



U.S. DEPARTMENT OF
ENERGY

Office of
Science

THE DOE OFFICE OF SCIENCE

Integrated Research Infrastructure Architecture Blueprint Activity

FINAL REPORT
2023



THE DOE OFFICE OF SCIENCE

Integrated Research Infrastructure Architecture Blueprint Activity

FINAL REPORT
2023

Cover Image: Using data transported via ESnet and analyzed in near-real time at NERSC from experiments performed at SLAC, this image represents small-molecule serial femtosecond crystallography, where a free-electron X-ray laser captures the structure of a nanocrystal just before the sample is destroyed. This collaboration is an example of a superfacility implementation that is revolutionizing scientific research.

Image Credit: J. Nathan Hohman. Illustration: Ella Maru Studio.

Creative: Samantha Trieu, Creative Services, Lawrence Berkeley National Laboratory

Editorial: Carol Pott, Computing Sciences Area, Lawrence Berkeley National Laboratory

DOI: 10.2172/1984466

23-CSA0-9059

Table of Contents

Executive Summary	1
The DOE Office of Science IRI Vision and Framework	3
1. Goals and Philosophy of the ABA	4
Goals and Objectives	4
Foundational Precursor Activities	4
Participants, Organization, and Leadership	5
Philosophy of Conduct of the ABA	6
2. Structured Steps Towards the IRI Culminating Insights	7
Timeline and Approach	7
Gathering Insight Phase	7
Design Phase	10
3. Culminating Insights and Artifacts	11
IRI Patterns and Their Architecture Blueprints	11
Time-Sensitive Patterns	11
Data Integration-Intensive Patterns	12
Long-Term Campaign Patterns	13
IRI Practice Areas Artifacts	14
Resource Co-Operations	14
Cybersecurity and Federated Access	14
User Experience	14
Workflows, Interfaces and Automation	14
Scientific Data Lifecycle	14
Overarching IRI Principles	15
Governance and Steering Considerations	15
Pattern Blueprint Compare and Contrast	15
Conclusions	15
References	15
Summary Presentations	16
Appendix A: Leadership and Participants	17
Roster of Leaders and Participants	18
Leadership Group	18
Participants	19
Headquarters Coordination Group	19
Appendix B: IRI ABA Participant Kickoff Webinar Presentation	22
Appendix C: Synthesis of Sprint 1 insights	33
Appendix D: IRI ABA Project Survey	38

Appendix E: The Integrated Research Infrastructure Architecture Blueprint Activity (IRI ABA) Launching the Design Phase	43
Appendix F: The Integrated Research Infrastructure Architecture Blueprint Activity (IRI ABA) Design Phase Convergence Iteration Event	53
Appendix G: IRI ABA Blueprint: Time-Sensitive Patterns	59
Appendix H: IRI ABA Blueprint: Data Integration–Intensive Patterns	83
Appendix I: IRI ABA Blueprint: Long-Term Campaign Patterns	90
Appendix J: IRI ABA Design Phase: Practice Group on Resource Co-Operations	101
Appendix K: IRI ABA Design Phase: Practice Group on Cybersecurity and Federated Access	110
Appendix L: IRI ABA Design Phase: Practice Group on User Experience	115
Appendix M: IRI ABA Design Phase: Practice Group on Workflows, Interfaces, and Automation	122
Appendix N: IRI ABA Design Phase: Practice Group on Scientific Data Life Cycle	128
Appendix O: IRI ABA Design Phase: Practice Group on Portable/Scalable Solutions	137
Appendix P: IRI ABA Design Phase: Focus Area on Overarching IRI Principles	143
Appendix Q: IRI ABA Design Phase: Focus Area on Governance/Steering Approaches	148
Appendix R: IRI ABA Design Phase: Focus Area on Comparing and Contrasting the Pattern Blueprints	152
Appendix S: Compendium of SC and National Reports Relevant To An Integrated Research Infrastructure/Ecosystems Approach, as of February, 2022	156

Executive Summary

The complexity of scientific pursuits is increasing rapidly with aspects that require dynamic integration of experiment, observation, theory, modeling, simulation, visualization, machine learning (ML), artificial intelligence (AI), and analysis. Research projects across the Department of Energy (DOE) are increasingly data and compute intensive. Innovative research teams are accelerating the pace of discovery by using high-performance computational and data tools in their research workflows and leveraging multiple research infrastructures.

Additionally, several recent high-level U.S. government reports underscore the necessity of a new advanced computing ecosystem for international competitiveness and national security. International competitors are moving forward with major research infrastructure integration efforts that seek to capture a competitive advantage in the global innovation race. Owing to its unparalleled constellation of world-class experimental and observational facilities and high-performance and extreme-scale computational, data, and networking infrastructure, DOE is positioned to be a global leader in this new era of integrated science. However, this new integration paradigm will demand continuing evolution to ensure the U.S. remains a global leader in research and innovation.

The DOE Office of Science (SC) has seized on the strategic importance of integration and has adopted a **vision for Integrated Research Infrastructure (IRI)**: *To empower researchers to meld DOE's world-class research tools, infrastructure, and user facilities seamlessly and securely in novel ways to radically accelerate discovery and innovation.* To respond to the evolving computational requirements of research and the competitive international innovation landscape, experimental facilities could be connected with high performance computing resources for near real-time analysis, and resources should be provided for merging enormous and diverse data for AI/ML techniques and analysis.

Implementing the IRI vision requires the creation of an integrated research ecosystem that transforms science via seamless interoperability. Today, many promising efforts and celebrated achievements prove the efficacy of “point-to-point” and lab-localized solutions to multifacility science problems. However, the enormous growth of integrative science requires a new holistic approach that minimizes duplication and maximizes efficiency to enable solutions to scale across disciplines and domains. Developing a comprehensive IRI strategy is essential to maximizing DOE investment and achieving the scientific potential of this emerging space.

In 2022, SC leadership directed the Advanced Scientific Computing Research (ASCR) program to conduct the Integrated Research Infrastructure Architecture Blueprint Activity (IRI ABA) to produce a reference framework to inform a coordinated, SC-wide strategy for IRI. This activity convened the SC science programs and over 150 DOE national laboratory experts from all 28 SC user facilities across 13 national laboratories to consider the technological, policy, and sociological challenges to implementing IRI.

Through a series of cross-cutting sprint exercises facilitated by the IRI ABA Leadership Group and peer facilitators, participants produced an IRI Framework based on the IRI Vision (see callout below) comprising IRI Science Patterns spanning DOE science domains, IRI Practice Areas needed for implementation, IRI blueprints that connect Patterns and Practice Areas, and overarching principles for realizing the DOE-wide IRI ecosystem. The resulting IRI framework and blueprints provide the conceptual foundations to move forward with organized, coordinated DOE implementation efforts.

Next Steps for Implementation

At the dawn of the exascale science era, many researchers and collaborations strive to meld data, simulation, and AI tools in novel ways, some with strict operational demands. Agency and program leaders feel the urgency to bring the best-integrated science approaches to bear on our greatest challenges. This Final Report of the cross-SC IRI ABA effort provides the scientific, technical, and organizational framework to create and sustain a more fully integrated DOE discovery and innovation ecosystem.

The following immediate and long-term steps form the basis of an implementation plan for the enhanced IRI computational and data infrastructure:

Governance and organization:

- Establish an IRI governance and steering structure to implement the IRI framework and ensure clear principles of engagement among IRI stakeholders, including DOE SC Programs, the DOE research and infrastructure communities, and related federal efforts.
- Establish field-level IRI practice groups responsible for implementing the technical and operational elements of the IRI framework.

Infrastructure:

- Planning: Develop the reference implementations for key IRI science and design patterns.
- Development: Build a test and development environment for IRI research and development.
- Deployment: Deploy high-performance data infrastructure that enables distributed and resilient operations to conduct IRI-integrated science.

Integrated operations:

- Enhance cross-SC operational integration and resilience of high-reliability computing and data infrastructure and services.
- Create common authentication/authorization security frameworks.
- Systematize interfaces across tools, infrastructure, and facilities.
- Standardize approaches across computing environments for allocations, application portability, and user services.

The DOE Office of Science IRI Vision and Framework

The IRI Vision: To empower researchers to meld DOE’s world-class research tools, infrastructure, and user facilities seamlessly and securely in novel ways to radically accelerate discovery and innovation. “Simple and powerful” is the mantra: researchers will benefit from an operational environment that is intuitive and simple to use yet extraordinarily powerful in accelerating discovery.

The IRI Framework is the product of the IRI ABA activity described in this Final Report. The IRI Framework provides a structured means to create an innovative and robust integrative scientific ecosystem for DOE researchers and the broader scientific community, leveraging and maximizing the impact of DOE’s world-class infrastructure, technologies, and expertise.

The key organizing elements of the IRI Framework are **Science Patterns** and **Practice Areas**.

IRI Science Patterns are broad classes of integrated research workflows with common driving features. Each Science Pattern represents a spectrum of DOE science domains and will benefit from a strategic and coordinated approach to design and solution. A given workflow case may span several Science Patterns.

- **Time-sensitive patterns** have urgency, requiring real-time or end-to-end performance with high reliability, e.g., for timely decision-making, experiment steering, and virtual proximity.
- **Data integration–intensive patterns** require combining and analyzing data from multiple sources, e.g., sites, experiments, and/or computational runs.
- **Long-term campaign patterns** require sustained access to resources over a long period to accomplish a well-defined objective.

IRI Practice Areas are cross-cutting communities of practice whose efforts will be essential to advance robust and extensible IRI designs and solutions.

- **User experience practice** will ensure relentless attention to user perspectives and needs through requirements gathering, user-centric (co)-design, continuous feedback, and other means.
- **Resource co-operations practice** is focused on creating new modes of cooperation, collaboration, co-scheduling, and joint planning across facilities and DOE programs.
- **Cybersecurity and federated access practice** is focused on creating novel solutions that enable seamless scientific collaboration within a secure and trusted IRI ecosystem.
- **Workflows, interfaces, and automation practice** is focused on creating novel solutions that facilitate the dynamic assembly of components across facilities into end-to-end IRI pipelines.
- **Scientific data life cycle practice** is focused on ensuring that users can manage their data and metadata across facilities from inception to curation, archiving, dissemination, and publication.
- **Portable/scalable solutions practice** is focused on ensuring that transitions can be made across heterogeneous facilities (portability) and from smaller to larger resources (scalability).

IRI Blueprints and Overarching Principles. IRI ABA produced blueprints – one for each Science Pattern that addresses all Practice Areas and a set of overarching principles and governance considerations for how IRI should be implemented.

1. Goals and Philosophy of the ABA

The ABA was the Office of Science’s first SC-wide convening on IRI. Linking different user facilities together is a difficult task because integration of operations across multiple research infrastructures poses interwoven technological, policy, and sociological challenges that counter conventional practices. For example, the 28 SC user facilities are each independent enterprises, sponsored, funded, managed, and operated as independent facilities. With these considerations in mind, the SC sought broad engagement with an eye to stimulating a variety of structured conversations that cut across established facility, program, institutional, and domain boundaries.

Goals and Objectives

The activity’s overarching objective was to produce the reference conceptual foundations to inform a coordinated “whole-of-SC” strategy for an integrative research ecosystem.

The organizers’ approach to achieving this goal was to:

1. Invite DOE experts across the SC user facilities, national laboratories, and key enterprise stakeholders, to participate in a series of activities and events.
2. Gather and analyze integrative use cases that inclusively span SC programs and user facilities.
3. Develop overarching design principles and one or more “architecture blueprints” that will address the chief IRI design patterns effectively.

Foundational Precursor Activities

Numerous research projects, demonstrations, pilots, workshop reports, infrastructure requirements reviews, and conversations across the SC enterprise over the past several years informed this activity:

- A compendium of SC and national reports (see [Appendix S](#)) identifies a large number of reports with IRI-relevant priority science and technology drivers.
- In June 2019, the directors of the ASCR and Basic Energy Sciences (BES) user facilities met for a one-day information-sharing session at Lawrence Berkeley National Laboratory (Berkeley Lab). The directors formed a joint ASCR-BES working group to explore integration concepts; that group delivered a white paper later that year. In parallel, the BES Light Sources Data Working Group developed concepts for integrated computation and data infrastructure.
- In FY 2021, an SC Integrated Computation and Data Infrastructure research funding activity was established. Subsequently, ASCR released a Funding Opportunity Announcement (FOA) and supported a small number of integration projects.
- Of particular importance, in 2020, DOE/ASCR convened an ASCR facilities IRI task force to develop a vision and principles for integrating ASCR facilities and connecting these to other facilities, capabilities, and users across the DOE complex to accelerate research and discovery. This effort produced a seminal white paper in March 2021 titled, *Toward a Seamless Integration of Computing, Experimental, and Observational Science Facilities: A Blueprint to Accelerate Discovery*¹. The IRI ABA built on the concepts and findings in this Final Report and expanded the scope to engage the whole of SC.

¹ “Towards a Seamless Integration of Computing, Experimental, and Observational Science Facilities: A Blueprint to Accelerate Discovery,” DOE ASCR IRI Task Force white paper, Mar 8, 2021, <https://doi.org/10.2172/1863562>.

- At the September 29, 2021, Meeting of the Advanced Scientific Computing Advisory Committee (ASCAC), ASCR Facilities Division Director Ben Brown laid out a vision for the ASCR facilities enterprise² that set the conceptual stage for an integrated ecosystem approach to SC research infrastructure coupling experimental, observational, computing, data, and networking facilities and resources.

Participants, Organization, and Leadership

The organization of effort for the activity was as follows:

- **ASCR executive leaders:** The ASCR Facilities Division provided executive leadership for the activity.
- **Leadership Group:** Chaired by the ASCR executive leaders, a core team of eight DOE national laboratory subject matter experts spanning SC program office mission areas and facilities coordinated and steered the IRI ABA activities.
- **Headquarters Coordination Group:** Chaired by the ASCR executive leaders, a group across SC program offices met virtually to share IRI-related programmatic priorities, align objectives, and review outcomes.
- **Participants:** A total of over 150 staff members across the DOE national laboratories participated in the IRI ABA activity, spanning science and technology backgrounds, research and facilities foci, and national laboratory membership.

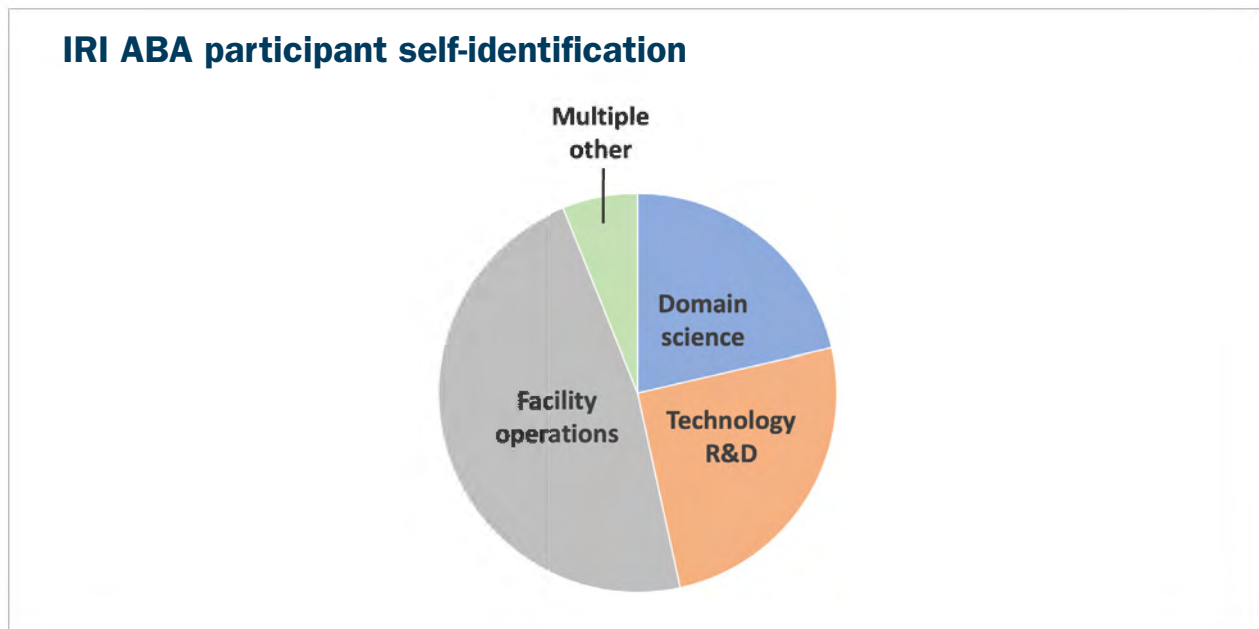


Figure 1: Participant demographics in terms of expertise/activity area

Please refer to [Appendix A](#) for the complete roster of leaders and participants, and their institutional affiliations.

² Benjamin Brown, “A Vision for the ASCR Facilities Enterprise,” presentation to the ASCAC, Sept 29, 2021, https://science.osti.gov/-/media/ascr/ascac/pdf/meetings/202109/Brown_ASCR_Facilities_Vision_202109.pdf.

