

DOE National Laboratories' Computational Facilities – Research Workshop Report

About Argonne National Laboratory

Argonne is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC under contract DE-AC02-06CH11357. The Laboratory's main facility is outside Chicago, at 9700 South Cass Avenue, Argonne, Illinois 60439. For information about Argonne and its pioneering science and technology programs, see www.anl.gov.

DOCUMENT AVAILABILITY

Online Access: U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free at OSTI.GOV (<http://www.osti.gov/>), a service of the US Dept. of Energy's Office of Scientific and Technical Information.

Reports not in digital format may be purchased by the public from the National Technical Information Service (NTIS):

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312
www.ntis.gov
Phone: (800) 553-NTIS (6847) or (703) 605-6000
Fax: (703) 605-6900
Email: **orders@ntis.gov**

Reports not in digital format are available to DOE and DOE contractors from the Office of Scientific and Technical Information (OSTI):

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
www.osti.gov
Phone: (865) 576-8401
Fax: (865) 576-5728
Email: **reports@osti.gov**

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor UChicago Argonne, LLC, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or 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 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. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, Argonne National Laboratory, or UChicago Argonne, LLC.

DOE National Laboratories' Computational Facilities – Research Workshop Report

February 2020

Executive Summary

DOE-ASCR researchers and computational facilities staff convened at Argonne National Laboratory for a two-day workshop on April 17–18, 2019. The intent of the workshop was to understand the requirements and challenges of the DOE computational facilities and to discuss and brainstorm research directions that can have an impact on DOE computational facilities. This report documents the background information on the workshop, its purpose and structure, the challenges the computational facilities are facing, and the key research directions of interest.

Background and Purpose of the Workshop

The workshop was a continuation of activities to enhance the working relationship between the research and facility programs at the eight multiprogram laboratories (Argonne National Laboratory, Brookhaven National Laboratory, Lawrence Berkeley National Laboratory, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, Oak Ridge National Laboratory, Pacific Northwest National Laboratory, Sandia National Laboratory), with a focus on joint initiatives that will build stronger connections between these programs. The goals of the workshop were twofold: (1) provide an opportunity for facilities and research groups to engage in discussions to identify future research needs of the facilities for the 10-year timeframe, beyond the Exascale Computing Project, and (2) highlight key research directions of interest.

Workshop Structure and Deliverables

The workshop began with a short general session on the first day, with the remainder of the two days devoted to six breakout sessions, each focused on broad challenge areas identified by the facilities:

- Future of Computer Architectures
- Workflows and Scheduling
- Smart Facilities and Systems
- OS Research, System Management and Networks
- Advances in HPC Storage
- Programming Environments and Portability

The primary outcome of the workshop was a list of key *research directions of interest*, corresponding to the community-level research goals synthesized from the ideas generated during the workshop. Another important result was a discussion of cross-cutting themes

Summary of Key Research Directions of Interest

Each of the six breakouts identified three key research directions of Interest.

Future of Computer Architectures

1. Developing simulation and modeling tools to determine the optimal architecture of a heterogeneous HPC resource for a given workload
2. Analyzing application workload performance characteristics and their relationship to heterogeneous hardware components
3. Identifying optimal design and configuration of heterogeneous components for production computing

Workflows and Scheduling

1. Identifying and managing workflow patterns and resource scheduling for emerging workloads, resource types, and architectures
2. Determining primitives and abstraction levels for scheduling, managing, and executing workflows
3. Enabling end-to-end workflow monitoring, modeling, optimization, and automation

Smart Facilities and Systems

1. Characterizing workflow for optimal resource allocation to maximize system efficiency
2. Developing failure prediction to reduce system and network disruption through proactive maintenance
3. Detecting new threats for emerging HPC and high-performance networks(HPN) architectures

OS Research, System Management and Networks

1. Addressing OS, system, and network management needs for emerging workloads using shared resources
2. Providing OS, system, and network management support for increased architectural complexity and heterogeneity
3. Enabling verifiable and predictable end-to-end behavior

Advances in HPC Storage

1. Providing more automated and intelligent storage systems
2. Ensuring robust, searchable, automatable metadata
3. Addressing the storage and I/O needs of emerging workloads

Programming Environments and Portability

1. Investigating and co-designing sets of programming model abstractions for systems
2. Mapping the programming model abstractions to a set of heterogeneous hardware resources
3. Investigating methodologies, capabilities, and models for debugging, correctness, and performance

Challenges Faced by DOE Computational Facilities

Prior to the workshop, staff from the DOE computational facilities were asked to describe their key challenges in both running a facility and meeting user requirements. Below are the details from each breakout session, along with the ideas that resulted in the key research directions of interests.

Challenge Name: Future of Computer Architectures

1. **Challenge:** Deploying DOE computational resources that deliver application performance increases year after year in the post-Moore era where computer architectures will resort to specialization in order to achieve performance gains.
2. **Brief description:** In the post-Moore era computer architectures will become further specialized for varying computational characteristics. Depending on the nature of the specialization, this will give rise to the use of field programmable gate arrays (FPGAs) and coarse-grained reconfigurable arrays, domain-specific accelerators, and/or instruction set architecture extensions. DOE computational facilities need to be able to navigate this era using quantitative techniques to understand the implications and develop mitigation strategies. Solutions should include the following:
 - a. Mechanisms for determining the optimal configuration of a heterogeneous HPC resource to meet the needs of a workload.
 - b. Methods for determining the optimal pairing of a heterogeneous computing element and an application or phase of an application/workflow.
 - c. Methods for generating quantitative data to influence the vendor community to incorporate DOE needs into their future hardware designs.
3. **Underlying causes and trends:**
 - a. Diminishing returns from future fab technologies as the feature size of microelectronics nears physical limits, resulting in a proliferation of heterogeneous accelerator technologies.
 - b. Need for increasing year-on-year application performance combined with a large installed code base of DOE applications.
 - c. Larger breadth of science use cases—simulation, data analysis, machine learning (ML), and deep learning (DL).
4. **Current strategies being employed:**
 - a. Deployment of heterogeneous HPC resources at DOE centers.
 - b. Increase in the power and size of HPC systems deployed at DOE facilities.
 - c. Funded vendor co-design engagements to incorporate the needs of DOE workloads.
 - d. Centers of excellence in application performance optimization.
5. **Key research directions of interest:**
 - a. *Developing simulation and modeling tools to determine the optimal architecture of a heterogeneous HPC resource for a given workload*
 - i. **What challenge(s) does this research direction address?**
Methodologies and tools are needed that allow facilities to quantitatively understand emerging architectures and systems and make more effective procurement decisions.
 - ii. **What are the key research questions?** How do we develop methodologies and capabilities for full-system and application dynamic modeling with positive impact on performance and power? How do we improve the composability of hardware simulator components? How do we

