Summary of the Advanced Scientific Computing Advisory Committee Facilities Subcommittee Report

Recommendations for the Future of ASCR Facilities

May 2024

Presented by

Ed Seidel, Amanda Randles

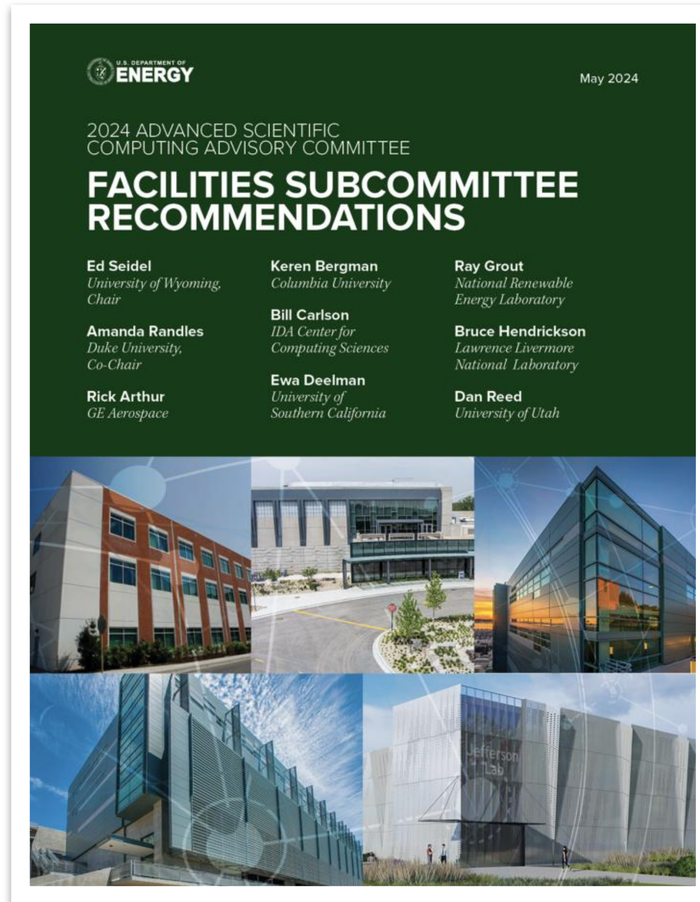

Committee Members

- Ed Seidel, University of Wyoming, Chair
- Amanda Randles, Duke University, Co-Chair

- Rick Arthur, GE Aerospace
- Keren Bergman, Columbia University
- Bill Carlson, IDA Center for Computing Sciences
- Ewa Deelman, University of Southern California
- Ray Grout, National Renewable Energy Laboratory
- Bruce Hendrickson, Lawrence Livermore National Lab
- Dan Reed, University of Utah

The Charge
SC Director Berhe Asked That We…

- Consider what new or upgraded facilities will be necessary in the coming decade (2024-34) to position the SC at the forefront of scientific discovery, including a list of five specific ASCR facilities: ALCF, OLCF, NERSC, HPDF, ESnet
- The potential of each to contribute to world-leading science in the next decade, considered in terms of the
  - readiness for construction
  - sufficiency of R&D performed to date to ensure technical feasibility
  - extent to which the cost to build and operate the facility is understood; and site infrastructure readiness.
- Please place each facility in one of three categories
  - ready to initiate construction
  - significant scientific/engineering challenges to resolve before construction
  - mission and technical requirements not yet fully defined

Top Level Summary of Report

(More detail in a minute)

1. Support and develop all five ASCR facilities
   - ALCF, OLCF, NERSC, HPDF, ESnet absolutely essential, but in different stages of development
   - Significant R&D required, depending on facility
2. Recognize, manage ASCR facilities into as an integrated ***Ecosystem***
   - Integration required to serve science
3. Launch a comprehensive R&D program
   - Cannot be business as usual
   - Requires prototyping on 5-year timescale that informs pathways to future systems on ten-year timescale
   - DOE cannot go it alone; other agencies, vendors, science communities all needed
   - New governance model needed, both within DOE and across agencies

Failure to follow these suggestions risks loss of global leadership in DOE and in broader US science, technology and economic development capacity

Based on Scientific Need

KEY FINDINGS

Key Findings of Committee (1)

ASCR advanced computing systems continue to be critical for SC to remain at the forefront of scientific discovery as science becomes more interdisciplinary, integrated, and digital.

