on-demand computing.

Public clouds are used increasingly often for science [15, 23, 24]. Whether they can be used effectively for at least some DOE on-demand workflows is largely a policy question, although connectivity between ESnet and public cloud providers will also be an issue.

**SaaS workflow:** The past five years have seen the most consequential change in the nature of workflow since the electronic computer: namely, cloud-hosted software-as-a-service (SaaS) as an enabler of large-scale automation and outsourcing.

Amazon illustrates the nature of this seismic shift. Amazon is at its heart an extreme-scale workflow automation company. It implements sophisticated workflows that encompass both consumer-facing activities (e.g., searching for, selecting, purchasing, reviewing products), and back office activities (e.g., inventory management, billing, shipping). It delivers these workflows in ways that are intuitive, reliable, and efficient for consumers. It also packages and delivers both the workflows themselves and their constituent building blocks in ways that permit easy adoption by myriad other companies. The impact of SaaS workflow on society as a whole, and especially small businesses, has been profound. Consumers and companies increasingly hand off time-consuming and error-prone activities (e.g., storing photos, booking travel, ordering products) to the likes of Amazon, who perform them far more reliably, efficiently, and cost-effectively than could any individual.

We expect to see SaaS workflow becoming increasingly important in science as well. Early examples such as Globus [20] are already improving important, providing data transfer and sharing [4, 7], identity and group management [9], and data publication [8] services to large numbers of DOE researchers. We anticipate many other labor-intensive but routine tasks being outsourced to cloud-hosted SaaS. Many such services will be concerned with research data management. The question then arises as to who should support these services. Does ESnet have a role?

## 10.9 Outstanding Issues

We identify several areas in which we see current and future challenges.

### 10.9.1 Network performance data

ASCR researchers seeking to understand, predict, and optimize the performance of scientific workflows require more information than is currently available regarding the status of the various elements involved in end-to-end network paths. Experience suggests that poor performance (e.g., see Figure 10.4) is often due to poorly understood interactions between components (e.g., LANs, WANs, firewalls, storage systems, file systems, network protocols, and competing activities) that are not typically studied together.

As part of the DOE dV/dt project [16], Pegasus has been extended to automatically capture resource usage metrics of workflow tasks. This functionality uses operating system monitoring facilities as well as system call and library call interposition to collect fine-grained profile data. To gather additional information about the infrastructure, the use of network performance monitoring tools such as perfSONAR [31] is crucial, for example, to discover *soft failures* in the network, where the network seems to be up, but is performing at just a fraction of its peak efficiency, which translates to poor performance for data-intensive workflows. However, correlating this application-level and infrastructure-level data remains challenging. It would also be beneficial to have models of network behavior that can predict the behavior of data transfers or help to automatically determine the best parameters to set.

### 10.9.2 Network infrastructure as a service

Data-driven workflows have become a centerpiece of modern computational and data-intensive science. Networked Infrastructure-as-a-Service (NIaaS) offers control interfaces for dynamic virtualization (e.g., circuits and

SDN). Networked cloud infrastructures link distributed resources into connected arrangements, sometimes referred to as *slices*, targeted at solving a specific problem. This slice abstraction is central to providing mutually isolated pieces of networked virtual infrastructure, carved out from multiple cloud and network transit providers, and built to order for guest applications like scientific workflows. The NSF ADAMANT project [1] uses the ExoGENI [5] NaaS system, which uses the ORCA control framework [10] to create mutually isolated slices of interconnected virtual infrastructure from multiple clouds and network providers. As a result the workflow can have more predictable performance.

The ability to provision resources ahead of and during workflow execution is critical to obtaining good workflow performance. However, today such provisioning is still *ad hoc*. It would be beneficial for workflow management systems to be able to reliably provision network paths and storage resources at the end points so that data can flow efficiently between workflow components in an efficient and reliable way. Another aspect of provisioning would be to provision within a specified timeline all the needed resources, network, storage, and compute.

### **10.9.3 On-demand computing**

Many experimental and data analysis workflows require on-demand computing: the ability to acquire a few hundred or in some cases many tens of thousands of cores rapidly, often with little precise information about when exactly they will be needed. While not a networking problem per se, the solution that DOE facilities ultimately make to this problem may have big implications for ESnet. If computing is performed on site, then ESnet requirements may be limited. If computing is performed remotely, then ESnet will face a major new bandwidth source.

### **10.9.4 Inconsistent end-to-end performance at facilities**

The complex internal network architecture of experimental facilities such as APS means that the end-to-end performance achieved at individual beamlines is often poor. Performance can vary widely across beamlines. ESnet could help with the design of improved network architectures.

## **10.10 Acknowledgments**

We thank Brian Toby for his review of this material, and Justin Wozniak, Mike Wilde, Hemant Sharma, Jon Almer, Ray Osborn, Ben Blaiszik, Kyle Chard, Tekin Bicer, Francesco De Carlo, and others for their input.