- determine practical methodologies for system development that span the lifecycle from design to development to production?
- iii. **What is new and novel about this approach?** Incorporation of new hardware trends, capabilities to simulate at greater scale than today, inclusion of power considerations, and a focus on meeting the needs of DOE facilities.
 - iv. **What is the impact of this work on the identified challenges?** Delivery of DOE HPC platforms with superior performance and energy efficiency for DOE applications beyond commodity roadmaps. Increasing cost effectiveness of future procurements for application performance.
- b. *Analyzing application workload performance characteristics and their relationship to heterogeneous hardware components*
- i. **What challenge(s) does this research direction address?** Research is needed in order to determine the optimal match of an application for a hardware component. Also needed is better understanding of how we can generate quantitative data to influence the vendor community to incorporate the needs of the DOE into their future hardware designs.
 - ii. **What are the key research questions?** Can we develop workload characterization methodologies for complex applications on heterogeneous hardware? Can we determine the correct abstraction for heterogeneous components: should they target applications, motifs, functions (molecular dynamics vs sqrt)? Does the user need to know the component exists if it can be abstracted away?
 - iii. **What is new and novel about this approach?** Focus on applications and workloads, thereby moving a step beyond proxy apps and loop nests. Incorporation of extremely heterogeneous hardware targets, and determination of optimal heterogeneous components for a particular application or class of applications.
 - iv. **What is the impact of this work on the identified challenges?** Provision of accurate, actionable, representative data to vendors, facilities, and software developers. Enablement of quantitative decision making informed by a greater understanding of application performance characteristics. Determination of the optimal pairing of a heterogeneous computing element and an application, or phase of an application or workflow.
- c. *Identifying optimal design and configuration of heterogeneous components for production computing*
- i. **What challenge(s) does this research direction address?** Researchers need to determine how to optimally explore the design space of heterogeneous components for processors, memory, networking, and/or storage for production computing,
 - ii. **What are the key research questions?** Which heterogeneous components should be developed and/or explored? How can we determine the balance and coupling between heterogeneous components from the node architecture through the system architecture? How can we influence

industry with quantitative analysis to produce heterogeneous components that meet the needs of the DOE?

- iii. **What is new and novel about this approach?** Addressing the emerging heterogeneous architectures; integration of approaches to design space exploration.
- iv. **What is the impact of this work on the identified challenges?** Greater alignment between the heterogeneous components deployed and the needs of DOE applications

Challenge Name: Workflows and Scheduling

1. **Challenge:** Enabling rich, performant, and reproducible workflows for science while effectively exploiting co-schedulable resources both within a facility and across facilities, such as those between experimental and computational facility, together with collaborative workspaces and interaction modalities to facilitate scientific discovery.
2. **Brief description:** We are witnessing an increasing use of machine learning and analytics and the coupling of these techniques with large-scale simulations for scientific discovery. Science workflows and campaigns consist of multiple stages, each with diverse and varying requirements. At the same time, system resources are increasingly being architected with the ability to support finer-grained provisioning and to be co-scheduled. We are also witnessing new modalities to interact with supercomputers, including interactive notebooks and human-in-the-loop experiments, and richer modes to facilitate collaborative teams are needed. Mechanisms to orchestrate workflows and effectively schedule these at a facility and across facilities will be critical for scientific productivity.
3. **Underlying causes and trends:**
 - a. Workflows that originate outside an HPC facility, including real-time and near-time steering of experiments and simulations.
 - b. Limited ability of existing HPC schedulers to co-schedule resources or handle streaming workloads. Often inflexible interfaces/API for coupling complex workflows, with few appropriate abstraction layers.
 - c. Complexity of workflows and HPC systems, along with limited monitoring tools and a lack of integrated tools or capabilities that can assess the performance of an end-to-end workflow, thus making it difficult to ensure workflow resiliency as well as throughput (such as is available for a single executable application with TAU and the HPCToolkit).
 - d. Increasing trends in systems architectures with the capability to provision resources at finer granularity.
 - e. Need to provide improved quality-of-service guarantees for workflows with faster turnaround time requirements; growing need to facilitate complex science workflows such as ensemble-based workflows, UQ, V&V, in situ analysis and visualization, analytics, and ML, together with simulations on large-scale supercomputers.

- f. Need to run application software stacks and workflows on multiple sites each with a different software environment.
 - g. Existence of multiple user accounts and at various facilities, making running workflows across facilities challenging without integrated security policies or FederatedID.
 - h. Growing need to facilitate persistent or long running services for science campaign.
 - i. Increased need for science campaigns and workflows orchestrated by distributed teams on the system.
4. **Current strategies being employed:**
- a. Ad hoc schedulers written by developers to manage various stages of a workflow with different resource requirements.
 - b. Human-in-the-loop model leveraged at multiple stages of a workflow to assure success.
 - c. Short-running jobs bundled in large chunks by developers, who manage the workflow software on their own.
 - d. Use of MPI to pack the smaller jobs and run them as a single MPI job.
 - e. Resources left unused to ensure availability or to overcome any faults
 - f. Containers increasingly used at HPC facilities because of the need to run complex software stacks and make them portable across sites.
5. **Key research directions of interest:**
- a. *Identifying and managing workflow patterns and resource scheduling for emerging workloads, resource types, and architectures.*
 - i. **What challenge(s) does this research direction address?** A prime challenge is the difficulty in efficiently running complex and emerging workflows at HPC facilities (streaming, AI, interactive usage, coupling experiments and facilities, edge devices).
 - ii. **What are the key research questions?** What are the viable scheduling models that can embrace these changes (for diverse resources, component types and across administrative and security domains)? What are representative workflow benchmarks and miniapps that characterize these new workloads?
 - iii. **What is new and novel about this approach?** Accounting for new science drivers and workloads (streaming, AI, coupling of simulation and analysis. HPC and experiments). New architectures—heterogeneous components, storage tiers, fine-granularity provisioning of resources, future architectures (quantum, neuromorphic). Dynamic requirements and interactions, resources not known a priori, routine frequent failure at scale.
 - iv. **What is the impact of this work on the identified challenges?** Better utilization of facilities. Opportunity to support new workloads and science areas.
 - b. *Determining primitives and abstraction levels for scheduling, managing, and executing workflows*