Science-driven imperative: DOE science programs require large experimental facilities, theory, and leading-edge computation, data analysis and storage, and advanced networking.

This includes new applications and integrated workflows for data management, artificial intelligence (AI)/machine learning (ML), and physics-based simulation and modeling.

The success of programs managed by ASCR, BES, BER, HEP, FES, and NP depends on the success of ASCR facilities, a point *reinforced by interviews with other SC subcommittees.*

Key Findings of Committee (2)

A combination of complementary facilities is needed to support DOE mission science, and through their integration should be thought of as a single overarching ecosystem with multiple components.

These components provide complementary functions and are integrated and synergistic. Together, they span the set of capabilities needed to support DOE mission science in the coming decade.

We do not see how one can be funded over another without risking science goals across all SC programs.

Key Findings of Committee (3)

ASCR facilities are also important to other organizations, including but not limited to, *NNSA, NSF, NIH, NIST, NASA, NOAA, and DoD, and to U.S. industrial competitiveness*.

The impact of ASCR facilities extends beyond SC to support scientific discovery, U.S. industry, and global economic competitiveness.

By virtue of its investment and expertise, DOE is, *de facto*, the lead federal agency for advanced scientific computing.

If DOE fails to lead R&D in the face of a changing computing world, the rest of the country will suffer as well.

Key Findings of Committee (4)

Continued success of DOE mission science requires an "all hands on deck" approach to developing next-generation computing infrastructure.

Because ASCR facilities operate in rapidly evolving economic and technical landscapes of the semiconductor and computing industries and changing research practices, *progress cannot be business as usual*.

This committee views it as critical for ASCR to leverage expertise beyond the boundaries of the Office of Science to conduct the necessary R&D and attract vendors in R&D partnerships.

Key Finding (1)

ASCR advanced computing systems continue to be critical for SC to remain at the forefront of scientific discovery as science becomes more interdisciplinary, integrated, and digital.

Recommendation 1

ALL FIVE ASCR FACILITIES ARE ESSENTIAL

Recommendation 1

Ensure the continued support and development of all five ASCR computational facilities reviewed—ALCF, OLCF, NERSC, HPDF, and ESnet—as they are central and essential to all SC science programs and broader national science and engineering research programs.

Each facility provides distinct and critical functionalities that are essential to achieve SC science goals.

Facility	Description	Importance to SC Science Mission	Readiness for Construction for 2034 deployment
ALCF	Leadership Class computing facility	Absolutely central; required for success of science mission	ALCF-4: Ready to initiate construction ALCF-5: Significant scientific and engineering challenges to resolve before construction
OLCF	Leadership Class computing facility	Absolutely central; required for success of science mission	OLCF-6: Ready to initiate construction OLCF-7: Significant scientific and engineering challenges to resolve before construction
NERSC	High performance production scientific computing center	Absolutely central; required for success of science mission	NERSC-10: Ready to initiate construction NERSC-11: Significant scientific and engineering challenges to resolve before construction
ESnet	High performance networking; connects computing & experimental facilities across SC	Absolutely central; required for success of science mission	ESnet-7: Ready to initiate construction ESnet-8: Significant scientific and engineering challenges to resolve before construction
HPDF	Distributed data-focused facility with hub-and-spoke architecture. Will provide unique data storage and management capabilities.	Absolutely central; required for success of science mission	HPDF Hub: Significant scientific and engineering challenges to resolve before construction HPDF Spokes 1 and 2: Mission and technical requirements not yet fully defined

Recommendation 1

Ensure the continued support and development of all five ASCR computational facilities reviewed—ALCF, OLCF, NERSC, HPDF, and ESnet—as they are central and essential to all SC science programs and broader national science and engineering research programs.

Each facility provides distinct and critical functionality that are essential to achieve SC science goals.