Philosophy of Conduct of the ABA

The organizers emphasized an all-of-SC approach throughout the conduct of the ABA. The organizers operated with explicit guiding principles of inclusivity, focus on science and the end user, emphasis on identifying frameworks and common underlying patterns rather than on devising technical solutions, and agility of purpose (see Figure 2):

- **Inclusive, cross-cutting participation.** To achieve success, the activity was designed to engage a broad set of stakeholders spanning the SC complex, including the research community, user facilities, and national laboratories, which have challenging computational/data workflows that are a high priority for SC programs. IRI ABA organized activities in ways that avoided groups of participants assembled around usual program/community silos.
- **Focus on cross-cutting scientific use cases and patterns.** The activity emphasized progress and pattern identification through iterative gathering and synthesis of information and perspectives. Each iteration was anchored on identifying canonical IRI patterns and IRI “modes” that spanned multiple science domains and integrative workflows. In turn, these IRI science patterns informed the last stage of the ABA: the framing of prospective IRI design patterns.
- **Emphasis on user needs and perspectives.** Prioritization of user perspectives was central to the activity to avoid premature fixation on prospective technical solutions. The organizers adopted a user-centered approach that emphasized listening to end-user, scientist-provider, and technologist perspectives to derive common IRI capability gaps, requirements, and design patterns and to inform the ultimate artifacts and conclusions.
- **Agile “sprint” activities emphasizing frameworks, not solutions.** The organizers adopted an urgency-oriented nimble approach to convening cross-cutting groups of participants to accomplish short sprint activities. The mantras were that 80% quality is good enough to keep moving forward, and the IRI ABA was about common framings, not “what gets built.”

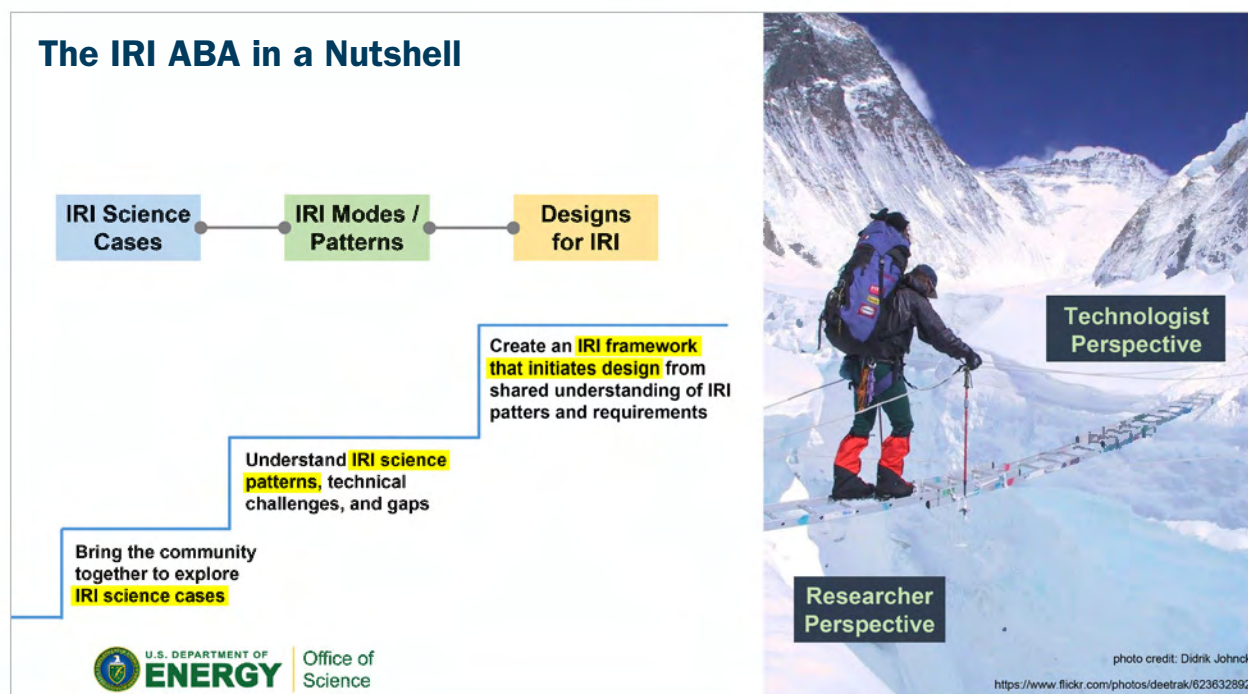


Figure 2: The IRI ABA aims to provide the conceptual foundations to inform strategy for an integrative research ecosystem

2. Structured Steps Towards the IRI Culminating Insights

Timeline and Approach

An intensive array of activities was undertaken from January to September 2022 (see Figure 3). Beginning with eliciting IRI science cases and user needs and challenges across the DOE national laboratory complex in a gathering insight phase, the ABA identified primary IRI modes and patterns. It carried them forward into a design phase to craft a set of framing blueprints and supporting documents.

Details are provided in the [IRI ABA Participant Kickoff Webinar Presentation](#).

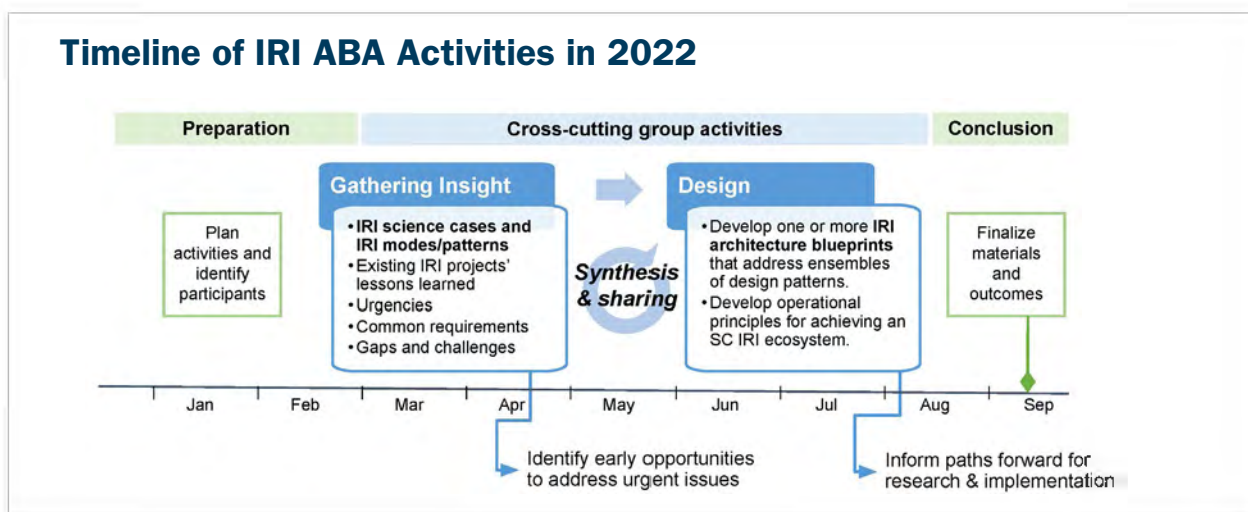


Figure 3: An overview of the IRI ABA activities

Gathering Insight Phase

This phase focused on gathering insights from across the DOE complex relevant to conceiving and developing an IRI framework. Over 110 participants from DOE national laboratories were organized into 10 cross-cutting groups, each composed of scientists, technologists, infrastructure developers and operators, and others from a range of SC programs and facilities.

The gathering insight phase consisted of two “sprint” activities.

- Sprint 1 was designed as a user-centered activity focused on identifying relevant IRI science cases and user experiences and challenges with integration. This led to vital insights that helped frame the next sprint.
- Sprint 2 gathered insights, lessons, and patterns from existing IRI-relevant SC integration projects, activities, and initiatives already underway or completed by IRI ABA participants.

For Sprint 1, a unique methodology was developed for eliciting and listening closely to “science case voices” related to the challenges of IRI. Care was taken in developing and testing the questionnaire to allow interviewees to express in their own words their expectations, wishes, and challenges in performing their scientific activity. Each of the 10 gathering insight groups identified and solicited candidates for interviews from among the SC scientific programs and facilities. Attention was given to ensuring broad coverage of a variety of science cases across SC science domains (see Figure 4). A total of 30 live scientist/user interviews were conducted using a standard questionnaire across the 10 groups, resulting in nearly 30 hours of recordings and many sets of notes. Each group developed insights from its interviews which were collectively summarized across groups by the

leadership group. These summaries resulted in proposed themes and topics for a deeper exploration of common and diverse patterns, issues, and challenges.

[Synthesis of Sprint 1 insights](#) (released to IRI ABA participants for feedback)

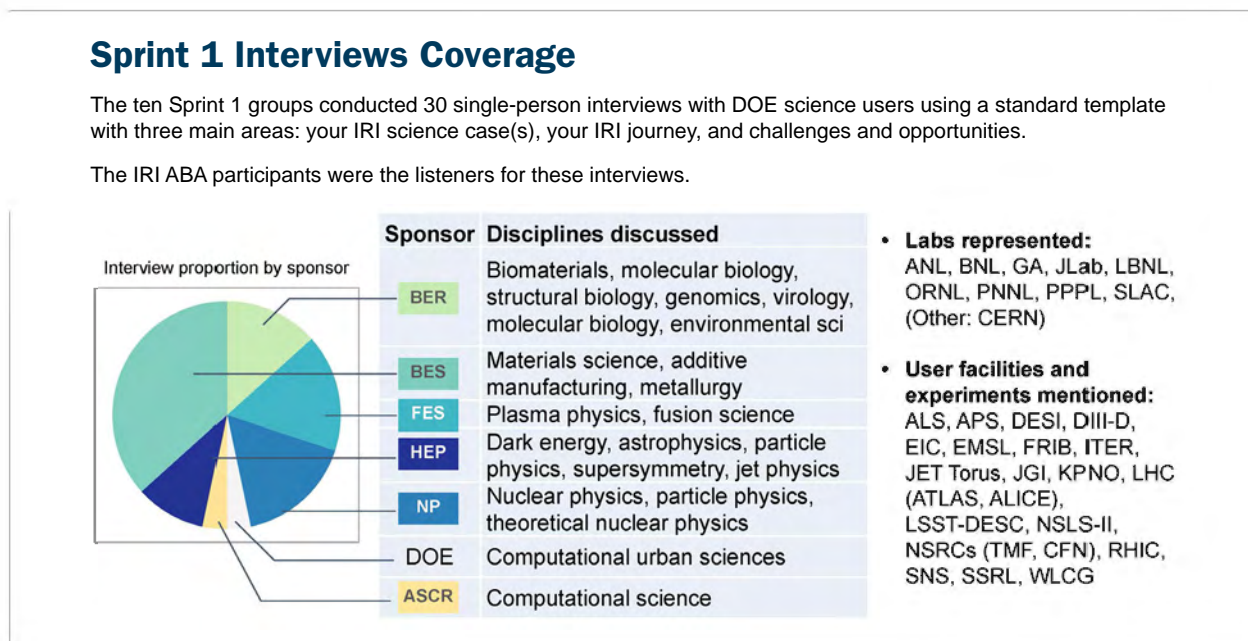


Figure 4: Summary of coverage in Sprint 1 interviews

For **Sprint 2**, a participant survey³ was first conducted to identify IRI-aligned projects. This resulted in 75 unique entries in the following categories: science cases, computing (distributed/grid, on demand, architecture, job management, middleware), data (real-time analysis, multi-modal analysis, management), discovery platforms and integration frameworks, identity management, networking (requirements gathering, traffic load balancing, wireless, APIs), workflows and performance monitoring, and software and applications.

Informed by the survey results, the cross-cutting groups (the same as for Sprint 1) then each held discussion sessions guided by the following focus questions:

- What scope/themes relevant for an IRI have been addressed by existing projects?
- Where do we have collective experience and where do we not?
- Lessons and takeaways from existing projects: Where were the problems, what went well, what was difficult?
- Where do people feel passionately? What ideas or problems elicit strong emotions?
- How do technologists feel they can best interact with the users?
- How can we sustain IRI and keep it alive through communities of practice and fruitful partnerships?

³ See the [IRI ABA Project Survey, Appendix D](#).

The results of these Sprint 2 discussions, also informed by Sprint 1 insights, were collectively synthesized into a concise set of thematic insights (see Figure 5):

- Overarching classes of cross-cutting IRI science patterns that span the SC science domain space.
- Top challenges to realizing the IRI vision.

These insights served as the principle references to structuring the activities of the follow-on design phase and ultimate IRI ABA artifacts.

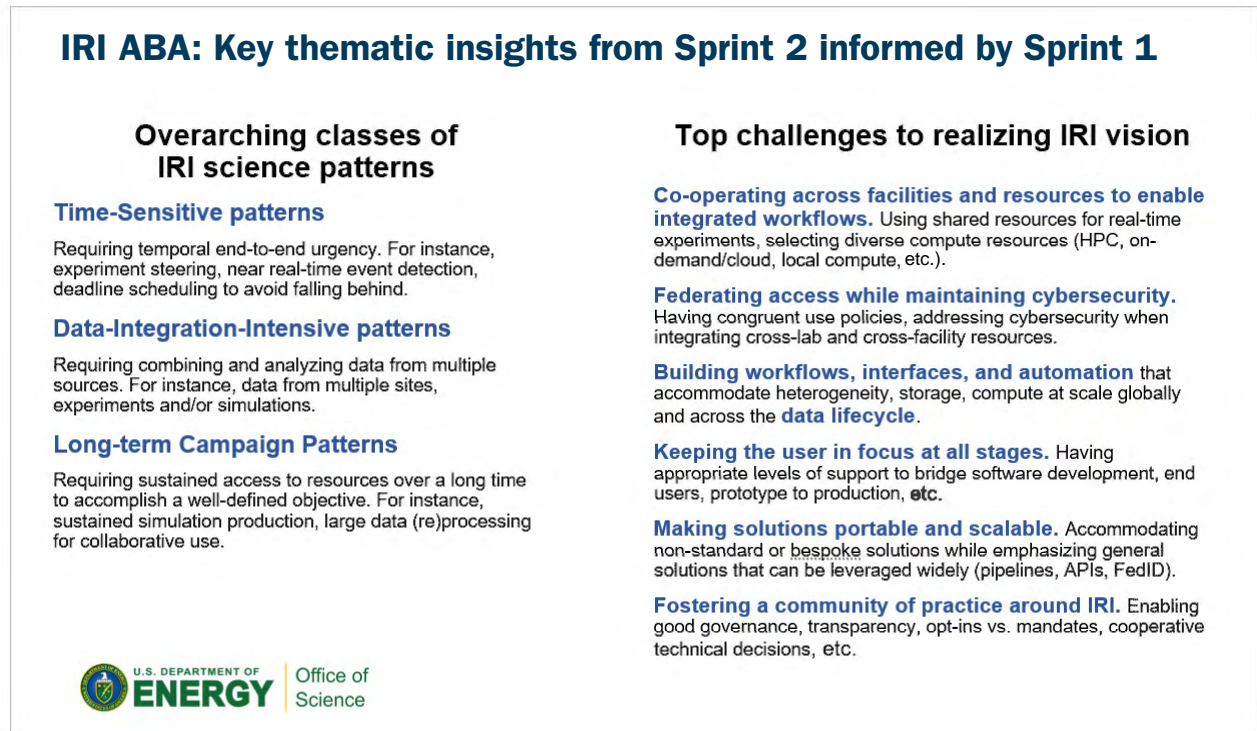


Figure 5: Summary of the key thematic insights from Sprint 2, informed by the results of Sprint 1

Design Phase

In this phase (see Figure 6), groups of participants, including 21 participants who joined after the gathering insight phase, were convened in new working groups representing three **IRI science patterns** and six **IRI practice areas** that were identified in the gathering insight phase:

- IRI science patterns: time-sensitive patterns, data integration–intensive patterns, and long-term campaigns.
- IRI practice areas: resource co-operations; cybersecurity and federated access; workflows, interfaces, and automation; science data life cycle; user experience; and portable/ scalable solutions.

Each group focused on drafting a definitional document for their science pattern or practice area; the three science pattern documents were effectively initial drafts of the science pattern-based architecture blueprints. A virtual “iteration/convergence event” was then held over four days, bringing together all participants of these nine working groups. The iteration/convergence event agenda promoted crosscutting conversations in which participants visited different working groups to refine the practice area documents and architecture blueprints. Inspired by these interactions, the organizers and participants developed three additional focus papers on cross-cutting IRI principles, commonalities and differences in the blueprints, and considerations of potential governance and steering structures.