- i. **What challenge(s) does this research direction address?** Existing workflow and scheduling solutions have limited ability to manage complex resources, often provide inflexible interfaces, and lack appropriate building blocks for composing workflows.
 - ii. **What are the key research questions?** What are the appropriate building blocks to compose scientific workflows, and what software abstractions are needed to effectively compose them? How do we offer resilience support for workflows? What are the appropriate abstraction levels to optimize performance of workflows and complexity of systems?
 - iii. **What is new and novel about this approach?** Determining the fundamental components of scientific workflows for new science drivers within and across facilities, including emerging architectures and workflows originating from experimental facilities, and deciding how to effectively compose them.
 - iv. **What is the impact of this work on the identified challenges?** Identification of correct primitives and abstractions, making the creation and execution of workflows that support the broad range of DOE mission science both possible and effective. If we do nothing, however, we will continue to inefficiently provide ad hoc support for some workflows at some facilities, and other workflows will receive less support from facilities. Humans-in-the-loop will be required for customized support.
- c. *Enabling end-to-end workflow monitoring, modeling, optimization, and automation*
- i. **What challenge(s) does this research direction address?** Workflows require manual intervention and human-in-the-loop operations. Limited ways exist to understand performance and bottlenecks in a complex workflow. Users cannot adapt workflow to current system conditions or dynamic requirements.
 - ii. **What are the key research questions?** How can users expose intent or plans about workloads instead of execution procedures, and how can schedulers/workflows act on that guidance? What are the appropriate telemetry tools and data to optimize workflows effectively? What are the effective workflow-level fault recovery models? How can we use the data to make short-term (real-time, near-time) and long-term decisions at the facilities? How do we model the end-to-end workflows and infrastructure and study their behavior? What is the impact of new facility policies on scheduling behavior, and what should facility policies be in order to meet emerging workloads?
 - iii. **What is new and novel about this approach?** Limited understanding of workflow behavior, preventing effective deployment of complex workflows. Deeper understanding of complex workflows, bottlenecks, and characteristics that will lead to more efficient execution at HPC facilities.
 - iv. **What is the impact of this work on the identified challenges?** Availability of HPC resources (e.g., edge devices) for new science communities to analyze data.

Challenge Name: Smart Facilities and Systems

1. **Challenge:** Determining how to leverage monitored information to characterize workloads, predict system and network failures, and manage new threats for emerging HPC and HPN architectures.
2. **Brief description:** Facilities can collect a wealth of data about their infrastructure ranging from environmental factors to details of jobs running on the machines and data movement on networks. HPC sites gather low-level OS metrics, application monitoring metrics, and information about job queues and job submissions along with the running codes and the libraries used, as well as facility information such as temperature and machine information. HPNs typically collect usage information such as link utilization and flow data, along with monitoring environmental conditions across the network. There is a perceived need to apply machine learning and analytics techniques that can take this information to characterize workloads that can drive *actionable* decisions: (i) characterizing workflows to facilitate better job allocation or intelligent resource (i.e., network, compute, storage) reservation based on usage and demand, identify applications behaving abnormally so that a system can take corrective action; (ii) predicting failures in order to minimize downtime and disruption in ever-evolving facilities with increasingly complex interdependencies between (sub)systems (e.g., compute, storage, network, power, cooling); and (iii) detecting threats and mitigating them to maintain operational integrity in today's facilities, which are growing in scale and complexity that is quickly exceeding the effective scope of human-scale management. In summary, tools are needed that will automate the tedious task of relating and analyzing data from many sources across an HPC site.
3. **Underlying causes and trends:**
 - a. Decisions made by resource managers without knowledge of an application's or workflow's behavior.
 - b. Variability and contention in an HPC system, which can arise for numerous reasons and from different components, including shared resources such as interconnects and filesystems, HPC application behavior that may starve other applications for resources when running, or poorly written code or nonoptimal user behavior.
 - c. Desire of facilities to be able to power cap HPC systems; scheduling for power allocation when desirable (lower energy costs at certain times, hot days).
 - d. Nontrivial task of real-time tuning of resources (network, filesystem) based on congestion. Jobs typically require a combination of network, compute, and storage resources. However, each of these resources is treated as distinct entities and managed independently, resulting in suboptimal resource planning and allocation.
 - e. Collection and curation of monitoring data by facilities in a manner that would not facilitate large-scale analytics. And because each facility is unique, there is not yet a standard or agreed-on set of data to collect. Limited expertise has been available in HPC centers to perform analysis of monitoring data; but even then the data has often been inaccessible to potential analysts and available to system administrators only.

- f. (Unpredictable) timeliness of equipment failure, representing a significant factor in the proportion of the service disruption. Predicting failures in advance can reduce disruption by facilitating proactively targeted scheduled maintenance.
 - g. Cybersecurity threats constant in everyday operations, particularly in networks, thus raising concerns for data confidentiality and integrity of the science experiments. A mitigation strategy is to identify threats by using resource usage and performance behavior (e.g., Identifying DDoS and port-scanning behaviors) but requires the appropriate data fidelity and analytic tools.
 - h. Lack of labeled data sets of what constitutes “normal” versus “abnormal” behavior, hampering further recognition efforts in trends.
4. **Current strategies being employed:**
- a. Use of local monitoring efforts at HPC facilities—but none that encompass the full HPC center in a way that is easily accessed by analysts, and none that effectively join the disparate sources of data to understand the full performance picture at a facility.
 - b. Several research efforts focusing on energy savings by power capping and power allocation (e.g., the GEOPM framework).
 - c. Ad hoc rules in monitoring frameworks to identify performance issues or anomalies and notify administrators, who then look and decide whether something can or should be done.
 - d. Statistical techniques to understand seasonal trends in time stamped data such as data surges across network links and interfaces in SNMP data.
 - e. Data-driven time series prediction techniques such as DL methods, enabling univariate prediction of failures.
 - f. Failure of equipment typically mitigated through redundancy.
 - g. Detection of threats typically based on explicit triggers such as known attack signatures or threshold settings.
 - h. Use of best practices for security—but no two facilities are alike, and exceptions apply.
 - i. Use of monitoring tools such as perfSONAR, Tstat, and Netflow to collect data, each at different time intervals and containing different features of the same transfers.
 - j. Use of only a small number of features to recognize potential threats to a system, due to limited ability to filter, analyze, and store the increasing amount of monitoring data.
5. **Key research directions of interest:**
- a. *Characterizing workflow for optimal resource allocation to maximize system efficiency*
 - i. **What challenge(s) does this research direction address?** Monitoring and optimization — We need to understand at a much finer fidelity how our machines are being used, for example, identify “abnormal” behavior, by correlating the vast, disparate quantities of information gathered on HPC systems to generate workload characterizations.
 - ii. **What are the key research questions?** How do we characterize workloads from a system and network perspective? Can we leverage