Necessary but not sufficient for success!

Key Finding (2)

A combination of complementary facilities is needed to support DOE mission science, and through their integration should be thought of as a single overarching ecosystem with multiple components.

Recommendation 2

THE ECOSYSTEM

Recommendation 2

Science demands integration. We advocate viewing ASCR facilities not as isolated entities, but as *integral components of a single, larger integrated computational Ecosystem,* with a single governance model.

- Will require new ways of governing and potentially funding the overall ***Ecosystem*** (e.g. this group of facilities and ongoing integration)
  - *should not be developed via individual site procurements*.
- Critical for supporting SC science programs, along with additional software, algorithm, workforce, and science application components, to serve science and engineering research.
- Importance to entire national scientific and technological capability, and its importance as a model internationally, cannot be overstated.

The *Ecosystem*

- **Science-driven imperative**: *Science increasingly requires combination of experiment, computation, data analysis*
  - Complementary capabilities of ALCF, OLCF, NERSC, HPDF, and ESnet should be viewed, integrated, and operated as integrated ASCR facilities ***Ecosystem***.
- **Supports National Priorities**: *Ecosystem* provides unique capabilities for *SC & other research agencies that cannot be found elsewhere,* essential to entire national scientific enterprise.
- **Serves U.S. industry, helping sustain global competitiveness for the nation**. Industry needs for ***Ecosystem*** will continue to grow over the coming decade.
- Is under ASCR, but we urge SC leadership to…
  - Work w/ASCR and other SC offices to manage and fund it
  - Work w/NNSA and key science agencies to collaborate on use and R&D to develop it
- We suggest a new governance model
  - **Form a higher-level coordination body**
  - Include lab directors and ensure that future generations are distinct and innovative, addressing risks of uniformity
- It must be dynamic to allow for new facilities

Supporting Exponential Data Growth

Relation of the Ecosystem with IRI

- The Ecosystem must complement Integrated Research Infrastructure (IRI)
- IRI will integrate *ASCR computational systems* with *SC's experimental facilities*
  - Serving science need: integrate complementary approaches
- ASCR already collaborating with BER, BES, HEP, NP, FES to develop IRI
  - Will connect all 28 SC user facilities across DOE national laboratories seamlessly
- ***Ecosystem*** is vital yet distinct from IRI
  - Stands on its own under ASCR
  - Also essential to support IRI; can't be realized w/out Ecosystem

Key Findings (3) and (4)

(3) ASCR facilities are also important to other organizations, including but not limited to, NNSA, NSF, NIH, NIST, NASA, NOAA, and DoD, and to U.S. industrial competitiveness.

(4) Continued success of DOE mission science requires an "all hands on deck" approach to developing next-generation computing infrastructure.

Recommendation 3

COMPREHENSIVE R&D PROGRAM

Recommendation 3

- **A comprehensive, coordinated R&D program delivering multiple prototype computing systems over a five-year timescale must be mounted to inform pathways for this integrated ecosystem, operational by 2034, due to**
  - rapidly evolving economic/technical landscapes of semiconductor and computing industries
  - changing research practices
- With the end of Moore's Law and with vendors focusing on other markets, the future of high-end computing for science is highly uncertain.
  - R&D is needed to chart this course, to influence vendors, and to prepare applications for future platforms.
  - The changing nature of interdisciplinary research requires a deeper integration of facilities, workflows, algorithms, software and application tools.
- The R&D program should involve a collaborative effort across DOE computational facilities spanning SC. To ensure comprehensive, multidisciplinary approach, it is critical to include contributions from
  - computing and cloud vendors, computer science, DOE experimental facilities, domain science and engineering communities, with deep collaborations with DOE NNSA and other federal agencies.
- A new governance model will be required to manage this.