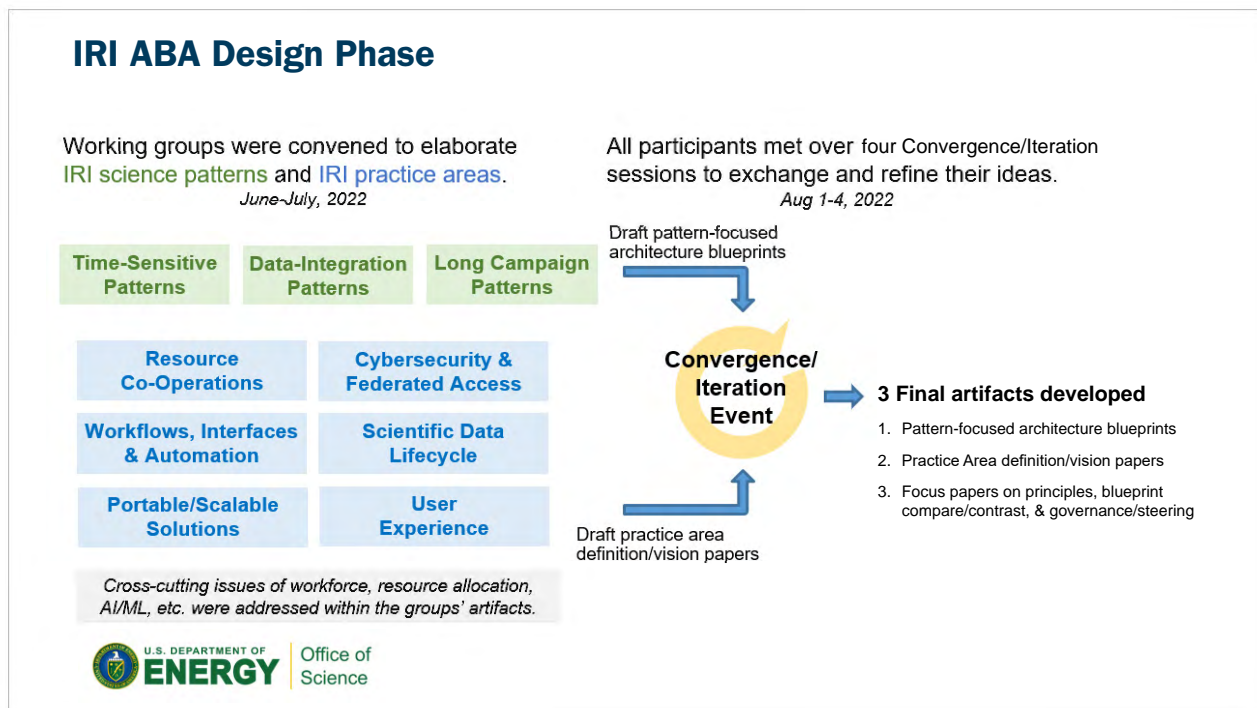


Figure 6: Overview of work performed in the IRI ABA Design Phase

3. Culminating Insights and Artifacts

The design phase produced artifacts in three interconnected areas: pattern blueprints, practice areas, and focus topics. The IRI pattern blueprints together serve as a reference framework for addressing IRI needs and approaches for the major categories of use patterns. The practice areas described the necessary technical and organizational activities and structures needed to support the patterns. The focus topics describe high-level considerations toward advancing and implementing the IRI effort.

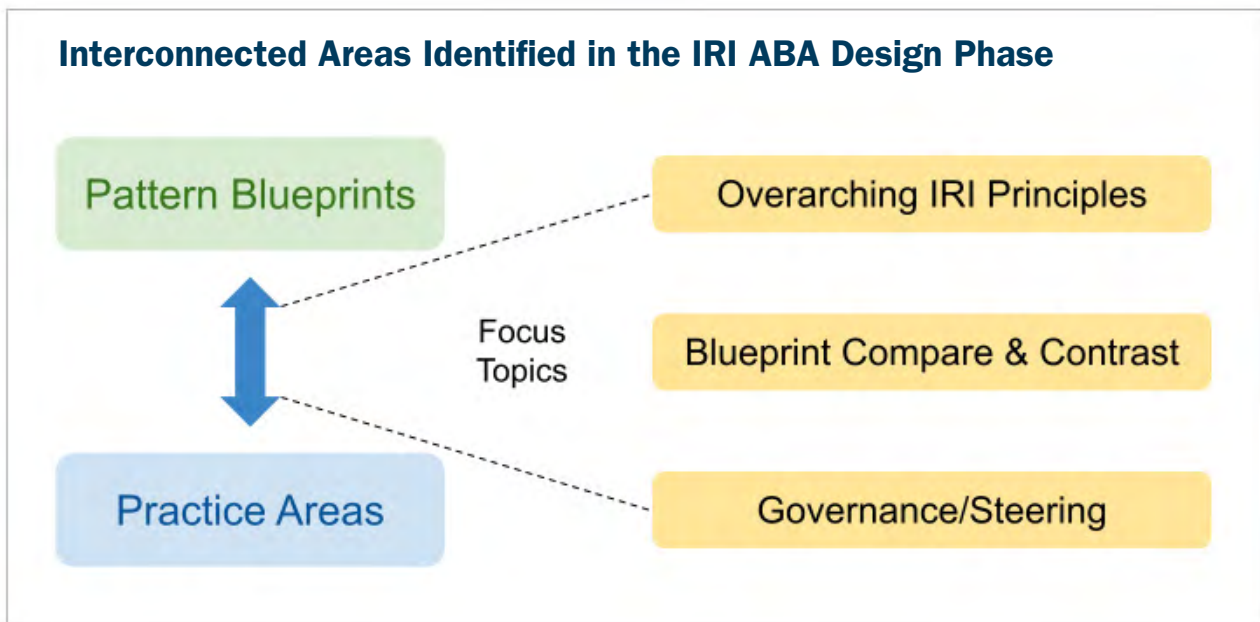


Figure 7: Relationship between the three areas (pattern blueprints, practice areas, and focus topics) identified in the design phase

IRI Patterns and Their Architecture Blueprints

Time-Sensitive Patterns

Time-sensitive IRI patterns comprise workflows with time critical/sensitive requirements (i.e., real time or near real time), which can be motivated by various factors such as timely decision making, experiment steering, virtual proximity, and loss of data fidelity. These time-sensitive workflows involve integration across multiple facilities and resources. They are found in many scientific domains such as beamline-based materials science, astronomy and astrophysics, observational science, and experimental fusion science.

The Time-Sensitive Patterns Group highlighted an ensemble of workflow areas that are important to address for these patterns (and which might also apply to the other IRI patterns): experiment control, distributed systems administration, and data management (see Figure 8).

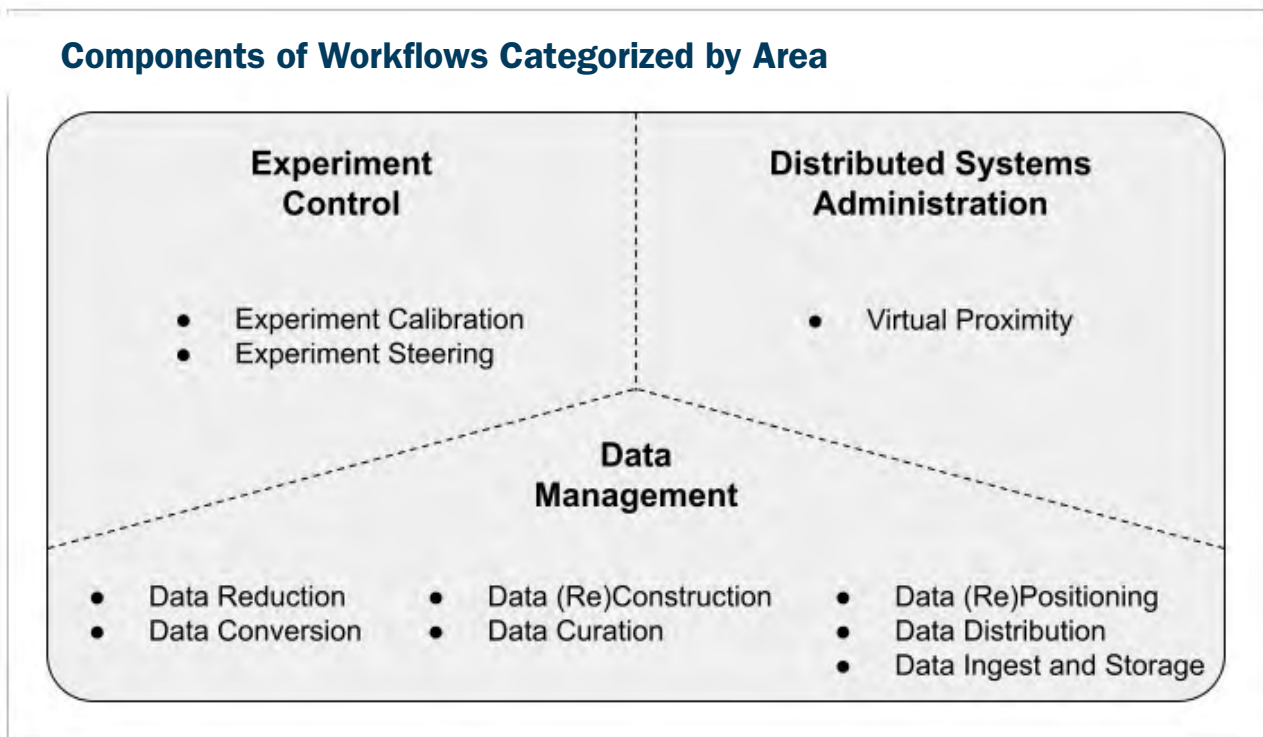


Figure 8: Important workflow areas for time-sensitive patterns

Some high-level takeaways and important next steps from the group discussion included:

- Highlighting “classes” of time sensitivities, e.g., by time periods (ms, sec, mins, hours, days, etc.) and by motivation (decisions that cannot wait, experiment control, loss or fidelity of data, etc.).
- Emphasizing the importance of user experience, e.g., usability, reliability, etc.
- Determining what resource needs to be local versus remote.
- Understanding that time-sensitive workflows may require security enforcement that is time sensitive as well.

Supporting document: [IRI ABA Blueprint: Time-Sensitive Patterns](#)

Data Integration-Intensive Patterns

Data integration-intensive patterns are characterized by a need to perform analysis of data combined from multiple sources, which can include data from multiple sites, experiments, and/or simulations. This can also include tracking metadata and provenance for reproducible science and interactive data analysis, possibly at scale.

The group recognized at least two broad pattern areas within this class:

- Integration of data from simulations and experiments/observations to generate new insight and subsequent direct actions.
- Cross-site data-driven discovery, which includes using similar, multimodal, or heterogeneous data already generated at different facilities, or running the same tool, e.g., simulation software, on different systems, or experimental/observational data originating at different sources, the results of which must be combined, processed, and analyzed.

Some gaps and opportunities from the group discussion included:

- **Gaps:** cross-facility APIs for resource co-operations; common/appropriate resource allocation models; standard abstracted workflow and automation tools; complex-wide data storage and searching capabilities; new models for “wide-area” cybersecurity; common or well-understood data policies; lack of FAIR data; user-focused user experience; lack of portable code; and cross-training of staff (scientific, engineering, support, administrative).
- **Opportunities:** many early-win science opportunities exist for this pattern; common APIs for facilities; standards for metadata; streaming data to/from compute and storage facilities; common and well-understood data policies; support for FAIR data; and templates for portable code.

Supporting document: [IRI ABA Blueprint: Data-Integration-Intensive Patterns](#)

Long-Term Campaign Patterns

This class of science use patterns is characterized by a need for sustained access to resources at scale over a longer time to accomplish a well-defined objective. Robustness, reproducibility, and reliability are important to accomplishing long-term science, and these patterns will likely involve significant logistical planning. Examples include sustained simulation production and large data (re)processing for collaborative use.

The group determined that the key overall challenge is to intentionally plan and coordinate resources between campaigns and facilities over time. Specific challenge areas include long-term storage past the end of a project; a present mismatch between the short-term nature of resource allocations and mechanisms (i.e., compute and instrument time) versus the long-term needs of a campaign; the continual evolution of technologies and approaches within facilities and campaigns (e.g., computing architectures, infrastructure and instruments, cybersecurity, workflow systems); avoiding interruptions in campaigns due to facility downtimes; and the present lack of holistic (all-of-SC) approaches to resource allocations.

Additional perspectives from the group included:

- The evolution of staffing duration of a campaign needs to be factored in.
- The facilities may provide infrastructure and also need to accommodate the varied data management requirements of the programs and research domains.
- Abstraction layers are likely to be a key and pervasive component of the solutions we need.
- A common machine-usable interface to facilities looks like it is a prerequisite, as is scheduling that does not have humans in the loop.
- Data re-use is an appealing idea but difficult to achieve, so data needs to be well-described and documented to be useful down the road.

Supporting document: [IRI ABA Blueprint: Long-Term Campaign Patterns](#)

IRI Practice Areas Artifacts

In addition to the three IRI pattern blueprints described previously, IRI ABA teams developed a set of briefs for six practice areas that were identified as critical technical and operational areas needed to enable the IRI science use patterns.

Resource Co-Operations

Allocations/provisioning of multiple heterogeneous resources across multiple facilities for large collections of scientific programs must be aligned in time and planned. IRI requires new levels of cooperation, collaboration, co-scheduling, and joint planning across facilities and across DOE programs.

Supporting document: [IRI ABA Design Phase: Practice Group on Resource Co-Operations](#)

Cybersecurity and Federated Access

Users require a distributed research infrastructure with seamless access and consistent services while the infrastructure must be operated according to cybersecurity requirements and policies set at the federal level. Operators of user facilities also have different missions, and thus different requirements, across the lab complex. Balancing these constraints can also lead to sources of impedance. Novel secure design patterns and architectures will be required to support open science-integrated architecture for seamless scientific collaboration.

Supporting document: [IRI ABA Design Phase: Practice Group on Cybersecurity and Federated Access](#)

User Experience

Understanding evolving users' needs and experiences is critical for technologists to develop effective IRI solutions. This area is central for building an effective IRI. Strategies for enabling users, including requirements gathering, user-centric (co)-design, liaising approaches, and related topics, have been proposed. This topic has implications for all other practice areas.

Supporting document: [IRI ABA Design Phase: Practice Group on User Experience](#)

Workflows, Interfaces and Automation

Users need to systematically and easily assemble system components to support IRI science cases in the form of end-to-end pipelines. Users should be able to manage these overlays and middleware effectively across facilities.

Supporting document: [IRI ABA Design Phase: Practice Group on Workflows, Interfaces, and Automation](#)

Scientific Data Lifecycle

Users need to manage their data (along with metadata) across facilities from inception to curation, archiving, dissemination, and publication. Technologists need to understand the requirements across different communities to develop solutions appropriate for an IRI and the principles of effective data management to provide a FAIR-based data pipeline with end user-focused interfaces.

Supporting document: [IRI ABA Design Phase: Practice Group on Scientific Data Lifecycle](#)

Portable / Scalable Solutions

Users and technologists need their applications to move/translate across heterogeneous facilities (be portable) and go from smaller to larger resources (be scalable).

Supporting document: [IRI ABA Design Phase: Practice Group on Portable/Scalable Solutions](#)

Focus Topics Artifacts

During the design phase activity, participants developed papers on three focus topics to provide additional context and high-level considerations for a future IRI strategy: overarching IRI principles, governance and steering considerations, and a comparative analysis of the three IRI patterns.

Overarching IRI Principles

For IRI projects to be successful, as they contribute to large-scale infrastructure, it is vital for the community to agree on a set of foundational principles. The principles articulated stem from respect for the users of the IRI and for the facilities. The overarching principles captured in the document reflect experience in understanding what makes for effective and persistent infrastructure that can be deployed, maintained, and used.

Supporting document: [IRI ABA Design Phase: Focus Area on Overarching IRI Principles](#)

Governance and Steering Considerations

Achieving an operational IRI depends on DOE's facilities and their users, researchers and their projects, and science communities having the right incentives, governance, and operating structure. We envision a governance structure that would include a policy body and working groups to cover technical aspects such as standards, evaluation, and cyber hygiene.

Supporting document: [IRI ABA Design Phase: Focus Area on Governance/Steering Approaches](#)

Pattern Blueprint Compare and Contrast

The three design patterns (time sensitive, data integration, and long-term campaigns) have unique attributes and areas of commonality. From the information gathered in the design phase, the group documented the commonalities and uniqueness of the patterns. The unique areas often pointed out a different strategic focus with respect to the various practice areas, which could arise due to a difference in the level of maturity needed by a pattern. The common areas can be a guide to areas of investment with broad cross-cutting benefits.

Supporting document: [IRI ABA Design Phase: Focus Area on Comparing and Contrasting the Pattern Blueprints](#)

Conclusions

The broad cross-cutting nature of the IRI ABA has demonstrated the specific value that an integrated approach can offer DOE program offices, DOE facilities' users, and staff. The artifacts produced by this activity offer specific directions and framing for what IRI operational and technical capabilities may look like in the future, what the focus areas need to be, and how such integrative capabilities and services might be stewarded and operated across the DOE's varied programs, national laboratories, and user facilities.