- advanced analytics and machine learning techniques to classify workloads by tying together performance data, network data, environmental signatures, executable names, and library tracking to identify a general *fingerprint* for an application or job? Can data be used to make automatic scheduling decisions, find bad hardware, identify abnormal behavior, reboot nodes, adapt resources such as power and networking to where they are needed most, identify failing jobs, and basically to have the machine run itself?
- iii. **What is new and novel about this approach?** Introduction of machine learning to tie together many different data sources. One current approach is manual instrumentation of libraries, modules, and binaries, but this by itself does not fully characterize an application, since different users may use the same code in many different ways (not sufficient for accurate classification).
 - iv. **What is the impact of this work on the identified challenges?** Insight into workload performance anomalies or failures and inform next-generation system design. Better (co-)scheduling of resources based on measured characteristics of particular workflows. Application anomaly detection and potentially automatic response.
- b. *Developing failure prediction to reduce system and network disruption through proactive maintenance*
- i. **What challenge(s) does this research direction address?** Failure prediction – There is a need to minimize disruptions to jobs on HPC systems, both by reducing the frequency of planned maintenances and by minimizing unplanned outages. Such actions would also reduce chances for data loss and lead to better understanding of the causes of failures.
 - ii. **What are the key research questions?** What approaches are most effective using data from the facilities? How can machine learning and advanced analytics be used on sparse heterogeneous data (within and across facilities) to drive facility operating decisions? What failure prediction algorithms will have to be developed or acquired? How can the effectiveness of predictions be demonstrated?
 - iii. **What is new and novel about this approach?** Use of machine learning for job failure detection, classification, and anomaly finding. Some examples exist, but production tools still are needed that can be deployed on supercomputers for optimizing how the resources are used.
 - iv. **What is the impact of this work on the identified challenges?** Reduced frequency of unplanned and planned outages, minimizing disruption to users and reducing workload for administrators.
- c. *Detecting new threats for emerging HPC and HPN architectures*
- i. **What challenge(s) does this research direction address:** Threat Identification – In the next 10 years, HPC systems and networks will include heterogeneous components with new devices and capabilities. We need to understand what unique threats these introduce for emerging compute architectures and networks and how can they be defeated?
 - ii. **What are the key research questions?** How do we determine the HPC and HPN threat model? What are the attack vectors that target HPC systems? How are emerging HPC and HPN architectures and system designs (e.g., opening up interconnects, deploying FPGAs, implementing AI-controlled infrastructure) enabling new attack vectors? How reliable

- are current detection methodologies, and can we speed up the identification of security vulnerabilities for facilities?
- iii. **What is new and novel about this approach?** Research motivated by the emergence of new technologies and investigation of heterogeneous architectures for future HPC systems.
 - iv. **What is the impact of this work on the identified challenges?** Minimization of 0-day vulnerabilities, and better security for facilities. Detection of attempted intrusions without degrading network performance for legitimate users.

Challenge Name: OS Research, System Management and Networks

1. **Challenge:** Determining how to provide new system and network management capabilities and techniques to support emerging workloads on increasingly complex and heterogeneous systems with verifiable, predictable, and transparent performance.
2. **Brief description:** DOE computing facilities face three primary challenges with respect to system and network management. The first is to support emerging workloads on HPC systems, such as data generated by experimental instruments or data, requiring a higher level of security. These workloads may be challenging because of the volume of data involved; the rate at which it must be processed; and the security and sensitivity requirements of transmitting, processing, and storing the data. The second primary challenge is to monitor and manage system components and networks effectively despite increased architectural complexity and heterogeneity. The third challenge is closely related to the second: to provide verifiable and predictable performance and transparent operation of all parts of end-to-end workflows running on federated experimental and computational resources. The capability to understand, control, and reproduce the behavior (functionality and performance) of such workflows is usually desirable and sometimes required. Developers and users want systems to adhere to the "principle of least surprise" so they can reason about the behavior of their software based on their knowledge of the software, hardware, and network of the system(s) they are using.
3. **Underlying causes and trends:**
 - a. Improvements in data analysis techniques, coupled with the inclusion of architectural features specifically designed to support such analyses, making HPC systems increasingly attractive targets for processing experimental data and data from non-traditional sources (e.g., health care data).
 - b. Increasing complexity of HPC system architectures that can deliver higher absolute performance—but they make predictable performance more challenging.
 - c. Decreased availability of data describing the behavior (e.g., performance) of the components of federated resources including within HPC systems, coupled with decreased availability of controls over that behavior (e.g., fine-grained power management). Lack of consistency across the data sources and control mechanisms that are available.