Backdrop

- Computing landscape (advanced computing especially) undergoing profound change
  - Outside DOE control…
    - DOE no longer big enough to influence nor go it alone
  - Generative AI and cloud hyperscalers shift market
  - Free lunch of Dennard scaling and Moore’s Law over
- DOE’s differentiated mission requires new model of public–private partnerships
  - Must design, develop, and deploy a next gen advanced computing infrastructure
  - Not naturally coming with above changes in market

Comprehensive R&D Program (1)

- Current model of periodic vendor system procurements will not work
- Must leverage collaborative R&D initiatives to *shape the future*
  - Must span all of DOE, leveraging expertise across SC and NNSA
  - Foster broad, whole-of-government partnerships with other federal agencies that have related missions and depend on infrastructure
    - International as well…
    - Increase capacity to address need, influence market, involve science stakeholders who depend on DOE facilities

Comprehensive R&D Program (2)

- More in-depth and foundational R&D approach than previous "PathForward" initiatives
  - Establish long-term collaborative partnerships with variety of technology partners
    - Traditional hardware vendors, startups, and cloud hyperscalers
  - Hardware and software development
  - Integrate tools and techniques for handling complex, distributed workflows and enabling multidisciplinary discovery
- Build substantial hardware–software prototypes (5-year timescale)
  - Test ideas and help de-risk promising technology paths for component technology providers and product vendors
- **Workforce essential**: R&D approach will help attract and retain the necessary talent to keep DOE and nation globally competitive
  - WF Dev vital for sustaining these capabilities and propelling ***Ecosystem*** forward
  - Developing and retaining top talent within DOE and the broader U.S. scientific community is crucial for maintaining global competitiveness and security

Comprehensive R&D Program (3)

Only once the feasibility and integration of these hardware–software prototypes have been validated should the procurement process commence (for later gen systems)

Summary: Must

- Identify and prototype viable computing technologies before committing to significant future procurements.
- Develop methods to integrate these technologies into a cohesive, multidisciplinary computing, data, and network infrastructure designed to support DOE science programs

Critical also to other agency programs reliant on ASCR facilities

Comprehensive R&D Program (4)

- DOE simply cannot "go it alone."
- **Failure to follow recommendations risks nation's advanced computing ecosystem**, including
  - Loss of U.S. global leadership in advanced computing
  - Further destabilization of the computing hardware vendor ecosystem due to premature technology choices
  - Inability to achieve DOE's science objectives, as well as collateral science effects at other agencies that depend on DOE
  - New generations of systems with even lower efficiency, with concomitant scientific, technical, and political risks
- **Failure to adopt a long-term, integrated R&D program may lead to erosion or loss of program funding**
- Success in this approach should lay the foundations for success across all such critical areas

All rated as ***absolutely essential,*** *although in different stages of development*

INDIVIDUAL FACILITY ASSESSMENTS

Leadership Class Facilities
ALCF, OLCF

- Strong conclusion of subcommittee
  - *Two* LCF facilities serve to anchor technology pathways, build, sustain critical workforce, enhance U.S. economy
    - one facility simply cannot
  - Unique science resources to DOE, agencies, industry
- Next gen systems OLCF-6, ALCF-4 in CD process
  - Aim for operation by late 2020's
    - 5x-10x improvement in applications performance
    - Improved energy efficiency
    - Integration of AI, modeling and simulation, and data-intensive capabilities in a single resource
    - Support for DOE's plans for IRI
  - Require R&D as part of the procurement process

Leadership Class Facilities (2)

- Future generation systems OLCF-7, ALCF-5
- Aim for operations by 2034
  - Required capabilities and the associated scientific opportunities much harder to predict
  - Need for computing resources certain to be very high
- Require R&D as described above
- Must be informed by prototype systems

NERSC

- Envisions transition from current emphasis on modeling and simulation, AI training and inference, and data analytics
  - Becoming a workflow-driven facility
- Stands out in supporting a diverse range of computational workflows
  - Approximately 10,000 scientists engaged in 1,000 projects per year
- Subcommittee chairs often named it as their essential workhorse

NERSC (2)

- NERSC-10 more evolutionary, can proceed to construction with primarily facility-specific R&D activities
  - Centered on creating workflow system components, APIs, and federated IDs
  - Anchor NERSC's integration with the ***Ecosystem***
- NERSC-11, expected around 2030
  - Aims to build on the achievements of NERSC-10
  - Expand connectivity within the DOE and to broader emerging technological ecosystem, encompassing initiatives like IRI/HPDF, pervasive AI, and advanced computing paradigms
  - Currently in preliminary planning phase, NERSC-11 strategically positioned to leverage the technological strides and insights gained from NERSC-10 and current R&D