References

[Compendium of SC and National Reports Relevant To An Integrated Research Infrastructure/Ecosystems Approach, as of February, 2022](#)

Summary Presentations

Event	Link
Kickoff Webinar Feb 17, Feb 22, 2022	The Integrated Research Infrastructure Architecture Blueprint Activity (IRI-ABA) Orientation for Participants and Stakeholders
Phase1 Reflection and Synthesis May, 2022	IRI Sprint 1 Synthesis Summary
Design Phase Webinar June 15, 2022	The Integrated Research Infrastructure Architecture Blueprint Activity (IRI-ABA) Launching the Design Phase
Convergence Event Aug 1-4, 2022	The Integrated Research Infrastructure Architecture Blueprint Activity (IRI-ABA) Design Phase Convergence Iteration Event

Appendix A

LEADERSHIP AND PARTICIPANTS

Roster of Leaders and Participants

Advanced Scientific Computing Research (ASCR) executive leaders:

Ben Brown, Director, ASCR Facilities Division

Bill Miller, Senior Technical Advisor, ASCR Facilities Division , on detail from the National Science Foundation

Leadership Group

Debbie Bard, National Energy Research Scientific Computing Center (NERSC), Lawrence Berkeley National Laboratory

Amber Boehnlein, Thomas Jefferson National Accelerator Facility

Kjiersten Fagnan, Joint Genome Institute, Lawrence Berkeley National Laboratory

Chin Guok, Energy Sciences Network (ESnet), Lawrence Berkeley National Laboratory

Eric Lançon, Scientific Data and Computing Center, Brookhaven National Laboratory

Sreeranjani “Jini” Ramprakash, Argonne Leadership Computing Facility, Argonne National Laboratory

Arjun Shankar, Oak Ridge Leadership Computing Facility, Oak Ridge National Laboratory

Nicholas Schwarz, Advanced Photon Source, Argonne National Laboratory

The image shows a group portrait of the leadership team for the SC Integrated Research Infrastructure Architecture Blueprint Activity. The title is "SC Integrated Research Infrastructure Architecture Blueprint Activity Leadership". The team is divided into two groups: "HQ Executive Leadership" and "IRI-ABA Leadership Group".

HQ Executive Leadership		IRI-ABA Leadership Group				
 Ben Brown Director ASCR Facilities Division	 Debbie Bard Group Lead for Data Science Engagement NERSC, LBNL	 Amber Boehnlein Chief Information Officer JLab	 Kjiersten Fagnan Chief Informatics Officer JGI, LBNL	 Chin Guok Group Lead for Planning and Architecture ESnet, LBNL		
 Bill Miller Senior Technical Advisor ASCR Facilities Division	 Eric Lancon Director, Scientific Data and Computing Center BNL	 Jini Ramprakash Deputy Division Director ALCF, ANL	 Arjun Shankar Section Head, Advanced Technologies OLCF/NCCS, ORNL	 Nicholas Schwarz Group Leader, Scientific Software Eng. & Data Mgmt., APS, ANL		

Figure 9: IRI ABA leadership group

Headquarters Coordination Group

Basic Energy Sciences: **Tom Russell**

Biological and Environmental Research: **Paul Bayer, Jay Hnilo, Resham Kulkarni**

Fusion Energy Sciences: **Josh King, Matt Lanctot**

High Energy Physics: **Jeremy Love, Eric Church**

Isotope Program: **Kristian Myhre**

Nuclear Physics: **Xiaofeng Guo, Jim Sowinski**

Participants

Ghaleb Abdulla, Lawrence Livermore National Laboratory

Corey Adams, Argonne National Laboratory

Ryan Adamson, Oak Ridge National Laboratory

William Allcock, Argonne National Laboratory

Rachana Ananthakrishnan, Argonne National Laboratory

Scott Atchley, Oak Ridge National Laboratory

Carl Bai, Oak Ridge National Laboratory

Nathan Baltzell, Thomas Jefferson National Accelerator Facility

Jaydeep Bardhan, Pacific Northwest National Laboratory

Pete Beckman, Argonne National Laboratory

Douglas Benjamin, Brookhaven National Laboratory

Tekin Bicer, Argonne National Laboratory

Doug Bishop, Princeton Plasma Physics Laboratory

Johannes Blaschke, Lawrence Berkeley National Laboratory

Reuben Budiardja, Oak Ridge National Laboratory

Nick Buraglio, Lawrence Berkeley National Laboratory

Paolo Calafiura, Lawrence Berkeley National Laboratory

Stuart Campbell, Brookhaven National Laboratory

Shane Canon, Lawrence Berkeley National Laboratory

Franck Cappello, Argonne National Laboratory

Adam Carlyle, Oak Ridge National Laboratory

Steve Chan, Lawrence Berkeley National Laboratory

Kyle Chard, Argonne National Laboratory

Mathew Cherukara, Argonne National Laboratory

Taylor Childers, Argonne National Laboratory

David Cowley, Pacific Northwest National Laboratory

Sydni Credle, National Energy Technology Laboratory

Eli Dart, Lawrence Berkeley National Laboratory

Zichao Di, Argonne National Laboratory

Markus Diefenthaler, Thomas Jefferson National Accelerator Facility

Bill Dorland, Princeton Plasma Physics Laboratory

Markus Eisenbach, Oak Ridge National Laboratory

Bjoern Enders, Lawrence Berkeley National Laboratory

Christian Engelmann, Oak Ridge National Laboratory

Peter Ercius, Lawrence Berkeley National Laboratory

Evan Felix, Pacific Northwest National Laboratory

Rafael Ferreira da Silva, Oak Ridge National Laboratory

Nicola Ferrier, Argonne National Laboratory

Gina Fisk, Department of Energy Office of Science

Chris Fuson, Oak Ridge National Laboratory

Vincent Garonne, Brookhaven National Laboratory

Lindsey Gray, Fermi National Accelerator Laboratory

William Gustafson, Pacific Northwest National Laboratory

Steve Hammond, National Renewable Energy Laboratory

Marcus Hanwell, Brookhaven National Laboratory

Tom Harper, Pacific Northwest National Laboratory

Damian Hazen, Lawrence Berkeley National Laboratory

Steven Henke, Argonne National Laboratory

Bryan Hess, Thomas Jefferson National Accelerator Facility

Alexander Hexemer, Lawrence Berkeley National Laboratory

Graham Heyes, Thomas Jefferson National Accelerator Facility

Rob Hovsopian, National Renewable Energy Laboratory

Michael Hu, Argonne National Laboratory

Bin Hu, Los Alamos National Laboratory

Shantenu Jha, Brookhaven National Laboratory

Wayne Joubert, Oak Ridge National Laboratory

Kate Keahey, Argonne National Laboratory

Kyle Kelley, Oak Ridge National Laboratory

Rajkumar Kettimuthu, Argonne National Laboratory

Mariam Kiran, Lawrence Berkeley National Laboratory

Michael Kirby, Fermi National Accelerator Laboratory

Ezra Kissel, Lawrence Berkeley National Laboratory

Christopher Knight, Argonne National Laboratory

Hari Krishnan, Lawrence Berkeley National Laboratory

Olga Kuchar, Oak Ridge National Laboratory

Yatish Kumar, Lawrence Berkeley National Laboratory

Jitendra Kumar, Oak Ridge National Laboratory

Shawn Kwang, Lawrence Berkeley National Laboratory

Jerome Lauret, Brookhaven National Laboratory

David Lawrence, Thomas Jefferson National Accelerator Facility

Tom Lehman, Lawrence Berkeley National Laboratory

John MacAuley, Lawrence Berkeley National Laboratory

Kurt Maier, Pacific Northwest National Laboratory

Addi Malviya-Thakur, Oak Ridge National Laboratory

David Martin, Argonne National Laboratory

Don Maxwell, Oak Ridge National Laboratory

Lee Ann McCue, Pacific Northwest National Laboratory

Apurva Mehta, SLAC National Accelerator Laboratory

Verónica Melesse, Oak Ridge National Laboratory

Orso Meneghini, General Atomics

Bronson Messer, Oak Ridge National Laboratory

Antonino Miceli, Argonne National Laboratory

Ben Mintz, Oak Ridge National Laboratory

Kathryn Mohror, Lawrence Livermore National Laboratory

John Morris, Brookhaven National Laboratory

Swagato Mukherjee, Brookhaven National Laboratory

Todd Munson, Argonne National Laboratory

Jeff Neel, Argonne National Laboratory

Matthew Norman, Oak Ridge National Laboratory

Peter Nugent, Lawrence Berkeley National Laboratory

Sarp Oral, Oak Ridge National Laboratory

Hannah Parraga, Argonne National Laboratory

Amedeo Perazzo, SLAC National Accelerator Laboratory

Eric Pouyoul, Lawrence Berkeley National Laboratory

Giri Prakash, Oak Ridge National Laboratory

Lavanya Ramakrishnan, Lawrence Berkeley National Laboratory

Georg Rath, Lawrence Berkeley National Laboratory

Katherine Riley, Argonne National Laboratory

Kelly Rose, National Energy Technology Laboratory

Rob Ross, Argonne National Laboratory

Simon Roux, Lawrence Berkeley National Laboratory

David Schissel, General Atomics

Malachi Schram, Thomas Jefferson National Accelerator Facility

James Sethian, Lawrence Berkeley National Laboratory

Elizabeth Sexton-Kennedy, Fermi National Accelerator Laboratory

John Shalf, Lawrence Berkeley National Laboratory

Stacey Sheldon, Lawrence Berkeley National Laboratory

Frederik Shulz, Lawrence Berkeley National Laboratory

Mike Skwarek, Argonne National Laboratory

Adam Slagell, Lawrence Berkeley National Laboratory

Sterling Smith, General Atomics

Cory Snavelly, Lawrence Berkeley National Laboratory

Adam Stone, Lawrence Berkeley National Laboratory

Christine Sweeney, Los Alamos National Laboratory

Holly Szumila-Vance, Thomas Jefferson National Accelerator Facility

William Tang, Princeton Plasma Physics Laboratory

Jonathan Taylor, Oak Ridge National Laboratory

Greg Tchilinguirian, Princeton Plasma Physics Laboratory

Vivek Thampy, SLAC National Accelerator Laboratory

Jana Thayer, SLAC National Accelerator Laboratory

Rollin Thomas, Lawrence Berkeley National Laboratory

Chris Tracy, Lawrence Berkeley National Laboratory

Thomas Uram, Argonne National Laboratory

Rama Vasudevan, Oak Ridge National Laboratory

Feiyi Wang, Oak Ridge National Laboratory

Andrew Wiedlea, Lawrence Berkeley National Laboratory

Qin Wu, Brookhaven National Laboratory

Kevin Yager, Brookhaven National Laboratory

Xi Yang, Lawrence Berkeley National Laboratory

Qingteng Zhang, Argonne National Laboratory

Jason Zurawski, Lawrence Berkeley National Laboratory

Appendix B

IRI ABA PARTICIPANT KICKOFF WEBINAR PRESENTATION

The Integrated Research Infrastructure Architecture Blueprint Activity (IRI-ABA)

Orientation for Participants and Stakeholders

February 17 and 22, 2022

Outline

- Setting the stage
About the IRI-ABA } **30 minutes**
- Discussion / Q&A **30 minutes**

SC Integrated Research Infrastructure Architecture Blueprint Activity Leadership

HQ Executive Leadership



Ben Brown
Director
ASCR Facilities Division



Bill Miller
Senior Technical Advisor
ASCR Facilities Division

IRI-ABA Leadership Group



Debbie Bard
Group Lead for Data
Science Engagement
NERSC, LBNL



Eric Lancon
Director, Scientific Data
and Computing Center
BNL



Amber Boehlein
Chief Information Officer
JLab



Jini Ramprakash
Deputy Division Director
ALCF, ANL



Kjersten Fagnan
Chief Informatics Officer
JGI, LBNL



Arjun Shankar
Section Head,
Advanced Technologies
OLCF/NCCS, ORNL



Chin Guok
Group Lead for Planning
and Architecture
ESnet, LBNL



Nicholas Schwarz
Group Leader,
Scientific Software
Eng. & Data Mgmt.,
APS, ANL

Code of conduct

The DOE Office of Science (SC) is fully and unconditionally committed to fostering safe, diverse, equitable, and inclusive work, research, and funding environments that value mutual respect and personal integrity.

If anyone involved in this event has questions or concerns, please approach the HQ Executive Sponsors.

more at >> <https://science.osti.gov/sc-2/Research-and-Conduct-Policies/Diversity-Equity-and-Inclusion/SC-Statement-of-Commitment>



4

Setting the stage



Architecture Blueprint Activity

5

What do we mean by “integration” and “architecture blueprints”?

Integrated Research Infrastructure (IRI)

“Integrating” means having the technical, operational, and policy means for *connecting infrastructure compositably and seamlessly*.

Such integration carries major operational implications for the variously connected research infrastructures. This complexity is why broad participation/representation across DOE is critical.

Architecture Blueprint

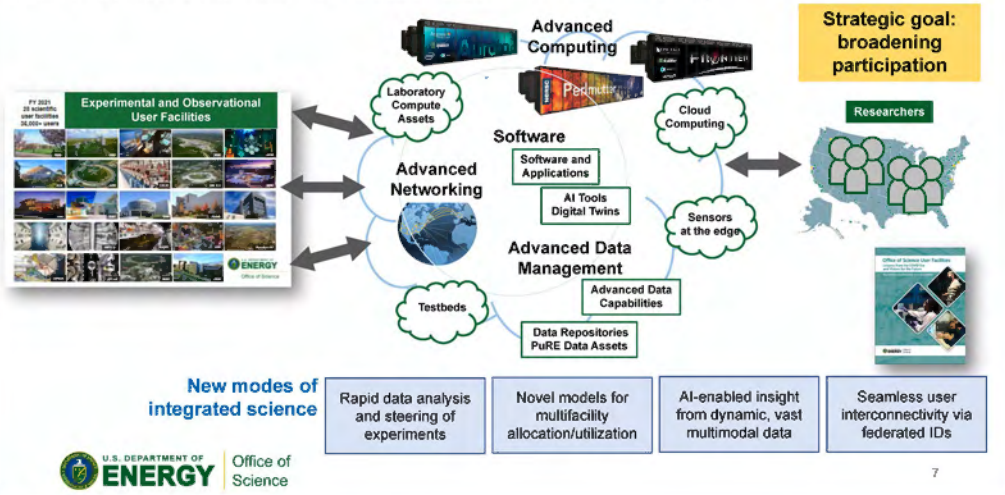
An architecture blueprint is a high-level design concept of the systems, interfaces, and standards required to implement the integration vision. Requirements and principles for operations, governance, allocations, and workforce will also be important.

Given the anticipated breadth of reference priority challenge use cases, it is likely that more than one architecture blueprint will be needed to address the vision comprehensively.



6

The vision: A DOE/SC integrated research ecosystem that transforms science via seamless interoperability



The 2020 SC User Facilities Roundtable Report set the stage

December 2020 Roundtable Report




Future All-Facility Meetings

“Facility directors have expressed an interest in hearing about other facilities’ challenges and strategies. During the SC roundtable, participants noted and appreciated that, for the first time, all SC user facilities were involved in the same meeting. The exchange of information was beneficial and enlightening, allowing facility directors and staff to better understand the entire DOE system, the challenges faced by other facilities, and the strategies employed to meet those challenges.”

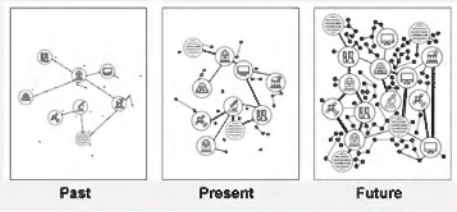
“Data, computing, and networking infrastructure are critical for scientific productivity; they are the substrates the research community uses to explore, create, and share information.”

Interconnectivity and integration of instrumentation, data, and computing have been explicitly recognized as strategic requirements for national R&D

The 2021 National Strategic Overview from the subcommittee on Research and Development Infrastructure formally redefined “federal R&D Infrastructure” to now include computing, data, and networking facilities, resources and services.



“R&D continues to shift from smaller to bigger science, driven in large part by advances in computing and other research cyberinfrastructure, which interlink[s] research data, analytics, ... and experimental instrumentation.”



National imperatives for US leadership in strategic areas all require an interlinked ecosystem of instruments, compute, data



Multi-Agency R&D Priorities for the FY 2023 Budget (August 27, 2021)

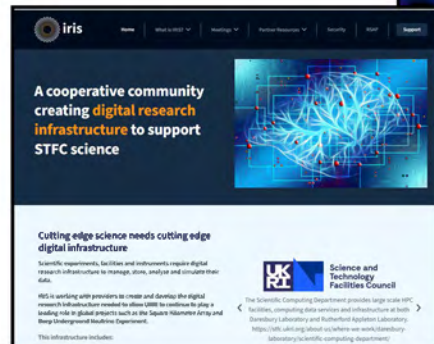
- Pandemic readiness and prevention.
- Tackling climate change.
- Catalyzing research and innovation in critical and emerging technologies.**
 - Artificial intelligence (AI)**
 - Quantum Information Science (QIS)**
 - High-performance computing**
- Advanced communications technologies, microelectronics, biotechnology, robotics, space technologies.
- Innovation for equity.
- National security and economic resilience.