- d. Increased desire or requirement for consistent, reproducible behavior that is transparent to system operators.
4. **Current strategies being employed:**
- a. Duplication of facility resources (e.g., deploying multiple, distinct systems with similar capabilities) in order to provide the required security levels with acceptable performance. Far preferable would be to colocate data storage and processing workflows to avoid the financial and effort costs caused by this duplication.
 - b. Ad hoc aggregation of coarse-grained monitoring data describing the behavior of HPC systems and federated resource components, with manual adjustments using the limited available controls.
 - c. Batch job scheduling with backfill, default OS scheduling within nodes (possibly with strategic elimination of unneeded system services). Manual scheduling at the federated resources level. Limited support for guaranteed interactive or near-real-time scheduling.
 - d. Technologies such as virtual machines and containers to provide reproducible, self-contained software stacks – but at the expense of transparency about the computations running in them.
5. **Key research directions of interest:**
- a. *Addressing OS, system, and network management needs for emerging workloads using shared resources*
 - i. **What challenge(s) does this research direction address?** There is an increased need to support emerging workloads on HPC systems.
 - ii. **What are the key research questions?** How can we manage security and sensitivity requirements in shared-resource environments? How can we provide acceptable performance, such as performance guaranteed to be within tight bounds, or an interactive response?
 - iii. **What is new and novel about this approach?** New (mixed) data sensitivity requirements and new throughput requirements (that facilities have faced before) such as the need for interactive or near-real-time response when processing experimental data, coupled with the desire to share HPC, network, and storage resources instead of duplicating them.
 - iv. **What is the impact of this work on the identified challenges?** Added capability to serve science domains that are not currently feasible to support. Limited resource procurement budgets can be used to procure more (or better) resources instead of duplicated resources.
 - b. *Providing OS, system, and network management support for increased architectural complexity and heterogeneity.*
 - i. **What challenge(s) does this research direction address?** The increasing complexity and heterogeneity in HPC system architectures require careful management.
 - ii. **What are the key research questions?** How can we automate, optimize, and manage complex and heterogenous systems? How can we provide quality-of-service guarantees? How can we support introspection

- (e.g., by management and performance tools) and ease porting of services and tools across systems?
- iii. **What is new and novel about this approach?** Holistic approach to resource management supporting new levels of automated resource monitoring and control in increasingly complex and heterogeneous environments.
- iv. **What is the impact of this work on the identified challenges?** Increased likelihood of achieving good system utilization with low (or at least more predictable) turnaround time for user computation. Decreased learning curve for users and administrators when approaching new systems due to familiarity with services and tools on existing systems. Decreased effort for providing services and tools on new systems.
- c. *Enabling verifiable and predictable end-to-end behavior*
 - i. **What challenge(s) does this research direction address?** A major challenge is attaining verifiable, predictable, transparent performance on HPN and HPC systems.
 - ii. **What are the key research questions?** How can we manage system configuration at the component level (hardware and software) to identify, explain, and mitigate regressions? How much can this be automated? How can we build, package, distribute, and run software in a way that provides verifiable behavior, predictable performance, and observability to system operators (e.g., to detect attempts to circumvent system security)? How can we control or guarantee performance of all resources used in end-to-end computation?
 - iii. **What is new and novel about this approach?** Expanded focus of research including end-to-end resources in complex workflows.
 - iv. **What is the impact of this work on the identified challenges?** Improved ability to explain end-to-end behavior and avoid end-to-end variability of complicated workflows involving HPC systems, thus enabling more users to diagnose and address performance regressions when they occur and to use HPC systems for complex workflows originating from DOE experimental facilities that are sensitive to workflow turnaround time.

Challenge Name: Advances in HPC Storage

1. **Challenge:** Providing HPC storage systems that translate improvements in capacity, bandwidth, and latency from new storage technologies into realizable performance from user workflows without adding additional user complexity.
2. **Brief description:** Currently deployed storage systems at HPC facilities rely on storage software technologies that were designed for a world dominated by spinning disk media and predominantly modeling/simulation workloads. The technology landscape and user diversity have changed dramatically, calling for new solutions that deliver performance outside of large, synchronous reads and writes and data management facilities beyond

complex directory structures and README files. These solutions should include the following:

- a. Better mechanisms for users and administrators to provision and manage their data across multiple storage tiers and namespaces through extensible, indexed, user-defined metadata.
- b. Ability to achieve predictable and reliable performance across a diverse set of competing I/O workloads without having to customize the workflow for each sites storage system.
- c. Support for new modalities of I/O such as streaming data analysis.
- d. Greater resilience to congestion and individual component failures.
- e. Better understanding of storage system behavior, ranging from end-to-end I/O performance to utilization patterns of diverse user communities, using modern tools and techniques.

3. Underlying causes and trends:

- a. Growing need to be able to search, filter, and find data within increasingly large datasets created by simulations and experimental/observational science from DOE user facilities.
- b. Proliferation of high-performance solid-state storage and emergence of nonvolatile memories exposing inefficiencies in software formerly hidden by the low performance of disk drives.
- c. Growth in user and scientific data volumes outpacing growth in media density, increasing the need for efficient but resilient storage systems that can be easily expanded or contracted over time.
- d. Growth in scale reaching fundamental limitations of maintaining strong consistency at extreme parallelism.
- e. Diversification and growth in user communities causing increased contention on globally shared storage systems.
- f. Requirement to preserve data from experiments and simulations for long periods of time and portably share data between facilities when accessing different in-storage system technologies.

4. Current strategies being employed:

- a. Incremental improvements and optimizations to existing file system data and metadata paths.
- b. Insertion of burst buffers and other non-POSIX buffering layers to improve performance at scale.
- c. Declustering parity, new parity algorithms, new consensus algorithms.
- d. Application of metadata-capable middleware such as HDF5; custom project-specific data management systems.

5. Key research directions of interest

- a. *Providing more automated and intelligent storage systems*
 - i. **What challenge(s) does this research address?** Facilities must provide HPC storage systems that deliver continually better performance to users while minimizing additional user complexity or management overhead.

- ii. **What are the key research questions?** What are the right metrics to gather to inform an automated and intelligent storage system? What are the right methods to analyze these metrics? How do we improve models to predict performance and behavior to inform dynamic adjustments the system's behavior? How do we develop metrics from expressions of user intent and facility requirements?
 - iii. **What is new and novel about this approach?** Incorporation and integration of user intent, metadata, and system telemetry to provide better adaptability and overall utility in HPC combines many prior streams of research.
 - iv. **What is the impact of this work on the identified challenges?**
Automated and intelligent storage systems, providing an overall better experience for all stakeholders (users and operators) with minimal added complexity, enabling facilities to support a much broader user base, and improving system efficiency and utilization in response to changes in workload and job mix.
- b. *Ensuring robust, searchable, automatable metadata*
- i. **What challenge(s) does this research direction address?** Users must be able to search and extract more insight from their data and facilitate the reproducibility of data, publishing of data, and sharing data and results.
 - ii. **What are the key research questions?** How do facilities offer a portable, scalable metadata "service" (either embedded within the file system or as separate infrastructure) that addresses these needs while being able to keep up with the performance requirements of storage systems? How do we retain this metadata as data flows within workflows, storage tiers, and across centers?
 - iii. **What is new and novel about this approach?** Application of existing techniques (e.g., those employed in the cloud) to the unique requirements of scientific data while maintaining scalability and portability across diverse computational environments remains unsolved.
 - iv. **What is the impact of this work on the identified challenges?**
Robust, searchable, and automatable metadata, enabling new scientific analyses by fusing multiple datasets, enabling data reuse by users and communities that were not involved in the data's genesis, and building greater trust in scientific outcomes through better provenance.
- c. *Addressing the storage and I/O needs of emerging workloads*
- i. **What challenge(s) does this research direction address?** A key challenge is staying ahead of the I/O requirements of the coming onslaught of data from streaming data sources and artificial intelligence workloads.
 - ii. **What are the key research questions?** Are these new workloads truly different from what HPC has historically encountered, and have their requirements approached a steady state in terms of code quality and data