High Performance Data Facility

- First-of-kind facility will support future DOE SC data-intensive applications
  - Critical to SC mission, cornerstone of IRI vision
  - Enable and accelerate scientific discovery by delivering state-of-the-art data management infrastructure, capabilities, and tools
  - Fills critical need for data storage/management solution to enhance support for DOE science
- At heart of IRI, provides high-performance data management solution
  - Hub and spoke model
    - Hub hosts centralized resources and services
    - Multiple mission-application spokes
  - Physically located at JLab and LBNL connected to extensible network of spokes via ESnet

High Performance Data Facility

- Site selection for HPDF only in October 2023; facilities design needs to be refined to reach CD-1
- While placed in the category as absolutely essential to achieve SC mission goals, we believe substantial R&D work is still required:
  - HPDF Hub is in category (b) significant scientific & engineering challenges to resolve before initiating construction
  - HPDF Spokes 1 and 2 are in category (c) mission and technical requirements not yet fully defined
- However, the ***Ecosystem*** and IRI depend crucially on HPDF to be successful

ESnet

- Crucial to DOE SC science programs
  - Engaged deeply in preparing for future challenges
  - Deep insights into networking needs as science disciplines evolve and next-gen instruments come online
    - HEP: LHC upgrades - tenfold increases in data rates, Rubin – 3 Gpixel images on minute time-scales
    - BES light sources and new capabilities in NP and BER necessitate robust network solutions
    - FES, real-time control of fusion experiments across continents - dynamic and high-speed networking

ESnet (2)

- Use cases show significant shifts in how science is done: new modes of integrated science will depend on advanced network services
  - Rapid, real-time data analysis, computational and experimental facilities, users widely distributed geographically
  - Novel AI-workflows using multiple user facilities
  - AI-enabled scientific inferences from widely distributed data sources
- ESnet-7 envisioned for 2027 timeframe
  - Net/data services require advances in AI techniques
  - More nimble, intelligent network functions and data-centric services
  - Dynamically defined network configurations required for experiments controlled by computation - rapidly cycle through parameter studies
  - ***Category (a) ready to initiate construction***
- Esnet-8
  - Continue above; infrastructure overhaul, optical fiber IRU lease renewals.
  - Further development of AI technologies to control networks for operations
  - ***Category (b) significant science/engineering challenges to resolve***

Without Esnet, the entire vision collapses; none of the facilities, nor the integrated Ecosystem, nor IRI could function.

Conclusions

- ASCR and SC programs internationally leading in both capabilities of the computational and networking facilities and breadth of the science programs that depend vitally on them
  - Both across DOE and other agencies
- Deep changes in research for all 5 SC offices are anticipated, due to the increasing integration of research disciplines needed to address
- All five ASCR facilities are ***absolutely essential*** to DOE SC, and other agencies. ***They cannot fail.***
- The entire set of facilities should be seen as integrated ***Ecosystem*** to support changing needs of science
  - New governance models are needed
- A **significant R&D and Workforce Development** effort is required to **maintain leadership** in DOE and broader national science competitiveness, industry, and overall economic competitiveness
  - DOE cannot go it alone
  - **All-of government approach for success**

Additional Comments from Subcommittee??

Backup

Process

January–February: Subcommittee holds virtual meetings to assemble, develop strategy to respond to charge, and provide guidance to the five labs.

February 23: Subcommittee co-chairs met with Jack Dongarra, author of prior ASCR report, to gather additional context.

February 29 and March 1: Subcommittee conducts virtual interviews with leadership from the five labs.

March 7: Subcommittee meets in Washington to debrief interviews and begin drafting report.

March 27 and 29: Subcommittee meets with other Office of Science subcommittees to discuss commonalities in charges and findings.

March–April: Committee refines report.

April 19: Subcommittee meets at ANL to continue refinement of report.