Emerging interagency concepts and initiatives towards a **National Research Ecosystem**

DOE is uniquely placed to lead and shape the national conversation and initiatives.

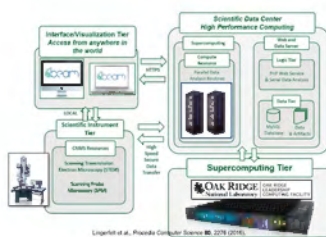
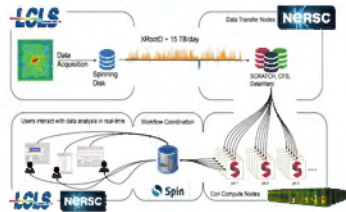
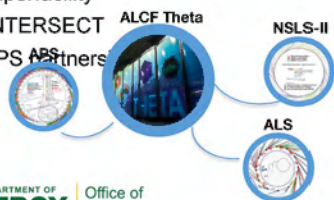
International efforts in integrated research infrastructure are expanding too

- China Science and Technology Cloud
- European Open Science Cloud
- IRIS UKRI SFTC initiative



Across DOE, innovators have been taking concerted steps towards integration through research, partnerships, and lab-level projects

- 2016 Data Pilots
- CAMERA
- BES Light Source Data Working Group DISCUS project
- BER joint EMSL-JGI FICUS joint-allocation program
- ASCR Integrated Research Infrastructure Task Force
- ECP ExaWorks, ExaFEL
- NERSC/LCLS LLANA project
- LBNL Superfacility
- ORNL INTERSECT
- ALCF-APS Partners
- ...



To prepare to engage the community, within ASCR an IRI Task Force contemplated the operational implications of integration concepts



ASCR Heterogeneous Research Infrastructure Task Force March 9, 2011

Toward a Seamless Integration of Computing, Experimental, and Observational Science Facilities: A Blueprint to Accelerate Discovery

About the ASCR Integrated Research Infrastructure Task Force

There is growing, limited recognition that integration of computational, data management, and experimental research infrastructure holds enormous potential to facilitate research and accelerate discovery. The complexity of data-intensive scientific research—whether modeling/simulation or experimental/observational—presents scientific opportunities and research challenges to the research community and large.

Within the Department of Energy's Office of Science (SC), the Office of Advanced Scientific Computing Research (ASCR) will play a major role in defining the SC vision and strategy for integrated computational and data research infrastructure. The ASCR facilities provide essential high end computing, high performance networking, and data management capabilities to address the SC mission and broader Departmental and national research objectives. Today the ASCR Facilities are already working with other SC stakeholders to explore novel approaches to complex, data-intensive research workflows, leveraging ASCR support.

- Corey Adams
 - Katie Antypas
 - Debbie Bard
 - Shane Canon
 - Eli Dart
 - Chin Guok
 - Ezra Kissel
 - Eric Lincorn
 - Bronson Messer
 - Sarp Oral
 - Jimi Ramprakash
 - Ajayn Shankar
 - Tom Uram
- U.S. DEPARTMENT OF ENERGY | Office of Science

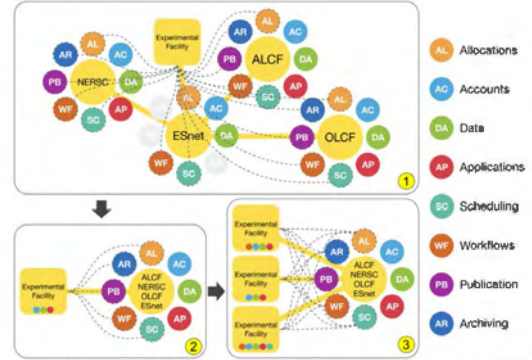


Figure 1. Depiction of the integration of experimental facilities with computational facilities, across the range of services provided, in contrast with the one-to-one approach required today. 1. Today, an experimental facility must arrange separate, bespoke interactions with individual HPC/HPIN facilities. 2. A future paradigm with common interfaces could simplify integration of an experimental facility with multiple HPC/HPIN facilities. 3. In turn, these common interfaces could support expansion and integration across multiple experimental facilities and HPC/HPIN facilities.

Why have an all-of-SC activity? Why now?

Our collective mission is to accelerate discovery and time-to-solution for our researchers and users.

There are a number of inspiring “hero-level” IRI demonstrations and projects.

But today, design of integrated workflows is point-to-point, not scalable, not composable/flexible, challenging for users and facilities, alike ...

Our challenge is to address this complexity with an SC strategy that maximizes the collective impact of DOE investments.

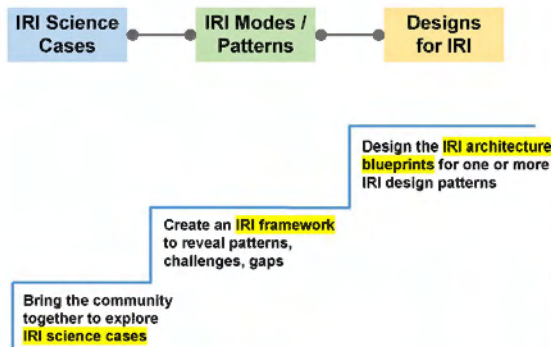
About the IRI-ABA

SC Integrated Research Infrastructure Architecture Blueprint Activity (IRI-ABA)

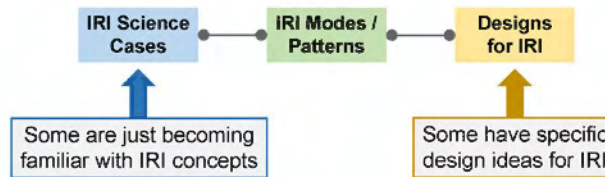
- **Aim:** Produce the **reference conceptual foundations** to inform a coordinated “whole-of-SC” strategy for an integrative research ecosystem.
- **Approach:**
 - Invite DOE experts across the SC user facilities, SC national laboratories, and key SC enterprise stakeholders to participate in a series of activities and events.
 - Gather and analyze **integrative use cases** that inclusively span SC programs and user facilities.
 - Develop overarching design principles and one or more “architecture blueprints” that will address the chief IRI **design patterns** in an efficient way.
 - Identify **urgent program and lab priorities** and early win opportunities.
- **Intended outcomes:**
 - Produce a shared understanding across SC and DOE of IRI requirements, operational and technical gaps and needed investments, and a common lexicon to describe these.
 - Position SC programs to contemplate future investment decisions.
 - Explore leveraging existing SC and ASCR resources and services as well as identifying new needs for research and capability gaps for new resources that do not yet exist.
- **Timeline:** February through September 2022.



The Activity in a Nutshell



Each user facility, stakeholder group, and individual is starting from a unique place vis-à-vis familiarity with IRI science cases, concepts, and design ideas

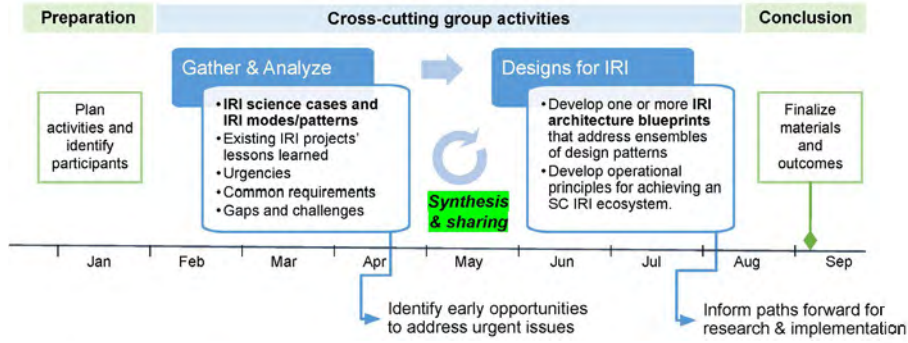


A fundamental goal of the activity is to bring everyone together to create a **shared IRI framework: a conceptual language of IRI science modes and design patterns.**

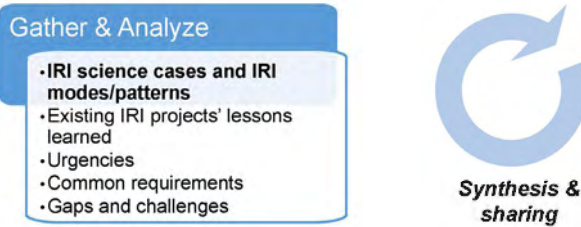
Why do we need this framework? In order to identify common requirements, gaps, and opportunities that will inform strategy.



SC IRI Architecture Blueprint Activity, FY 2022



About the IRI-ABA: Initial Activities



Summary of initial Events and Sprints

Sprint name	Goal	Who will participate?	Mode/method
Gather Insight: IRI Science Cases	Identify/describe representative IRI science cases with an ear to patterns spanning SC science. Focus areas: A. Existing use cases. B. Future/blue sky use cases. C. Urgent priority cases	Small cross-cutting groups, ideally representing at least two SC or DOE science domains Facilitators and note-takers	Facilitated interview & templated insight capture <i>Starting materials:</i> Set of "seed" IRI science cases, template, themes/modes
Gather Insight: Lessons Learned from Existing Projects	Capture and summarize insight, lessons learned, and patterns across relevant existing integration projects and initiatives across SC. Gather notions of IRI designs. Describe each project, design patterns represented, and lessons.	Small cross-cutting groups of leaders and participants of the projects, IRI technologists, others. Facilitators and note-takers	Facilitated interview & templated insight capture <i>Starting materials:</i> Set of "seed" IRI science cases, template, themes/modes.
Gather Insight: Facility Director & Other Stakeholder Conversations	Develop a cross-cutting understanding of IRI drivers, perspectives, opportunities, urgencies, blockers, lessons, and concerns regarding IRI; identify IRI science cases to investigate next.	Facility directors and other key lab management stakeholders. Facilitators and note-takers	Facilitated conversations via ad hoc scheduled meetings

Logistics for participants: What's next

The initial set of gather & analyze sprint activities will take place March 7 to 18.

- We will be reaching out to provide additional information and to confirm your interest to participate.
- Participants should expect a time commitment of ~4 hours (not necessarily contiguous) in this timeframe.

Individual participants have agency and will decide whether to opt into the different "sprint" activities.

Intake survey – please respond if you have not done so, as this is crucial to Slack onboarding.

Slack – you will receive an invitation soon to the IRI-ABA Slack instance!

Shared filespace – no action, but please be aware that there will be a platform for document sharing.

Inquiries – We have established the email address doe-iri-aba@science.doe.gov for inquiries, but Slack will be the preferred platform so that all can see Q&A.



22

What this activity is, and is not

The IRI-ABA is...

- Open** to all parts of SC and DOE.
- Investigative and conceptual** in nature, to create tools for our future selves and the S&T community.
- Human-centered** design thinking, principally informed by the value propositions for scientists.
- Expanding** on the progress to date.
- Cross-cutting**, which is fundamental to the work of seeking and naming IRI patterns.
- Urgent**. Some stakeholders need insight ASAP; 80% quality is good enough to keep moving forward.
- Iterative** and agile to enable rapid learning.
- Novel** in its organization and deliverables, which are likely to be a collection of artifacts, insights, and tools.
- Flexible** to accommodate busy people.



The IRI-ABA is **NOT**...

- Exclusive to SC, to "facilities people," to ASCR.
- Deterministic; it will not recommend "what gets built." It will not result in "final technical designs."
- Idealized design thinking.
- Replacing the progress to date.
- Organized along the usual program/community silos, by intention.
- Perfectionist. The insights and tools we create are to be continuously improved upon in perpetuity.
- Linear and rigid in process.
- A traditional SC workshop resulting in a single final report.
- Strict in its demands, else it would wither.

23

Thank you!

Discussion / Q&A

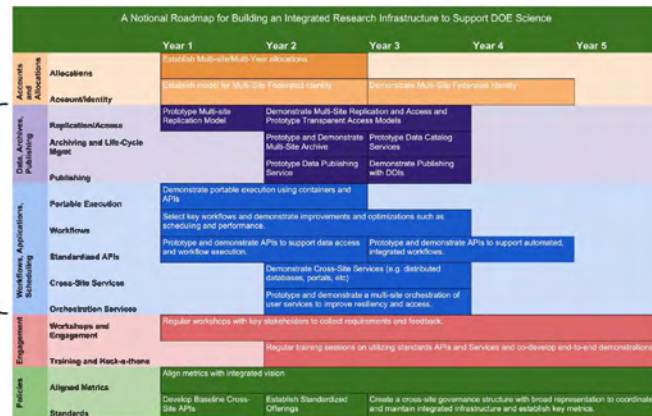


24

Background slides

The ASCR Facilities Task Force white paper offered a notional roadmap for contemplating core design elements of a performant IRI ecosystem

These core technical areas offer a map of the next stage of architecture blueprint development




Guiding principles for a performant ecosystem that allows facilities to build solutions while maintaining their core operational identities

- Flexibility**.....assembly of resource workflows is facile; complexity is concealed
- Performance**.....default behavior is performant, without arcane requirements
- Scalability**.....data capabilities without excessive customizations
- Transparency**.....security, authentication, authorization should support automation
- Interoperability**....services should extend outside the DOE environment
- Resiliency**.....workloads are sustained across planned and unplanned events
- Extensibility**.....designed to adapt and grow to meet unknown future needs
- Engagement**.....promotes co-design, cooperation, partnership
- Cybersecurity**.....security for facilities and users is essential

The roundtable report contained key insights regarding the need for distributed and interoperable computing resources

“Facilities may be able to deploy a distributed network of connected and interoperable computing resources that enable all scales of computing, data exploration, and analysis.”



- The pandemic highlighted the importance of **secure remote collaboration and facilitated access to data and computing resources**.
- **Data management and data stewardship** present another critical opportunity. User facilities generate exabytes of unique, irreplaceable data which must be managed, curated, and made available for analysis and computation.
- This requires a host of **user-connecting operational approaches and technologies** such as high-quality interfaces, collaboration tools, federated identity management, automation of experiments and workflows, and more.
- “With collaboration among all its user facilities, DOE SC is in a position to facilitate all aspects of the data lifecycle across its facility complex, including simulations, experiment design, data generated at scientific instruments, data analysis, and data archiving for future use.”
- “Seamlessly connecting a user with data and computing enables more uniform and egalitarian data exploration and analysis capabilities.”

Today the ASCR facilities enterprise is contending with new complexity. We are entering a new era of advanced scientific computing.

The practice of science is evolving. Couplings between modeling/simulation, experimental/observational data, advanced algorithms, and AI/ML tools have the power to accelerate discovery and innovation.

Where we once focused on batch jobs and bulk data transfer, we now have **complex workflows**.

Computing technology is evolving along multiple trajectories. General purpose computing is but one market segment. **Managing risk and opportunity in our hardware choices is increasingly complex.**

The **people** of the ASCR facilities enterprise are making extraordinary impacts today; their expertise and efforts are sought by many. And yet many talented individuals do not participate.

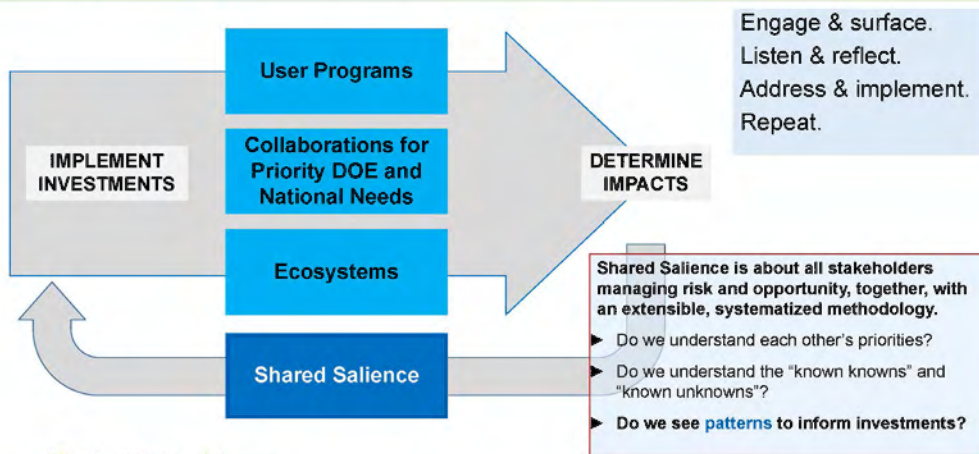
Our workforce challenges are significant.

Institutions, programs, and researchers are under pressure to provide/obtain computing and data resources.

Our users, our partners, and we ourselves crave shared clarity of insight and intent.

Our challenge today is to confront this complexity and arrive at a strategy that maximizes the impact of ASCR, Office of Science, and DOE investments—to be greater than the sum of the parts.