access requirements? How will HPC storage systems have to adapt to meet these needs? How can these workloads be best adapted to work well on HPC storage systems?

- iii. **What is new and novel about this approach?** Determination of a way to bring the needs of traditional HPC and new workloads together and shape future HPC storage systems to effectively support both.
- iv. **What is the impact of this work on the identified challenges?** Design of storage systems suited to diverse and new workloads, resulting in better cost efficiency through multipurpose and broadly useful systems and enabling new workflows that mix and match tools and methods to better address scientific challenges.

Challenge Name: Programming Environments and Portability

1. **Challenge:** Developing and widely deploying programming models for complex node architectures that are expressive and performant.
2. **Brief description:** Without a set of common programming models for current and future HPC node types, developers are left with an ad hoc assemblage of possibilities. Some of the most popular current choices are proprietary (e.g., CUDA, TBB), rendering them nonportable. Other approaches (e.g., OpenMP 5.x and beyond) rely on uptake by compiler writers of new standards. The ASCR facilities see a clear need to prioritize performance and productivity above other concerns, with performance being especially important. These twin aims have been best served to this point via incremental approaches (e.g., through libraries), and the need to preserve the ability to make slower, secular changes to application software remains a vital concern.
3. **Underlying causes and trends:**
 - a. Node architecture design largely a product of individual vendor innovation and business strategy.
 - b. Proprietary approaches allowing maximum performance for particular hardware but leading to heavily branched development.
 - c. Standards bodies for popular languages moving reactively for the most part.
 - d. Abstractions about near and far memories, vector and scalar processors, and other node components eschewed in favor of architecture-specific notions.
4. **Current strategies being employed:**
 - a. Bifurcation of code bases by developers, producing multiple, independent branches that are tied to specific architectures.
 - b. Incremental improvements by developers to specific computational kernels as access is gained to new architectures.
 - c. Frameworks being adopted (e.g., Kokkos, RAJA) – but these are typically available only for specific languages (e. C++).
 - d. Libraries relied on to do compute-intensive work—but even this use requires code re-architecting in many cases.

5. **Key research directions of interest:**

a. *Investigating and co-designing sets of programming model abstractions for systems*

- i. **What challenge(s) does this research direction address?** Developers must be able to cope with the complexity and rate of change in current and future systems. Moreover, they must be able to tell the computer what to compute.
- ii. **What are the key research questions?** How do we identify the abstractions and the correct level of abstraction for useful work? What data/interfaces are needed to debug – to ensure both correctness and performance? Can interoperability and incremental progress be maintained for impact and effectiveness?
- iii. **What is new and novel about this approach?** The sheer volume of considerations raised by modern heterogeneous system architectures, to the point that new approaches are required. The disappearance of fixed points of resource cost due to the rapid rate of change in this complexity, requiring new ideas on how to reason about these costs.
- iv. **What is the impact of this work on the identified challenges?** Ability of facility users to succinctly describe the problem they want to solve, the parallelism, and so on, and to effectively exploit a variety of architectures. Failure to act could lead to decreased scientific output and reduced competitiveness.

b. *Mapping the programming model abstractions to a set of heterogeneous hardware resources*

- i. **What challenge(s) does this research direction address?** Two challenges arise: (1) how can dynamically changing system resources be effectively marshaled by developers of both “traditional” HPC workloads and emerging data and learning applications? and (2) how does the system know what to do?
- ii. **What are the key research questions?** How do we perform the mapping to the hardware, and who or what has that responsibility? How can we capture and reason about runtime/programming model interaction? How can abstractions be made to directly support composability? How can one reason about energy consumption, for example, the interaction with runtime performance in an energy-constrained and dynamic environment?
- iii. **What is new and novel about this approach?** Improvements in managing and tracing program execution, for example by using methods such as JIT and dynamic model generation that address the current problems of separate, independent runtimes (e.g., OpenMP, MPI) and multiple levels of abstraction (e.g., internally threaded libraries).

- iv. **What is the impact of this work on the identified challenges?** More effective monitoring of complex systems and more effective use of dynamic resources to overcome the complexity barrier. Failure to successfully address the challenges would significantly decrease the efficiency of future applications and decrease science productivity.
- c. *Investigating methodologies, capabilities, and models for debugging, correctness, and performance*
 - i. **What challenge(s) does this research direction address?** Research is needed to determine how we can provide introspection across various components and how we (or the runtime, or the OS) can know what the system is doing and act on that information.
 - ii. **What are the key research questions?** How do we formulate performance and power modeling and cost models while confronting reproducibility, predictability, and variability, even at the level of an entire workflow? Will new debugging models be required? At what level and how will integrated analyses be required?
 - iii. **What is new and novel about this approach?** The concept of introducing multilevel hierarchical analysis (both static and dynamic) to produce a single picture. The coordination of multiple decision-makers in the system and methods to make sense of the results are frontiers.
 - iv. **What is the impact of this work on the identified challenges?** Effective tools for users and necessary information for facilities. Failure would mean that performance and correctness would be almost impossible to ascertain and reason about.

Cross-Cutting Themes

Five common themes arose in the breakouts:

- Addressing increasing heterogeneity in HPC systems
- Understanding and supporting emerging workloads such as streaming, data analytics and machine learning
- Enabling automation and eliminating human-in-the-loop requirements
- Developing abstractions for expressing user intent
- Improving modeling, with deeper understanding of application, workflow, and workload characteristics

The connections between the cross-cutting themes and the six breakouts are given in the table below. Each of the themes is then described more fully in the text.

