STORAGE SYSTEMS AND INPUT/OUTPUT: **ORGANIZING, STORING, AND ACCESSING DATA FOR SCIENTIFIC DISCOVERY**

REPORT FOR THE DOE ASCR WORKSHOP ON STORAGE SYSTEMS AND INPUT/OUT

Gaithersburg, Maryland
September 19–20, 2018

U.S. DEPARTMENT OF ENERGY | Office of Science

Sponsored by the Office of Advanced Scientific Computing Research

The Impact of **Storage and I/O**

The success of the U.S. Department of Energy (DOE) computational, experimental, and observational sciences is inextricably tied to the usability, performance, and reliability of emerging storage systems and input/output (SSIO) technologies. In September 2018, therefore, DOE convened a workshop in which experts from the national laboratories, industry, and academia were asked to consider SSIO technologies from the viewpoint of both the scientific drivers, such as metadata and learning applications, and the technological drivers, including deepening storage hierarchies and architectural changes in future high-performance computing platforms.

SSIO technologies involve the organization, movement, and placement of data to enhance computation and discovery. Key questions arise, for example, from applications that are generating extreme-scale data and from use cases requiring simultaneous analysis of data from a variety of sources—often in real time. Effective exploitation of diverse but finite resources will require decisions about data retention times and access needs. Ultimately, the goal is to allow researchers to maintain their conceptual representations of scientific data, while simultaneously optimizing for performance and productivity across a variety of platforms.

The comments and thoughts of the workshop participants about these challenges and questions were distilled into a set of four research opportunities with the potential for high impact on DOE science. These research opportunities are summarized in the following pages, with more details available in the full workshop report (DOI: 10.2172/1491994).

DRIVERS	KEY QUESTIONS	RESEARCH OPPORTUNITES
Scientific □ Increasing need to support big data and learning applications □ Rapid growth of scientific dataset sizes **Technological** □ New, solid-state storage and tight integration in platforms □ New accelerators, sensors, and networks	□ How do we maintain scientists' productivity while leveraging complex and multi-layer storage systems? □ How can AI/ML assist in managing complex storage environments? □ What new software will be needed to adapt to streaming data sources? □ How do we enable storage to cooperate with workflow and scheduling systems? □ How do we motivate and maintain user trust?	□ Enabling science understandability and reproducibility through rich data formats, metadata, and provenance □ Accelerating scientific discovery through support of in situ and streaming data analysis □ Enhancing SSIO usability, performance, and resilience through monitoring, prediction, and automation □ Improving efficiency and integrity of data movement and storage through architecture of systems and services

Workshop participants identified a set of scientific and technological drivers that motivate research and development in this area, and they grappled with questions related to productivity, efficiency, and trust. Ultimately four high-value research opportunities emerged from the discussion.

Research Opportunities

Enabling Science Understandability and Reproducibility

Scientific research increasingly involves analyses over extremely large and complex data. Petabytes of data may be distributed across billions of files. Traditional approaches for organizing data are not meeting researcher needs.

Rich data formats may enable science communities to publish, search, and share datasets more—essential capabilities for reproducibility.

Storage systems with scalable, standard, extensible **metadata support** are needed in order to capture and store information about scientific data and its relationships. Decoupling science metadata storage from the underlying file system will allow each to evolve independently.

Improved and scalable **provenance tracking** is needed to document the science data lifecycle. Such tracking is essential for reproducibility and validation and can be leveraged to improve software fault detection and performance measurement.

Accelerating Scientific Discovery

Experimental and observational data is being produced at unprecedented rates and volumes. SSIO technologies will need to absorb vast amounts of data quickly, enabling critical data to be delivered and processed while the simulation is running and on the same platform where a simulation is executing.

To facilitate these new processing paradigms, researchers must answer several questions.

For **in situ analysis**, what data should remain in the system, and where should it reside? For **streaming analysis**, how does one handle the increasing heterogeneity between the data source and the destination computational facility?

Workflows explicitly describe dependencies between data and tasks. Understanding is needed about how these workflows can best inform reliability, performance, data movement, and storage requirements.

Enhancing SSIO Usability, Performance, and Resilience

Scientific data storage systems must orchestrate thousands of hardware components and dozens of system software components to serve the data management needs of a wide variety of applications.

Progress can be made along three research thrusts: (1) new scalable **methods for monitoring data** from HPC systems and the applications that execute on them; (2) **predictive models** trained to interpret this wealth of data to guide application optimizations and facility procurement decisions; and (3) **automated methods** to consistently translate monitoring and prediction into practical, actionable solutions without exacerbating the demands on facility staff and resources.

Improving Efficiency and Integrity of Data Movement and Storage

SSIO systems are becoming increasingly diverse in both type and architecture. Research is needed to increase interoperability across these systems, so that both node-local storage resources and shared storage devices can be coordinated and leveraged.

New architectures of systems and services are needed that better integrate SSIO systems with schedulers, resource managers, and monitoring and workflow systems, for example, mechanisms that inform the system about the storage hierarchy and sharing and lifetime requirements. Middleware leveraging these capabilities will maximize performance, minimize data movement, and ensure integrity of scientific data sets.

A New and Productive Storage and I/O Environment

Providing an environment in which scientists can productively operate on vast troves of data is fundamental to the success of numerous scientific endeavors using computing, whether the approach is through simulation, big data analysis, or the application of machine learning. Storage and I/O technologies provide the foundations for successful data storage, management, and access.

New research in storage and I/O is motivated by two distinct trends. First, new approaches to scientific discovery in traditional and emerging science domains drive the rapid growth of dataset size, increasing the importance of supporting data in motion. These new approaches present new requirements for storage and I/O technologies, including the need to traverse very large datasets for learning applications and management of streams of data in real time. Second, new computing, storage, and sensor technologies require a rethinking of our I/O models and software in order to advance scientific discovery. For example, new storage technologies are providing a wealth of options with different performance and resilience capabilities. Careful and orchestrated use of collections of varied devices will be critical to meeting performance and data retention goals in a cost-effective way.

The workshop identified four key research opportunities with the potential to advance DOE science discoveries: making storage and I/O capabilities reproducible, accelerated, resilient, and efficient, through new developments in extensible metadata, in situ data analysis and streaming, automation, and coordination of both local and shared storage devices.

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 any of their employees, 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.