Vision for the ASCR Facilities: How we will thrive together



Appendix C

SYNTHESIS OF SPRINT 1 INSIGHTS

IRI Sprint 1 Synthesis Summary

For validation and feedback

5/11/22

Summary of Sprint 1: Gathering insight from users on IRI science cases and associated challenges

- IRI ABA participants (~110) were organized into 10 cross-cutting facilitated groups each including scientists from at least 2 SC program domains, and infrastructure developers and operators and others from different DOE laboratories and facilities.
- Between January and March 2022, the groups met, invited, and conducted several ~1-hour standardized interviews of DOE National Lab scientist and engineers:
 - The goal of each interview: Gain understanding of the interviewee's science goals, workflow and/or computing/data environment; the journey needed to achieve the goals (how do they get it done?) and associated challenges/gaps/blockers to getting their work done.
 - Each group conducted a synthesis exercise to surface key insights across their interviews.
- The Leadership Group assembled and consolidated/summarized the ensemble data and syntheses from the groups to capture a proposed set of overall outcomes from Sprint 1.
 - Data from 30 interviews (30 hours of zoom recordings and transcripts, ~60 sets of interviewer raw observation notes and spreadsheets), and 10 sets of synthesis slides.
 - The LG additionally included a look at patterns across the cases and performed some mapping to considerations in the ASCR IRI Taskforce Report.

Sprint 1 Gathering Insight was focused on user voices

- Sprint 1 was designed as a user-centered/user activity. Our primary intentions were to convene cross-cutting groups and create a methodology for listening to "science case voices" related to the challenges of integrated research infrastructure.
- The interviews produced many cross-cutting perspectives on IRI user journeys such as (a very small sample from many hours of recording...):
 - "... most teams turn around projects [results] in a couple of years.."
 - "... we have a pretty complex workflow in terms of what depends on what.."
 - "... I/O being slow doesn't count as a formal outage and so sometimes you know it's like the system is up, what's the problem?"
 - "... [our] data is highly heterogeneous.."
- Cautionary note: Sprint 1 was not built to be comprehensive. While the 30 interviews span a broad representative spectrum of SC IRI science cases, we acknowledge there are many other IRI cases!

➤ *We invite reflections and feedback on this summary of Sprint 1 (directly to your facilitator, or slack channels) and in responses to future activities as we move forward.*

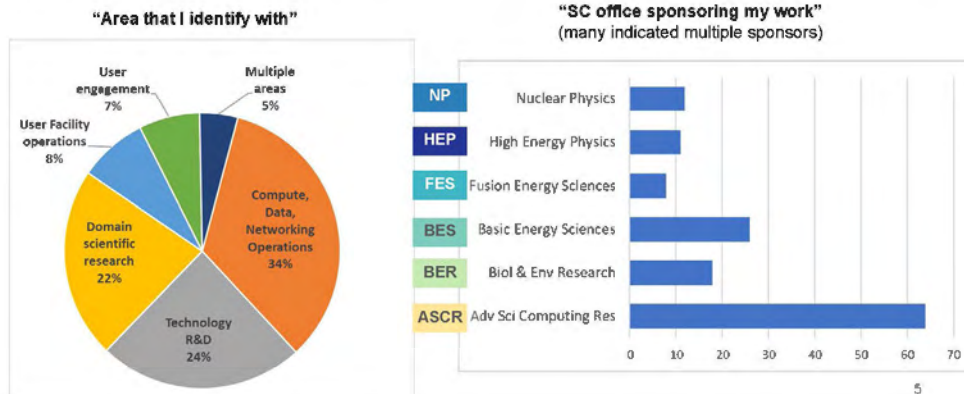
We employed different approaches for the synthesis

- I. User sentiments expressed as phrase/term fragments extracted from facilitator-driven synthesis for each group
Identified commonalities and recurring user sentiments; tweaked tense/words for clarity
- I. Interpreting requirements from each group's synthesis
Grouped by requirements; linked with Integrated Research Infrastructure Task Force Report themes

Snapshot of IRI ABA Participants (as of April 2022)

- ✓ 112 subject matter experts with diverse backgrounds from across the SC lab complex. (Does not include additional 30 science user interviewees who provided input during Sprint 1. The next two slides show that the interviewees were speaking to science use cases that used multiple facilities.)

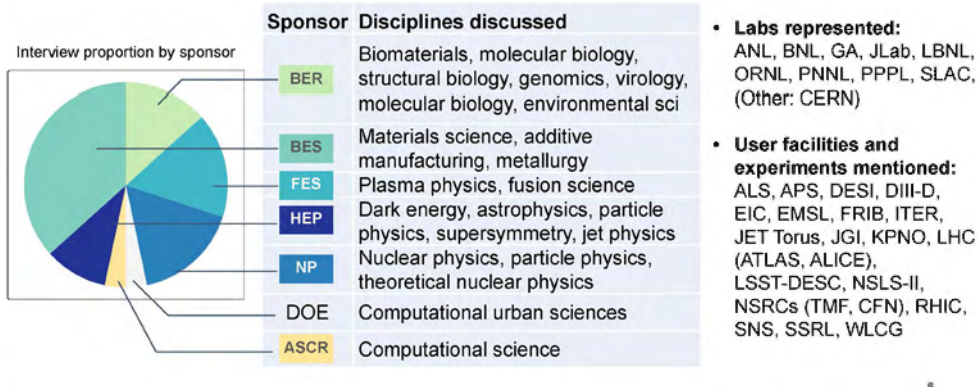
Interest survey responses



Sprint 1 Interviews Coverage

The 10 Sprint 1 groups conducted 30 single-person interviews with DOE science users using a standard template with 3 main areas: your IRI science case(s), your IRI journey, and challenges and opportunities.

The IRI ABA participants (shown in the previous slide) were the listeners for these interviews.



Sprint 1 interviews touched on many IRI cases

NP	Large-scale simulations integrating experimental data from DOE and international nuclear physics facilities	FES	DIII-D Fusion experiments + simulations, near-real-time analysis to determine/predict plasma conditions
BES	Autonomous AI/ML-driven experiments such as for ALS and NSLS-II	HEP	LHC simulation workflows, event generation for integration of VLCCG with HPC
NP	AI/ML incorporated into simulations to drive exploration of parameter space using codes like LAMMPS in the workflow	NP	GRETA spectrometer online computing pipeline construction - for 1 to 5 day campaigns
BER	Applying AI/ML to EMSL data for analysis	FES	DIII-D Tokamak plasma physics experiments - diagnostic data collection. Integrating ASCR HPC facilities with the VLCCG, for LHC computing.
BER	Combining multimodal data from simulations and experiments for molecular scale imaging workflows (instrument + compute + storage)	HEP	Event generation for complex simulation pipeline. Distributed, compute-intensive but not data-intensive.
HEP	Astronomical spectroscopic observations at DESI/KPNO, survey data daily streaming to HPC, target selection pipeline	BES	APS 8-ID-I Small-angle XPCS x-ray spectroscopy, high-frame rate camera, data management workflow including HPC. 0.2 PB of unsparisified data generated/day.
FES	Heterogeneous data handling and real-time analysis for fusion power experiment steering at DIII-D and JET for fusion	BES	Metallurgy, MIDAS x-ray analysis software for APS high-energy X-ray diffraction microscopy beamline data. High data throughput. Pipeline development to NERSC and ALCF for on-demand computing at scale.
HEP	Astronomy data management - DESI workflows - processing 100s of TBs of data utilizing HPC	BES	Materials science. High data throughput, facility data transfer from APS to PNNL for analysis.
BER	Crystallography microscopy and synchrotron light sources workflows and local and HPC computing	BES	Light source data processing workflows for large datasets (TBs to PBs), particularly ptychography, using AI/ML to help reduce data quickly.
FES	Fusion workflows: 10s -45GB per pulse, about 500TB total using custom SQL and NoSQL databases, computing with local and institutional resources, data exposed via APIs also for future ITER workflows	FES	Data management and real-time analysis for fusion experiments at PPPL.
BER	HPC-enabled high-throughput sequencing, large-scale sequence data analysis, sample and data life cycle, data product development (e.g., at the JGI)	BES	Data pipelines for high-speed detectors used at lights sources and NSRCs. Workflow development using NERSC
BES	Beamline data transfer at scale to HPC for real-time analysis and experiment steering, data transfer at scale to HPC (e.g., at SSRLL)	BES	ML Autonomous materials characterization workflows using data (100s GBs per day) collected at light sources, neutron sources, and NSRCs.

I. Common recurring sentiments across the user interviews

Data Management	<ul style="list-style-type: none"> Users are overwhelmed with large and growing amounts of data to manage, reduce, analyze. <ul style="list-style-type: none"> Users need to move data across facilities and use different systems at different steps of data processing chain. Users need bespoke data movement and workflow solutions, and long-duration support for data/metadata.
Automation/AI	<ul style="list-style-type: none"> Users need reliable automation & seamless access, and try to compensate via human effort. <ul style="list-style-type: none"> Users need automation, and anticipate AI, but struggle with skills and application of these novel technologies.
Heterogeneity	<ul style="list-style-type: none"> Users face mismatches between resources, tools, and needs. <ul style="list-style-type: none"> Users need heterogeneity in scale and type of resources but have platform fatigue learning many different platforms. Users need workflows to be at the center but need software APIs and standardization/uniformity. Users have a spectrum of computing needs from elastic computing (matching need to available resources) to urgent computing (near real-time/just-in-time, on-demand).
Ease of Use	<ul style="list-style-type: none"> Users find infrastructure hard to use. <ul style="list-style-type: none"> Users encounter a lack of transparency about workflow tools and resources, and many different use policies and cybersecurity barriers. Users need infrastructure to be easier to use and be more coordinated across resources and facilities.
Workforce Skills Gap	<ul style="list-style-type: none"> Users and teams struggle with workforce and training needs. <ul style="list-style-type: none"> Users (and their organizations) struggle with lack of skills, oversubscribed staff, recruiting, and retention. Users experience gaps between their working knowledge and skills and those of infrastructure experts. Users need support and expertise in data science.

User sentiments as recurring areas of interest

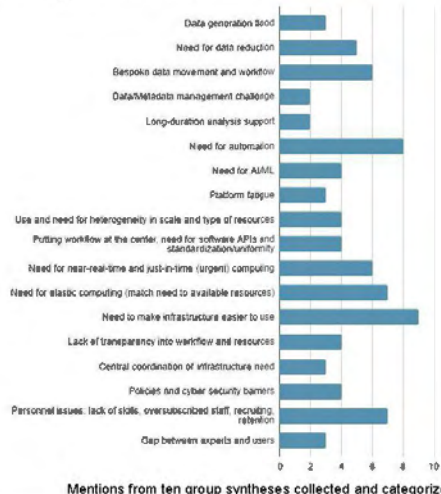
Data Management Considerations

Automation/AI

Heterogeneity

Ease of Use/ Empowering the User

Workforce Skills Gap



II. Interpreting “requirements” from group synthesis

- Need for scaling of long-term storage, compute, and associated networking
- Need for some level of workflow/software standardization (including APIs to access resources, data formats, etc.) and associated support model
- No “one size fits all” solution, need to support different compute models, storage solutions, allocation types, etc.
- Need for automation, but also AI/ML to help provide intelligent automation (and not just automating manual/mundane things)
- Need for better personnel resource management, training, career development, recruiting, etc.
- Need for different “classes” of policies and allocation
- Need for better security models to balance protection of data and ease of access
- Need for common/compatible AuthN/AuthZ models
- Need for better metadata and data management solutions
- Need for widely distributed (heterogeneous) compute and storage solutions, and networks to connect them
- Need resources to be available, reliable, and performant
- Need insights into how workflows/resources are performing, for troubleshooting or performance tuning
- Need for appropriate engagement and support to bridge gap between scientist, and tools developers and infrastructure operations
- Need for better data reduction techniques, potentially leveraging AI/ML

Requirements crosswalk with IRI Task Force WP areas

Allocations

- Need for scaling of long-term storage, compute, and associated networking
- No “one size fits all” solution, need to support different compute models, storage solutions, allocation types, etc.
- Need for different “classes” of policies and allocation

Accounts/Access

- Need for common/compatible AuthN/AuthZ models
- Need for better security models to balance protection of data and ease of access

Data/Archives/Publishing

- Need for scaling of long-term storage, compute, and associated networking
- Need for better metadata and data management solutions
- Need for better data reduction techniques, potentially leveraging AI/ML

Policies and Governance

- Need for different “classes” of policies and allocation

Applications/Scheduling/Workflow

- Need for scaling of long-term storage, compute, and associated networking
- No “one size fits all” solution, need to support different compute models, storage solutions, allocation types, etc.
- Need for some level of workflow/software standardization (including APIs to access resources, data formats, etc.) and associated support model
- Need for automation, but also AI/ML to help provide intelligent automation (and not just automating manual/mundane things)
- Need for widely distributed (heterogeneous) compute and storage solutions, and networks to connect them
- Need resources to be available, reliable, and performant
- Need insights into how workflows/resources are performing, for troubleshooting or performance tuning

Engagement and Partnerships

- Need for better personnel resource management, training, career development, recruiting, etc.
- Need for appropriate engagement and support to bridge the gap between scientists and tools developers and infrastructure operations

Summary and Next Steps

- Sprint 1’s user-centered interviews have been completed and we have collected a wealth of detailed information.
- Initial synthesis identifies and calls out specific themes and needs by the science facility-user community.
- **Please send feedback and/or questions on these slides to your group facilitator.**
- Sprint 2 begins soon and explores technology and operations, especially in the context of IRI-aligned projects and activities already underway or completed by our participants.
- These two sprints set the stage for the architecture blueprint development activities later in the year.

Appendix D

IRI ABA PROJECT SURVEY

Technical Domain	Project name	Reference
Software/applications	ECP	www.exascaleproject.org
https://exaworks.org	APS/ALCF Internal Project on on-demand computing	N/A
Workflows	ExaWorks (ECP)	https://exaworks.org
Data transfer and management	Globus	globus.org
Computing, grid computing job management	Distributed Resource Management Application API (DRMAA)	https://en.wikipedia.org/wiki/DRMAA
Computing, grid computing middleware	UNiform Interface to COmputing Resources (UNICORE)	https://en.wikipedia.org/wiki/UNICORE
Computing, grid computing architecture	Legion	http://www.anandnatrajan.com/papers/IBMJRD03.pdf
Workflows	Workflows Community Initiative	https://workflows.community
Workflows	ExaWorks	https://exaworks.org
Workflows	WorkflowsRI	https://workflowsri.org
Real-time analysis of facility experimental data (fusion)	Automatic Between-Pulse Analysis at ALCF to support DIII-D Operations	https://doi.org/10.1080/15361055.2017.1390388
Software/applications	Large-Eddy Simulation Atmospheric Radiation Measurement Symbiotic Simulation and Observation (LASSO) Activity	https://doi.org/10.1175/BAMS-D-19-0065.1
Integration framework	LBNL Superfacility Project	https://www.nersc.gov/research-and-development/superfacility
Workflows	Globus Architecture for Data-Intensive Experimental Research (Glacier)	glacier.readthedocs.io
Networking, requirements gathering	ASCR ESnet Requirements Reviews	https://www.es.net/science-engagement/science-requirements-reviews/requirements-review-reports/#
Networking, wireless	ESnet Wireless Edge	https://docs.google.com/presentation/d/12G037DMotqon2wLJ9HdOF74sObLvqpCkXSRYcwcWV7M/edit?usp=sharing
Integration framework	MLExchange: Bringing AI to Beamlines	https://www.osti.gov/biblio/1812187
Science case	Self-Driving Field Laboratories	https://acess.onlinelibrary.wiley.com/doi/full/10.2136/vzj2018.03.0061
Data transfer and management	Globus	globus.org
Workflows	Globus Automate	https://docs.globus.org/globus-automation-services
Workflows	Braid	https://anl-braid.github.io/braid

Technical Domain	Project name	Reference
Integration framework	Gladier	https://github.com/globus-gladier/gladier
Software/applications	funcX	funcx.org
Software/applications	Parsl	parsl-project.org
Science case	Pacific Northwest Cryo-EM Center	https://pncc.labworks.org
Identity management	ASCR DCDE	https://science.osti.gov/-/media/ascr/ascac/pdf/meetings/202001/Federating_DOE-SC_Facilities-ASCAC202001...
Science case	Belle II computing 2011-2018	https://www.belle2.org
Computing	EMSL-JGI computing resilience collaboration	N/A
Science case	COMPASS	https://ess.science.energy.gov/compass-coastal-systems-pilot-project
Workflows	ScienceCapsule	https://escholarship.org/uc/item/2709n3mt
Workflows	SCIRA	https://www.osti.gov/biblio/1772907-towards-interactive-reproducible-analytics-scale-hpc-systems
Workflows	ScienceSearch	https://ieeexplore.ieee.org/document/8588644
Science case	AI4ESP	https://www.ai4esp.org/files/AI4ESP1136_Varadharajan_Charuleka2.pdf
Science case	“NESAP for Learning (N4L)” project on Perlmutter	N/A
Science case	Integrated Simulation of Energetic Particles in Burning Plasmas (ISEP) Project	https://www.scidac.org/projects/2018/fusion-energy-sciences/isep.html
Science case	2021-2022 ALCC “AI/Deep Learning Prediction & Real-Time Control for Fusion Energy Systems”	N/A
Science case	2022 INFUSE “Improving Plasma Control Capabilities in Magnetically-Confining Tokamak Systems with Transformer Neural Networks”	N/A
Science case	2022 SUMMIT INCITE on “Exascale Simulation and Deep Learning Model for Energetic Particles in Burning Plasmas”	N/A
Integration framework	Project INTERSECT	https://www.ornl.gov/intersect
Science case	Advanced Plant Phenotyping Laboratory (APPL) project	https://www.ornl.gov/content/advanced-plant-phenotyping-laboratory-appl
Data analysis	Project ICEMAN	https://sns.gov/content/iceman-a-heterogeneous-platform-analysis-neutron-scattering-data
Workflows	Balsam	https://balsam.readthedocs.io