Breakout/Crosscut Theme	Heterogeneity in HPC Systems	Understanding and Supporting Emerging Workloads	Automation	Developing Abstractions	Improved Modeling
Future of Computer Architectures	X	X			X
Workflows and Scheduling	X	X	X	X	X
Smart Facilities and Systems	X	X	X		X
OS Research, System Management and Networks	X	X	X		
Advances in HPC Storage		X	X	X	
Programming Environments and Portability	X	X		X	X

Addressing Increasing Heterogeneity in HPC Systems

A cross-cutting theme described by the breakouts was the need for additional research into how to adapt to, utilize, and manage changes due to the deployment of more heterogeneous architectures. In the *Future of Computer Architectures* breakout this was highlighted as the need for better understanding of how complex workflows and applications could use heterogeneous components, the need for new simulation and modeling tools to determine the optimal architecture of a system for a given workload, and the need for additional research to explore the design space for heterogeneous components. In the *Workflows and Scheduling* and the *Programming Environment and Portability* breakouts the common theme was finding the right abstraction layers in order to be able to best take advantage of new heterogeneous architectures. In the *Smart Facilities and Systems* breakout, a focus was on how to identify new threat vectors for more heterogeneous systems. In the *OS Research, System Management and Networks* breakout a theme was the need for research into how to optimize and manage increasingly complex and heterogeneous environments. In all the breakouts where heterogeneity of computing architectures was addressed, the common element was that without core research activities, end-user applications and workflows will not be able to take advantage of key hardware advancements.

Understanding and Supporting Emerging Workloads

Emerging workloads such as machine learning, deep learning, analytics, and coupling of experiments and observations with HPC facilities are expected to be increasingly deployed on HPC facilities for scientific discoveries. All six working groups identified and discussed the need for effective understanding of these workloads and support for performant execution of these on HPC systems.

In particular, the *Workflows and Scheduling* breakout highlighted the need for end-to-end performance analysis tools in order to understand the entire application, the *Programming Environments and Portability* breakout highlighted the need for hierarchical analysis to produce a global picture, and the *Future of Computer Architecture* breakout highlighted the need for appropriate simulators to account for these new applications and novel systems and architectures.

To effectively map these workloads on the current platforms and to design future architectures to meet emerging needs for performant execution, the groups agreed that new R&D was essential across the software stack—spanning programming models, workflows, operating systems, resource management and runtimes, and data management—together with system architecture R&D. Toward this end, the *Workflows and Scheduling*, *OS Research*, *System Management and Networks*, and *Programming Environments and Portability* groups highlighted the need for R&D into effective runtimes and execution of these workloads while facilitating reproducibility and predictable performance.

Enabling Automation and Eliminating Humans-in-the-Loop

The desire to tune, adjust, or redirect resources while a complex process is running is a common theme throughout HPC. Users do these tasks while executing scientific workflows to ensure that the results from each step along the path from experimental setup to final data analysis is proceeding within normal parameters. Facilities operators do these tasks to fine-tune performance, availability, and utilization as workloads change. Unfortunately, such tuning and adjustment often require a human-in-the-loop, to evaluate whether the workflow or HPC subsystem is operating as expected and, if not, to make decisions about how best to steer the overall process toward a better outcome.

As a result, a cross-cutting desire emerged in the groups to develop techniques to automate control loops. Supporting complex workflows and optimizing hierarchical storage systems require systematic methods to (1) take information from the user's intent and metadata and the system's available system resources, (2) decide what overall configuration would be best suited for executing on the user's intent given the constraints of the diversity of system resources, and (3) measure the effects of applying that configuration to inform (1). A tremendous opportunity exists to accelerate end-to-end scientific workflows and vastly increase the efficiency of storage systems management by developing the tools and techniques to automate this process and remove the human in the loop. Doing so will likely require improving the tools and interfaces for collecting data about workflows and storage systems, combined with developing techniques to apply artificial intelligence to translate such telemetric data into dynamic configuration changes.

In the *Smart Facilities and Systems* and *OS Research, System Management and Networks* breakouts a key theme concerning automation was the need to automate the analysis of data from different system and facility sources in order to create actionable decisions such as identifying a threat, a poor-performing application, or a system component that may soon fail. In the *OS Research, System Management and Networks* breakout the emphasis was on automating the complexity of managing heterogeneous systems, including system configuration, resource management, and scheduling. All of the breakout groups who identified automation and reducing humans-in-the-loop as key themes noted the need for additional research in order to make the operation of next-generation heterogeneous systems supporting new workflows more efficient and requiring less manual intervention.

Developing Abstractions for Expressing User Intent

A recurring theme in several of the breakout sessions was the need to develop useful abstractions for ecosystem components that would allow users to express their intent rather explicitly marshaling resources to carry out an aim. Perhaps the central commonality expressed in these discussions was the need to manage the increasing complexity of programming models. Details of the underlying architecture need to be exposed to programmers to some extent, but the need to rely on explicit management of these features is too great a burden if portability to other platforms is important (as it almost always is for ASCR facility users). For scientific data workflows, many of the same concerns with respect to architectural complexity are present, but the possibly dynamic nature of workflow needs adds additional complexity to these needs. Some uncertainty also exists about whether practitioners may or may not be familiar with “traditional” HPC ecosystems. Future storage systems extend similar concerns to both ends of the canonical workflow: advancing storage performance without exposing all the requisite complexity to achieve that performance is seen as essential. Furthermore, users must be able to manage the complexity of the data themselves and to communicate to a wider audience the methods and techniques they use in this reasoning from the data.

Improving Modeling and Enabling Deeper Understanding of Application, Workflow, and Workload Characteristics

A desire for better and deeper understanding of the performance characteristics of the workload at DOE facilities was present in several breakout sessions, including the *Workflows and Scheduling, Future of Computer Architectures, and Programming Environments and Portability* breakouts. These questions were often raised at the same time as the theme of developing better modeling techniques to understand and quantify the impact of upcoming technology trends. Taken in combination, these concerns reflect a desire to be able to understand and, potentially, model the amenability of particular components of a workload to different heterogeneous components. Increased efforts in this direction could benefit numerous areas, and the groups expressed the hope that any future research in this space will encompass the cross-cutting nature of this topic.