April 20–May 1: ASCAC member Roscoe Giles provides preliminary review and feedback.

April 20–May 17: Subcommittee finalizes report.

May 22: Final report delivered to ASCAC

May 29: ASCAC meets in Washington, D.C. and votes on report.

By May 31: Report delivered to Office of Science.

However…

We regard Recommendation 1 as *necessary but not sufficient for success*.

- Additional points must be taken into account
  - the five ASCR computational facilities should be viewed and further developed as an integrated **ecosystem**.
  - as the *de facto* high-end national computing infrastructure, they together serve other parts of DOE and many other national research agencies
  - finally, a comprehensive multi-agency R&D program will be needed, with vendors and industry, if we are to be successful in maintaining international leadership in science and technology.

Science-driven Imperative (2)

- Conclusion
  - R&D program needed to define
    - Next generation systems
    - How they are integrated
    - Services needed to support
  - Must involve science communities
    - These communities are not only DOE SC communities
    - Include those primarily supported by NNSA, NSF, NIST, NASA, NIH, …
  - All hands on deck….

Science-driven Imperative

- Interviewed chairs of all SC subcommittees
  - BES, BER, HEPAP, NP, FES
- Results:
  - All ASCR facilities essential for their science
  - Methodologies of science are rapidly changing
    - Collaborative, interdisciplinary
    - Convergence of theory, experiment, data
    - AI/ML now as essential as Mod/Sim
  - Experimental facilities will generate orders of magnitude more data
    - New approaches are needed to facilitate integration

All Hands On Deck Approach Needed

TOWARD AN INTEGRATED COMPUTING AND DATA FACILITY (THE ECOSYSTEM)

The ***Ecosystem***

- Is under ASCR, but we urge SC leadership to…
  - Work with ASCR and other SC offices
    - on how best to manage and fund it
  - Work with NNSA and key agencies outside DOE, on
    - how best to collaborate on use and R&D needed to further develop it.
- We suggest new governance model
  - Form a higher-level coordination body
  - Include lab directors and ensure that future generations are distinct and innovative, addressing risks of uniformity
- The Ecosystem should be dynamic
  - If major new computing facilities are contemplated beyond the five, we urge that their integral role be considered

IRI (2)

- IRI plans developed via collaborative taskforce from ALCF, OLCF, NERSC, and ESnet
  - More than 170 experts from national laboratories joined to draft the IRI Architecture Blueprint Activity
  - Coordinated SC-wide strategy
  - High-Performance Data Facility (HPDF), spanning two sites, emerged as a crucial node in linking the ***Ecosystem*** and supporting all SC experimental and science programs.
- Initial IRI planning aligns with key aspects of our recommendations
  - Testbed activities anchoring R&D to find optimal scientific solutions
  - Common governance model for the ASCR Ecosystem's facility components

The Ecosystem

RECOMMENDATION 2

Recommendation 2

- Science demands integration. We advocate viewing ASCR facilities not as isolated entities, but as *integral components of a single, larger integrated computational **Ecosystem***
  - with a single governance model.
- This will require new ways of governing and potentially funding the overall ***Ecosystem***
  - should not be developed via individual site procurements.
- Rather it should be designed, developed, built, and operated as an integrated facility Ecosystem for DOE science. It is critical for supporting SC science programs, along with additional software, algorithm, workforce, and science application components, to serve science and engineering research.
- Further, this integrated ***Ecosystem*** is required for programs of other agencies, and industry.
- Its importance to the entire national scientific and technological capability, and its importance as a model internationally, cannot be overstated.

Vision and R&D Pathways

- Driven by needs of science, ALCF, OLCF, NERSC, HPDF, and ESnet should be viewed as, deeply integrated into, and operated as an integrated ASCR facilities ecosystem
- This transformation of five component facilities into a single integrated ecosystem has begun, *but substantial R&D will be required as it evolves over the next decade*
- We refer to this group of facilities and their ongoing integration as the ***Ecosystem***
  - Perhaps should be named in some way

Informs integrated facility on decadal timescale

COMPREHENSIVE R&D PROGRAM