Technical Domain	Project name	Reference
Networking	SENSE	https://arxiv.org/abs/2004.05953
Workflows	JAWS	https://jaws-docs.readthedocs.io/en/latest/Intro/how_jaws.html
Workflows	ALCC "Towards Resilient and Portable Workflows across DOE's Facilities"	https://crossfacilityworkflows.github.io/BestPractices/index.html
Data analysis	LCLS-II data analysis	N/A
Science case	Gamma-Ray Energy Tracking Array (GRETA)	https://greta.lbl.gov/
Science case	DUNE experiment	https://lbnf-dune.fnal.gov
Science case	ATLAS experiment	https://atlas.cern
Data transfer and management	SciStream	https://scistream.github.io
Data transfer and management	AI-Steer	N/A
Workflows, performance monitoring	RAMSES	https://ramsesproject.github.io
Computing, distributed	JLab Scientific Computing Environment	http://scicomp.jlab.org/scicomp/
Networking	EJFAT Data Steering Project	https://wiki.jlab.org/epsciwiki/index.php/EJFAT
Workflows	Environment for Realtime Streaming Applications (ERSAP)	https://wiki.jlab.org/epsciwiki/index.php/ERSAP
Software/applications	AToM	https://atom.scidac.io
Software/applications	OMFIT	https://omfit.io
Data transfer and management	A Framework for International Collaboration on ITER Using Large-Scale Data Transfer to Enable Near-Real-Time Analysis	https://doi.org/10.1080/15361055.2020.1851073
Science case	DUNE use of Google GPUs	https://doi.org/10.3389/fdata.2020.604083
Science case	LHC use of ASCR HPC	N/A
Science case	NERSC processing of SARS-CoV-2 data acquired at LCLS	https://www.nersc.gov/science/covid-19-research/c3-ai-digital-transformation-institute/nersc-and-lcls-team-up-on-sars-cov-2-research-article-page/
Data transfer and management	APS Beamline Data Pipeline Project	N/A
Data transfer and management	APS Data Management System	N/A
Science case	ATLAS/CMS Experiments at LHC	home.cern

Technical Domain	Project name	Reference
Networking, traffic load balancing	EJ-FAT (ESnet/JLAB - FPGA Accelerated Transport)	N/A
Networking, traffic load balancing	HECATE	N/A
Networking	ESnet 5G pilot program	N/A
Data management	Rucio	https://rucio.cern.ch
Data transfer and management	FTS	https://wlcg-ops.web.cern.ch/fts
Data transfer and management	Xrootd	https://xrootd.slac.stanford.edu
Software/Applications	CAMERA	camera.lbl.gov
Workflows	MLExchange	mlexchange.lbl.gov
Discovery platform	National Microbiome Data Collaborative (NMDC)	microbiomedata.org
Discovery platform	KBase	kbase.us

Appendix E

THE INTEGRATED RESEARCH INFRASTRUCTURE ARCHITECTURE BLUEPRINT ACTIVITY (IRI ABA) LAUNCHING THE DESIGN PHASE

The Integrated Research Infrastructure Architecture Blueprint Activity (IRI-ABA)

Launching the Design Phase

June 15, 2022

Outline

<p>IRI ABA Overall Goals and Plan Outcomes of Gathering Insight Phase Plan for the Design Phase</p>	}	30 minutes
<p>Discussion / Q&A</p>		30 minutes

SC Integrated Research Infrastructure Architecture Blueprint Activity Leadership

HQ Executive Leadership



Ben Brown
Director
ASCR Facilities Division



Bill Miller
Senior Technical Advisor
ASCR Facilities Division

IRI-ABA Leadership Group



Debbie Bard
Group Lead for Data
Science Engagement
NERSC, LBNL



Eric Lancon
Director, Scientific Data
and Computing Center
BNL



Amber Boehlein
Chief Information Officer
JLab



Jini Ramprakash
Deputy Division Director
ALCF, ANL



Kjersten Fagnan
Chief Informatics Officer
JGI, LBNL



Arjun Shankar
Section Head,
Advanced Technologies
OLCF/NCCS, ORNL



Chin Guok
Group Lead for Planning
and Architecture
ESnet, LBNL



Nicholas Schwarz
Group Leader,
Scientific Software
Eng. & Data Mgmt.,
APS, ANL

IRI ABA SC Headquarters Coordination Group

BER	Paul Bayer, Jay Hnilo, Resham Kulkarni
BES	Tom Russell
FES	Josh King, Matt Lanctot
HEP	Jeremy Love, Eric Church
IP	Kristian Myhre
NP	Xiaofeng Guo, Jim Sowinski

- Exchanges on urgent IRI needs, priorities, & commonalities across programs.
- Input on IRI science patterns and representative use cases to address.
- Engages and provides feedback at key points on Activity progress and outputs.



4

Code of conduct

The DOE Office of Science (SC) is fully and unconditionally committed to fostering safe, diverse, equitable, and inclusive work, research, and funding environments that value mutual respect and personal integrity.

If anyone involved in this event has questions or concerns, please approach the HQ Executive Sponsors.

more at >> <https://science.osti.gov/sc-2/Research-and-Conduct-Policies/Diversity-Equity-and-Inclusion/SC-Statement-of-Commitment>



5

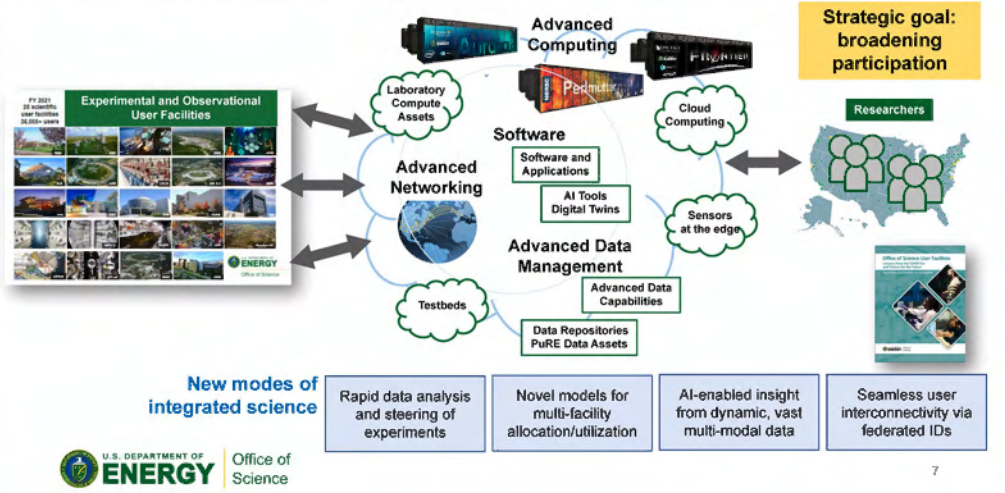
Reminder about the IRI ABA goals and plan



Architecture Blueprint Activity

6

The vision: A DOE/SC integrated research ecosystem that transforms science via seamless interoperability

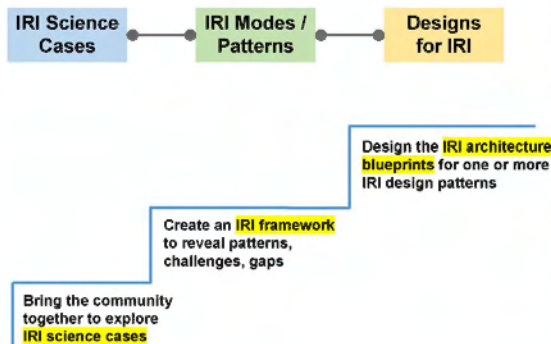


SC Integrated Research Infrastructure Architecture Blueprint Activity (IRI-ABA)

- **Aim:** Produce the **reference conceptual foundations** to inform a coordinated “whole-of-SC” strategy for an integrative research ecosystem.
- **Approach:**
 - Invite DOE experts across the SC user facilities, SC national laboratories, and key SC enterprise stakeholders to participate in a series of activities and events.
 - Gather and analyze **integrative use cases** that inclusively span SC programs and user facilities.
 - Develop overarching design principles and one or more “architecture blueprints” that will address the chief IRI **design patterns** in an efficient way.
 - Identify **urgent program and lab priorities** and early win opportunities.
- **Intended outcomes:**
 - Produce a shared understanding across SC and DOE of IRI requirements, operational and technical gaps and needed investments, and a common lexicon to describe these.
 - Position SC programs to contemplate future investment decisions.
 - Explore leveraging existing SC and ASCR resources and services as well as identifying new needs for research and capability gaps for new resources that do not yet exist.
- **Timeline:** February through September 2022.

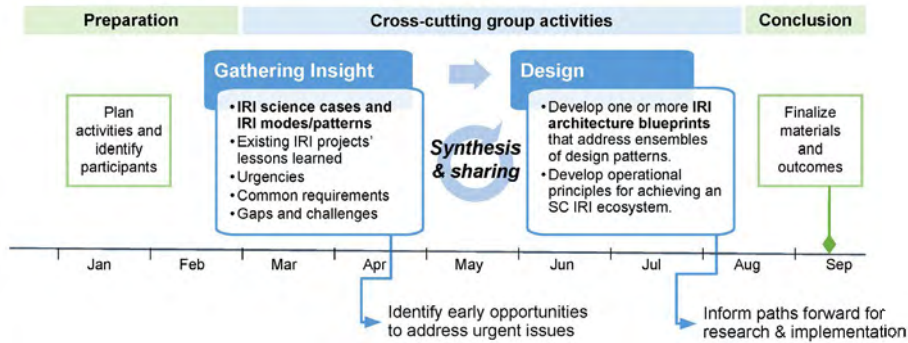


The Activity in a Nutshell



SC IRI Architecture Blueprint Activity, FY 2022

Goal: Produce the reference conceptual foundations to inform a coordinated “whole-of-SC” strategy for an integrative research ecosystem.



What this Activity is, and is not

The IRI-ABA is...

- Open** to all parts of SC and DOE.
- Investigative and conceptual** in nature, to create tools for our future selves and the S&T community.
- Human-centered** design thinking, principally informed by the value propositions for scientists.
- Expanding** on the progress to date.
- Cross-cutting**, which is fundamental to the work of seeking and naming IRI patterns.
- Urgent.** Some stakeholders need insight ASAP; 80% quality is good enough to keep moving forward.
- Iterative** and agile to enable rapid learning.
- Novel** in its organization and deliverables, which are likely to be a collection of artifacts, insights, & tools.
- Flexible** to accommodate busy people.

The IRI-ABA is **NOT**...

- Exclusive to SC, to “facilities people,” to ASCR.
- Deterministic; it will not recommend “what gets built.” It will not result in “final technical designs.”
- Idealized design thinking.
- Replacing the progress to date.
- Organized along the usual program/community silos, by intention.
- Perfectionist. The insights and tools we create are to be continuously improved upon in perpetuity.
- Linear and rigid in process.
- A traditional SC workshop resulting in a single final report.
- Strict in its demands, else it would wither.

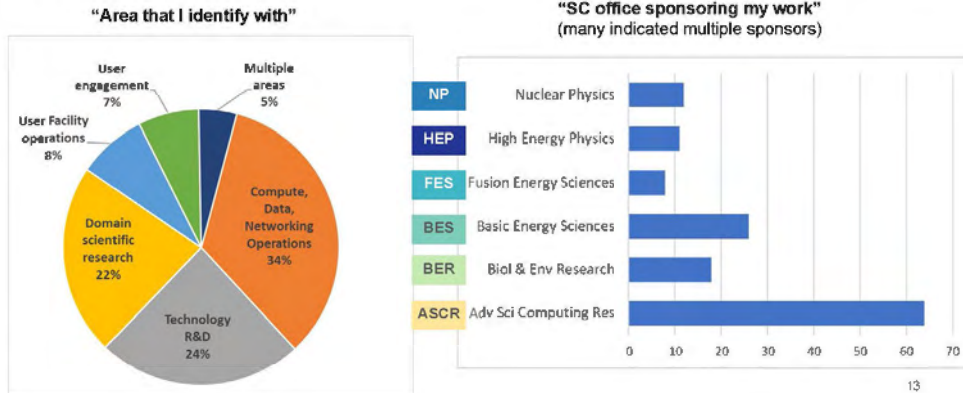
The Gathering Insight Phase: Accomplishments and Outcomes



Snapshot of IRI ABA Participants (as of May 2022)

- ✓ 112 subject matter experts with diverse backgrounds from across the SC lab complex. (Does not include additional 30 science user interviewees who provided input during Sprint 1. The next two slides show that the interviewees were speaking to science use cases that used multiple facilities.)

Interest survey responses



Summary of Gathering Insight Phase Sprints

Activity	Goal	Approach
✓ Sprint 1: IRI Users and Science Cases	Describe representative IRI science cases and understand the user experience and journey. Listen for common IRI patterns and challenges across SC domains.	<ul style="list-style-type: none"> • Small cross-cutting groups, representing several SC or DOE science domains and IRI technology areas • Structured interviews of scientists and other users • Within-group and overall synthesis of key learnings*
✓ Sprint 2: Lessons Learned from Existing Projects	Capture insights, lessons, and patterns from relevant existing SC integration projects and initiatives.	<ul style="list-style-type: none"> • Same cross-cutting groups • Unstructured discussions of focus questions • Survey on projects aligned with IRI • Synthesis of patterns and practice areas informing design phase.



* Feel free to continue to provide feedback on the Synthesis: *[Sprint1_Synthesis_For_Feedback.05-11-22.pdf](#)

14

Our thanks to the Sprint 1 and Sprint 2 facilitators!

Amber Boehnlein	Jini Ramprakash
Andrew Wiedlea	Kjiersten Fagnan
Arjun Shankar	Nicholas Schwarz
Bronson Messer	Sarp Oral
Chin Guok	Shane Canon
Debbie Bard	Tom Uram
Eric Lancon	



15

Sprint 1 interviews touched on many IRI cases

NP	Large-scale simulations integrating experimental data from DOE and international nuclear physics facilities	FES	DIII-D Fusion experiments + simulations, near-real-time analysis to determine/predict plasma conditions
BES	Autonomous AI/ML-driven experiments such as for ALS and NSLS-II	HEP	LHC simulation workflows; event generation for Integration of WLCG with HPC
NP	AI/ML incorporated into simulations to drive exploration of parameter space using codes like LAMMPS in the workflow	NP	GRETA spectrometer online computing pipeline construction - for 1 to 5 day campaigns
BER	Applying AI/ML to EMSL data for analysis	FES	DIII-D Tokamak plasma physics experiments - diagnostic data collection. Integrating ASRC HPC facilities with the WLCG, for LHC computing.
BER	Combining multimodal data from simulations and experiments for molecular scale imaging workflows (instrument + compute + storage)	HEP	Event generation for complex simulation pipeline. Distributed, compute-intensive but not data-intensive.
HEP	Astronomical spectroscopic observations at DESI/KPNO, survey data daily streaming to HPC, target selection pipeline	BES	APS 8-ID-I Small-angle XPCS x-ray spectroscopy, high-frame rate camera, data management workflow including HPC. 0.2 PB of unsparsified data generated/day
FES	Heterogeneous data handling and real-time analysis for fusion power experiment steering at DIII-D and JET for fusion	BES	Metallurgy, MIDAS x-ray analysis software for APS high-energy X-ray diffraction microscopy beamline data. High data throughput. Pipeline development to NERSC and ALCF for on-demand computing at scale
HEP	Astronomy data management - DESI workflows - processing 100s of TBs of data utilizing HPC	BES	Materials science. High data throughput, facility data transfer from APS to PNNL for analysis
BER	Crystallography microscopy and synchrotron light sources workflows and local and HPC computing	BES	Light source data processing workflows for large datasets (TBs to PBs), particularly ptychography; using AI/ML to help reduce data quickly
FES	Fusion workflows: 10s ~45GB per pulse, about 500TB total using custom SQL and NoSQL databases, computing with local and institutional resources, data exposed via APIs also for future ITER workflows	FES	Data management and real-time analysis for fusion experiments at PPPL
BER	HPC-enabled high-throughput sequencing, large-scale sequence data analysis, sample and data life cycle, data product development (e.g., at the JGI)	BES	Data pipelines for high-speed detectors used at light sources and NSRCs. Workflow development using NERSC.
BES	Beamline data transfer at scale to HPC for real-time analysis and experiment steering, data transfer at scale to HPC (e.g., at SSRL)	BES	ML Autonomous materials characterization workflows using data (100s GBs per day) collected at light sources, neutron sources, and NSRCs

Common recurring sentiments across the user interviews

Data Management	<ul style="list-style-type: none"> Users are overwhelmed with large and growing amounts of data to manage, reduce, analyze. <ul style="list-style-type: none"> Users need to move data across facilities and use different systems at different steps of the data processing chain. Users need bespoke data movement and workflow solutions, and long-duration support for data/metadata.
Automation/AI	<ul style="list-style-type: none"> Users need reliable automation & seamless access and try to compensate via human effort. <ul style="list-style-type: none"> Users need automation, and anticipate AI, but struggle with skills and application of these novel technologies.
Heterogeneity	<ul style="list-style-type: none"> Users face mismatches between resources, tools, and needs. <ul style="list-style-type: none"> Users need heterogeneity in scale and type of resources but have platform fatigue learning many different platforms. Users need workflows to be at the center but need software APIs and standardization/uniformity. Users have a spectrum of computing needs from elastic computing (matching need to available resources) to urgent computing (near real-time/just-in-time, on-demand).
Ease of Use	<ul style="list-style-type: none"> Users find infrastructure hard to use. <ul style="list-style-type: none"> Users encounter a lack of transparency about workflow tools and resources, and many different use policies and cybersecurity barriers. Users need infrastructure to be easier to use and be more coordinated across resources and facilities.
Workforce Skills Gap	<ul style="list-style-type: none"> Users and teams struggle with workforce and training needs. <ul style="list-style-type: none"> Users (and their organizations) struggle with lack of skills, oversubscribed staff, recruiting, and retention. Users experience gaps between their working knowledge and skills and those of infrastructure experts. Users need support and expertise in data science.