Conclusions

The DOE Computational Facilities – Research workshop brought together ASCR researchers and DOE computational and networking facilities staff to discuss the challenges that facilities are facing and how ASCR research could address these challenges with research in a 10-year timeframe. The workshop was divided into six breakout groups: *Future of Computer Architectures*; *Workflows and Scheduling*; *Smart Facilities and Systems*; *OS Research, System Management and Networks*; *Advances in HPC Storage*; and *Programming Environments and Portability*. Each group briefly described current trends and strategies being employed in their area and then discussed three research directions of interest to that area, including the challenges each addressed, the novelty of the research approach, and the potential impact. Also emerging from the discussions was identification of common cross-cutting themes across the breakout areas: system heterogeneity, emerging workloads, automation, abstractions, and improved modeling.

Appendix

DOE Computational Facilities – Research Workshop

Workshop Organizers

David Brown (LBNL)

Barney Maccabe (ORNL)

Valerie Taylor (ANL)

Workshop Organizing Committee

Katie Antypas (LBNL) [chair]

Katherine Riley (ANL) [co-chair]

Ken Alvin (SNL)

James Ang (PNNL)

Chin Guok (LBNL)

Mike Lang (LANL)

Bronson Messer (ORNL)

John Shalf (LBNL)

Becky Springmeyer (LLNL)

Breakout Leads

Future of Computer Architectures

- Matt Leininger (LLNL) leininger4@llnl.gov
- Nick Wright (LBNL) njwright@lbl.gov

Workflows and Scheduling

- Katie Antypas (LBNL) kantypas@lbl.gov
- Venkat Vishwanath (ANL) venkat@anl.gov

Smart Facilities and Systems

- Mark Fahey (ANL) mfahey@anl.gov
- Chin Guok (LBNL) chin@es.net

OS Research, System Management and Networks

- Phil Roth rothpc@ornl.gov
- Trent D’Hooge tdhooge@llnl.gov

Advances in HPC Storage

- Bill Allcock (ANL) allcock@anl.gov
- Glenn K. Lockwood (LBNL) glock@lbl.gov
- Sarp Oral (ORNL) oralhs@ornl.gov

Programming Environments and Portability

- Bronson Messer (ORNL) bronson@ornl.gov
- Katherine Riley (ANL) riley@anl.gov

Agenda

Day 1

7:30 Bus to pick up at Guest house, transport to building 241

8:00 Registration

8:45 - 9:00 Welcome and Introduction, Valerie/David/Barney + Charge + Q&A

9:00 - 10:00 Joint Presentation from ASCR facilities on our biggest challenges

10:00-10:30 - Break: Move to breakout sessions

10:30-12:00 - Breakout Sessions 1 (1.5 hours)

A) Advances in HPC Storage

B) Smart facilities and systems

C) Programming Environments and Portability

12:00-1:00 Working Lunch

1:00 - 3:00 Breakout Sessions 1 Continues (2 hours)

3:00 - 3:30 - Break

3:30 - 4:30pm - Breakouts Sessions 1 Continues (1 hour)

4:30pm - Breakout reports

Day 2

7:30 Bus to pick up at Guest house, transport to building 241

8:00 Registration

8:30-8:45 Charge for the day and Q&A

8:45-9:00 Move to breakouts

9:00-10:30 Breakout Session 2 (1.5 hours)

A) Workflows and Scheduling

B) Future of Computer Architectures

C) OS Research, System Management and Networks

10:30-11:00 - Break

11:00-12:00 - Breakout Session 2 Continues (1hr)

12:00-1:00 - Working Lunch

1:00 - 3:00 - Breakout Session 2 Continues (2hrs)

3:00 - 3:30 - Break

3:30 - 4:30 - Breakout reports

Attendees

Name	Organization
William Allcock	ANL
Ben Allen	ANL
Pavan Balaji	ANL
Mark Fahey	ANL
Carissa Holohan	ANL
Paul Hovland	ANL
Rob Ross	ANL
Venkat Vishwanath	ANL
Stefan Wild	ANL
Eli Dart	ESnet
Mariam Kiran	Esnet/CRD
Yatish Kumar	ESnet
Eric Pouyoul	ESnet
Latchesar Ionkov	LANL
Pat McCormick	LANL
Brad Settlemyer	LANL
Ann Almgren	LBNL
Esmond Ng	LBNL
Lavanya Ramakirshnan	LBNL
Sam Williams	LBNL
Dong Ahn	LLNL
Abhinav Bhatele	LLNL
Trent D'Hooge	LLNL
Matt Leininger	LLNL
Kathryn Mohror	LLNL
Barry Rountree	LLNL
Brandon Cook	LBNL-NERSC
Lisa Gerhardt	LBNL-NERSC
Tavia Stone Gibbins	LBNL-NERSC
Glenn Lockwood	LBNL-NERSC
Eric Roman	LBNL-NERSC
Chris Samuel	LBNL-NERSC
Nick Wright	LBNL-NERSC
David Bernholdt	ORNL
Markus Eisenbach	ORNL
Austin Harris	ORNL
Judy Hill	ORNL
Dmitry Liakh	ORNL
Sarp Oral	ORNL

Norbert Podhorszki	ORNL
Tom Potok	ORNL
Phil Roth	ORNL
Jeff Vetter	ORNL
Kevin Barker	ORNL
Sriram Krishnamoorthy	PNNL
Nathan Tallent	PNNL
Chris Zimmer	PNNL
Ron Brightwell	SNL
Robert Clay	SNL
Rob Hoekstra	SNL
Jay Lofstead	SNL

Division Directors	
David Brown	LBNL
Sudip Dosanjh	LBNL-NERSC
James Hack	ORNL
Barney Maccabe	ORNL
Inder Monga	ESnet
Mike Papka	ALCF
Valerie Taylor	ANL
Organizing Committee	
Katie Antypas [chair]	LBNL-NERSC
Katherine Riley [co-chair]	ALCF
Ken Alvin	SNL
James Ang	PNNL
Chin Guok	ESnet
Mike Lang	LANL
Bronson Messer	ORNL
John Shalf	LBNL



Argonne National Laboratory
9700 South Cass Avenue
Argonne, IL 60439

www.anl.gov



Argonne National Laboratory is a U.S. Department of Energy
laboratory managed by UChicago Argonne, LLC