# References

- [1] ADAMANT: Adaptive Data-Aware Multi-domain Application Network Topologies. <http://renci.org/research/adamant>.
- [2] William Allcock et al. "The Globus Striped GridFTP Framework and Server". In: *Proceedings of the 2005 ACM/IEEE Conference on Supercomputing*. SC '05. Washington, DC, USA: IEEE Computer Society, 2005, pp. 54–. ISBN: 1-59593-061-2. DOI: 10.1109/SC.2005.72. URL: <http://dx.doi.org/10.1109/SC.2005.72>.
- [3] Bryce Allen et al. "Software As a Service for Data Scientists". In: *Commun. ACM* 55.2 (Feb. 2012), pp. 81–88. ISSN: 0001-0782. DOI: 10.1145/2076450.2076468. URL: <http://doi.acm.org/10.1145/2076450.2076468>.
- [4] Bryce Allen et al. "Software as a Service for Data Scientists". In: *Communications of the ACM* 55.2 (2012), pp. 81–88.
- [5] Ilia Baldine et al. "Exogeni: A multi-domain infrastructure-as-a-service testbed". In: *Testbeds and Research Infrastructure. Development of Networks and Communities*. Springer, 2012, pp. 97–113.
- [6] *BaBar Copy*. <http://slac.stanford.edu/~abh/bbcp/>.
- [7] Kyle Chard, Steven Tuecke, and Ian Foster. "Efficient and Secure Transfer, Synchronization, and Sharing of Big Data". In: *Cloud Computing, IEEE* 1.3 (2014), pp. 46–55.
- [8] Kyle Chard et al. "Globus Data Publication as a Service: Lowering Barriers to Reproducible Science". In: *Submitted for review*. 2015.
- [9] Kyle Chard et al. "Globus Nexus: Research Identity, Profile, and Group Management as a Service". In: *10th IEEE International Conference on e-Science*. 2014, pp. 31–38.
- [10] Jeff Chase et al. "Beyond virtual data centers: Toward an open resource control architecture". In: *Selected Papers from the International Conference on the Virtual Computing Initiative (ACM Digital Library)*. 2007.
- [11] J. Y. Choi et al., eds. *ICEE: Wide-area in Transit Data Processing Framework for Near Real-time Scientific Applications*. 4th SC Workshop on Petascale (Big) Data Analytics: Challenges and Opportunities in conjunction with SC13. 2013.
- [12] Eli Dart et al. "The Science DMZ: A Network Design Pattern for Data-intensive Science". In: *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. SC '13. Denver, Colorado: ACM, 2013, 85:1–85:10. ISBN: 978-1-4503-2378-9. DOI: 10.1145/2503210.2503245. URL: <http://doi.acm.org/10.1145/2503210.2503245>.
- [13] Francesco De Carlo et al. "Data intensive science at synchrotron based 3D x-ray imaging facilities". In: *E-Science (e-Science), 2012 IEEE 8th International Conference on*. IEEE. 2012, pp. 1–3.
- [14] Ewa Deelman et al. "Pegasus, a Workflow Management System for Science Automation". In: *Future Generation Computer Systems* 46 (2015), pp. 17–35. DOI: 10.1016/j.future.2014.10.008.
- [15] Ewa Deelman et al. "The cost of doing science on the cloud: the Montage example". In: *2008 ACM/IEEE Conference on Supercomputing*. IEEE Press, Year, p. 50.
- [16] *dV/dt: Accelerating the Rate of Progress Towards Extreme Scale Collaborative Science*. <https://sites.google.com/site/acceleratingexascale>.
- [17] Ricky Egeland, Tony Wildish, and Chih-Hao Huang. "PhEDEx Data Service". In: *Journal of Physics: Conference Series* 219.6 (2010), p. 062010. URL: <http://stacks.iop.org/1742-6596/219/i=6/a=062010>.
- [18] J. Farthing et al., eds. *Data management at JET with a look forward to ITER*. International Conference on Accelerator and Large Experimental Physics Control Systems. 2006.
- [19] S. Floyd. *HighSpeed TCP for Large Congestion Windows RFC 3649, Experimental*. 2003. URL: <http://www.faqs.org/rfcs/rfc3649.html>.

- [20] I. Foster. "Globus Online: Accelerating and democratizing science through cloud-based services". In: *IEEE Internet Computing* May/June (2011), pp. 70–73.
- [21] Y. Gu and R. Grossman. "UDT: An Application Level Transport Protocol for Grid Computing". In: *PFLD-Net'2004*. Argonne, IL.
- [22] Tom Kelly. "Scalable TCP: improving performance in highspeed wide area networks". In: *SIGCOMM Computer Communication Review* 33.2 (2003), pp. 83–91. DOI: <http://doi.acm.org/10.1145/956981.956989>.
- [23] D Lifka et al. *XSEDE cloud survey report*. Tech. rep. Technical report, National Science Foundation, USA, 2013.
- [24] Ravi K. Madduri et al. "Experiences Building Globus Genomics: A Next-Generation Sequencing Analysis Service using Galaxy, Globus, and Amazon Web Services". In: *Concurrency - Practice and Experience* (2014).
- [25] F Marone and M Stampanoni. "Regridding reconstruction algorithm for real-time tomographic imaging". In: *Journal of synchrotron radiation* 19.6 (2012), pp. 1029–1037.
- [26] TE Mason et al. "The Spallation Neutron Source in Oak Ridge: A powerful tool for materials research". In: *Physica B: Condensed Matter* 385 (2006), pp. 955–960.
- [27] *MG-RAST: Metagenome Analysis*. <https://metagenomics.anl.gov>.
- [28] *PANORAMA: Predictive Modeling and Diagnostic Monitoring of Extreme Science Workflows*. <https://sites.google.com/site/panoramaofworkflows>.
- [29] A. Sim et al., eds. *ICEE: Enabling Data Stream Processing For Remote Data Analysis Over Wide Area Networks*. Supercomputing Frontiers 2015. 2015.
- [30] Kyle L Spafford and Jeffrey S Vetter. "Aspen: a domain specific language for performance modeling". In: *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. IEEE Computer Society Press. 2012, p. 84.
- [31] Brian L Tierney. "Instantiating a Global Network Measurement Framework". In: *Lawrence Berkeley National Laboratory* (2009).
- [32] G. Yun et al., eds. *Development of KSTAR ECE imaging system for measurement of temperature fluctuations and edge density fluctuations*. Vol. 81. Review of Scientific Instruments 10. 2010.