Requirements by IRI Task Force white paper areas

Allocations <ul style="list-style-type: none"> Need for scaling of long-term storage, compute, and associated networking No "one size fits all" solution, need to support different compute models, storage solutions, allocation types, etc. Need for different "classes" of policies and allocation 	Applications/Scheduling/Workflow <ul style="list-style-type: none"> Need for scaling of long-term storage, compute, and associated networking No "one size fits all" solution, need to support different compute models, storage solutions, allocation types, etc. Need for some level of workflow/software standardization (including APIs to access resources, data formats, etc.) and associated support model Need for automation, but also AI/ML to help provide intelligent automation (and not just automating manual/mundane things) Need for widely distributed (heterogeneous) compute and storage solutions, and networks to connect them Need resources to be available, reliable, and performant Need insights into how workflows/resources are performing, for troubleshooting or performance tuning
Accounts/Access <ul style="list-style-type: none"> Need for common/compatible AuthN/AuthZ models Need for better security models to balance protection of data and ease of access 	Engagement and Partnerships <ul style="list-style-type: none"> Need for better personnel resource management, training, career development, recruiting, etc. Need for appropriate engagement and support to bridge the gap between scientists and tools developers and infrastructure operations
Data/Archives/Publishing <ul style="list-style-type: none"> Need for scaling of long-term storage, compute, and associated networking Need for better metadata and data management solutions Need for better data reduction techniques, potentially leveraging AI/ML 	
Policies and Governance <ul style="list-style-type: none"> Need for different "classes" of policies and allocation 	

Summary of Sprint 2

Goals: Gather insights, lessons, patterns from existing IRI-relevant SC integration projects and initiatives. Incorporate reflection on user challenges from Sprint 1.

Participant survey on IRI-aligned projects

↳ 75 unique entries in these categories:

- Science cases
- Computing (distributed/grid, on demand, architecture, job management, middleware)
- Data (real-time analysis, analysis, management)
- Discovery platforms, integration frameworks
- Identity management
- Networking (requirements gathering, traffic load balancing, wireless, APIs)
- Workflows, performance monitoring
- Software and applications

Group discussion focus questions:

- What scope/themes relevant for an IRI have been addressed by existing projects?
- Where do we have collective experience and where do we not?
- Lessons and takeaways from existing projects: Where were the problems, what went well, what was difficult?
- Where do people feel passionately? What ideas, problems elicit strong emotions?
- How do technologists feel they can best interact with the users?
- How can we sustain IRI and keep it alive through communities of practice and fruitful partnerships?

Key thematic insights from Sprint 2 informed by Sprint 1

Overarching classes of IRI science patterns

Time-sensitive patterns

Requiring temporal end-to-end urgency. For instance, experiment steering, near real-time event detection, deadline scheduling to avoid falling behind.

Data integration-intensive patterns

Requiring combining and analyzing data from multiple sources. For instance, data from multiple sites, experiments, and/or simulations.

Long-term campaign patterns

Requiring sustained access to resources over a long time to accomplish a well-defined objective. For instance, sustained simulation production, large data (re)processing for collaborative use.

Top challenges to realizing IRI vision

Co-operating across facilities and resources to enable integrated workflows. Using shared resources for real-time experiments, selecting diverse compute resources (HPC, on-demand/cloud, local compute, ...).

Federating access while maintaining cybersecurity. Having congruent use policies, addressing cybersecurity when integrating cross-lab and cross-facility resources.

Building workflows, interfaces, and automation that accommodate heterogeneity, storage, compute at scale globally and across the **data life cycle**.

Keeping the user in focus at all stages. Having appropriate levels of support to bridge software development, end users, prototype to production, ...

Making solutions portable and scalable. Accommodating non-standard or bespoke solutions while emphasizing general solutions that can be leveraged widely (pipelines, APIs, FedID).

Fostering a community of practice around IRI. Enabling good governance, transparency, opt-ins vs. mandates, cooperative technical decisions, etc.

We will build the Design Phase around these key insights

Launching the Design Phase

Design

- Develop one or more IRI **architecture blueprints** that address ensembles of design patterns.
- Develop operational principles for achieving an SC IRI ecosystem.

IRI ABA Design Phase - Late June to Early August

- Focus groups will form on key themes from the Insight Gathering Phase: IRI patterns and IRI "practice" areas.
- Architecture blueprints will be created for each IRI pattern – each considering the IRI practice areas.

1 IRI Pattern and Practice Groups meet virtually over 4 weeks to develop the structured inputs to the Architecture Blueprints.



2 Group facilitators and the LG draft blueprints drawing directly from the group inputs.



3 Iteration event (virtual) in late July/early August with all participants to review the blueprints, exchange on principles, and revise their respective artifacts.

4 **End products**



Cross-cutting issues of workforce, resource allocation, AI/ML, governance, etc. will be addressed within the groups' artifacts.



22

What's next

A survey will be released shortly for you to select your first and second choices of the group you would like to participate on. This will be a quick turn-around.

- We will do our best to accommodate interests in building balanced groups.

Plan to meet several times (virtually) with your group in late June through July.

- Total time commitment of ~4-6 hours.
- Each group will draft structured inputs for the architecture blueprints.

Plan to attend the virtual iteration event in late July/early August (to be announced shortly)

- Total time commitment of ~4-8 hours over two days (not necessarily contiguous).

New IRI ABA Participants

- Please respond to our participant interest survey if you have not done so, as this is crucial to onboarding (slack, google drive of materials, etc.).



23

Thank you!

Discussion / Q&A



24

Additional Materials

IRI Science Pattern Areas

IRI Science Patterns	Description
Time-Sensitive Patterns	Workflows that have time critical/sensitive requirements, e.g., experiment steering, near real-time event detection, deadline scheduling to avoid falling behind.
Data Integration-Intensive Patterns	Analysis of combined data from multiple sources. Can include data from multiple sites, experiments, and/or simulations. Tracking metadata and provenance for reproducible science. Interactive analysis of data, possibly at scale.
Long-term Campaign Patterns	Sustained access to resources at scale over a longer time needed to accomplish a well-defined objective. Robustness, reproducibility, and reliability are important to accomplish. Likely to involve significant logistical planning. Examples include sustained simulation production and large data (re)processing for collaborative use.

IRI Practice Areas

IRI Practice Area	Description
Resource Co-Operations	Allocations/provisioning of multiple heterogeneous resources must be aligned in time and planned in advance to enable integrated workflows. IRI requires new levels of cooperation, collaboration, co-scheduling, and joint planning across facilities and across DOE programs.
Cybersecurity and Federated Access	Users need seamless access and consistent services from distributed research infrastructure, while lab cyber personnel operate under federal cybersecurity requirements and policies and facility operators have different missions and requirements, across the lab complex. Balancing these constraints can also lead to sources of impedance.
User Experience	Understanding users' needs and experiences is critical to technologists' ability to develop effective IRI solutions. This group will engage on approaches for enabling users: requirements gathering, user-centric (co)-design, liaising approaches, etc. (There are implications for all other practice areas.)
Workflows, Interfaces & Automation	System components need to be composablely assembled into end-to-end pipelines across facilities to support IRI science cases. Users should be able to manage these overlays and middlewares effectively across facilities.
Scientific Data Life Cycle	Users need to manage their data across facilities and time from creation (incl. metadata), staging, movement, storage, dissemination, curation, archiving, publishing, etc. Technologists need to develop IRI solutions that accommodate diverse requirements of different research communities.
Portable / Scalable Solutions	Users and technologists need to move/translate their efforts across heterogeneous facilities (be portable) as well as go from smaller to larger resources (be scalable).

Appendix F

THE INTEGRATED RESEARCH INFRASTRUCTURE ARCHITECTURE BLUEPRINT ACTIVITY (IRI ABA) DESIGN PHASE CONVERGENCE ITERATION EVENT

The Integrated Research Infrastructure Architecture Blueprint Activity (IRI-ABA)

Design Phase Convergence Iteration Event

August 1-4, 2022

Code of conduct

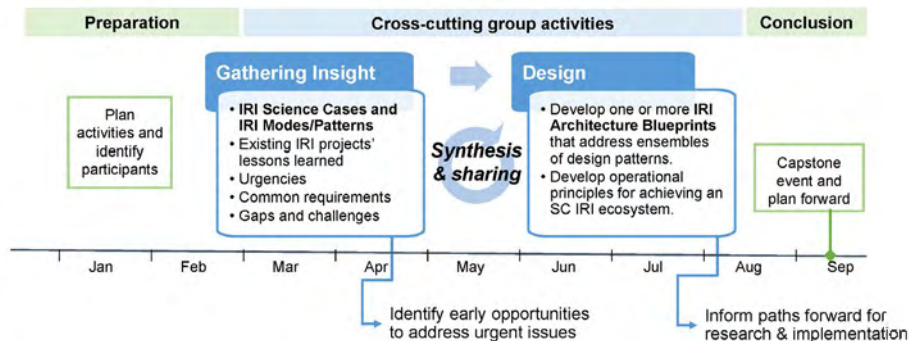
The DOE Office of Science (SC) is fully and unconditionally committed to fostering safe, diverse, equitable, and inclusive work, research, and funding environments that value mutual respect and personal integrity.

If anyone involved in this event has questions or concerns, please approach the HQ Executive Sponsors.

more at >> <https://science.osti.gov/sc-2/Research-and-Conduct-Policies/Diversity-Equity-and-Inclusion/SC-Statement-of-Commitment>

SC IRI Architecture Blueprint Activity, FY 2022

Goal: Produce the **reference conceptual foundations** to inform a coordinated "whole-of-SC" strategy for an integrative research ecosystem.



The Activity in a Nutshell

IRI Science Cases → IRI Modes / Patterns → Designs for IRI

Bring the community together to explore **IRI science cases**

Create an **IRI framework** to reveal patterns, challenges, gaps

Design the **IRI architecture blueprints** for one or more IRI design patterns

Research Infrastructure Perspective

Science Perspective

U.S. DEPARTMENT OF ENERGY Office of Science

photo credit: Didrik Johndk
https://www.flickr.com/photos/deetrik/623632892

IRI ABA Design Phase - Late June to Early August

- Focus groups will form on key themes from the Insight Gathering Phase: IRI patterns and IRI "practice" areas.
- Architecture blueprints will be created for each IRI pattern – each considering the IRI practice areas.

- 1 IRI pattern and practice groups meet virtually over 4 weeks to develop the structured inputs to the architecture blueprints.
- 2 Group facilitators and the LG draft blueprints drawing directly from the group briefs.
- 3 Virtual event in early August with all participants to review the Blueprints, exchange on principles, and revise their respective artifacts.
- 4 Design Phase Products

IRI Pattern Group: Time-Sensitive Patterns	IRI Practice Group: User Experience
IRI Pattern Group: Data-Integration Patterns	IRI Practice Group: Resource Co-Operations
IRI Pattern Group: Long-Campaign Patterns	IRI Practice Group: Scientific Data Lifecycle
IRI Practice Group: Workflows, Interfaces & Automation	IRI Practice Group: Cybersecurity & Federated Access
IRI Practice Group: Portable/Scalable Solutions	

Convergence - Iteration Event

Architecture Blueprints:
Time-Sensitive Patterns
Data-Integration Patterns
Long-Campaign Patterns

Practice Group Briefs*:
addressing key IRI practice areas

Overarching IRI Principles, Issues, Opportunities

* IRI practice groups are the seeds of an IRI community of practice

U.S. DEPARTMENT OF ENERGY Office of Science

Updated for Aug 1-4, 2022

5

Goals for the Convergence/Iteration Event

- **Maximum cross-pollination!**
 - Practice groups and patterns groups interact and exchange
- Advance the **IRI blueprint drafts** together
 - These provide a framing *not* the implementing solutions
 - "80% is good" continues to be the rule
- Advance the **practice group briefs**
 - Incorporate what was learned from the blueprint exchanges
- Capture **overarching IRI principles, issues, opportunities**
 - Also a home for ideas that don't fit in the blueprints and briefs

U.S. DEPARTMENT OF ENERGY Office of Science

6

IRI Architecture Blueprint Activity Artifacts

- **Gathering Insight**
 - Synthesis of IRI Science Cases and User Perspectives (Sprint 1)
 - Existing IRI Projects Survey and Lessons Learned (Sprint 2)

- **Design Phase**
 - IRI architecture blueprints organized by IRI pattern classes
 - IRI practice areas – group briefs
 - Overarching IRI principles, concepts, and issues
 - *Commonalities and differences in the blueprints for different pattern classes?*
 - *Ideas and options for IRI governance/steering?*
 - *Other important issues to address.*

- **Capstone event** (September)
 - Plan forward for IRI



7

Our thanks to the LG and Design Phase facilitators!

Debbie Bard	Eric Lançon
Johannes Blaschke	John MacAuley
Amber Boehnlein	Jini Ramprakash
Shane Canon	Nicholas Schwarz
David Cowley	Arjun Shankar
Kjiersten Fagnan	Tom Uram
Chin Guok	



8

Day 1 Convergence/Iteration Aug 1, 2022

Plenary <ul style="list-style-type: none"> • Welcome. Goals for the event and today. • Quick scroll through the blueprint drafts. 	All	30 min
Blueprint Breakouts <ul style="list-style-type: none"> • Time Sensitive • Data-Integration Intensive • Long-Term Campaigns Each breakout reviews & refines draft blueprint. Capture issues to resolve. Breaks during the session as desired.	<ul style="list-style-type: none"> • Patterns group members go to respective blueprint session • Practice group members <u>divide up</u> to cover each blueprint breakout 	1:15
Break		5-10 min
Practice Group Meetings <ul style="list-style-type: none"> • Practice groups update their contributions based on the blueprint discussions. 	<ul style="list-style-type: none"> • Practice groups meet • Patterns group members <u>divide up</u> to cover each practice group meeting 	1 hr



9

Day 2 Convergence/Iteration Aug 2, 2022

Plenary <ul style="list-style-type: none"> Welcome. Goals for the event and today. Q&A and top issues identified in Day 1 	All	30 min
Blueprint Breakouts <ul style="list-style-type: none"> Time Sensitive Data-Integration Intensive Long-Term Campaigns Each breakout reviews & refines draft blueprint. Each has 1-2 people capture issues to resolve.	<ul style="list-style-type: none"> Patterns group members go to their respective blueprint session Practice group members should divide up and attend each breakout 	1:15 hr
Break	All	5-10 min
Plenary <ul style="list-style-type: none"> Lightning talks from groups (5 min ea.) Vision, top goals, top issues/challenges. Open mic, Q & A 	All	1 hr

Day 3 Convergence/Iteration Aug 3, 2022

Plenary <ul style="list-style-type: none"> Welcome. Goals for the event and today. Today's focus: higher-level IRI questions <ul style="list-style-type: none"> Overarching IRI principles Blueprint compare/contrast Ideas and options for IRI governance/steering 	All	30 min
Breakouts on the focus questions <ul style="list-style-type: none"> Overarching IRI principles Blueprint compare/contrast IRI governance/steering ideas/options 	Participants choose among the topics	1:15 hr
Break		15 min
Plenary <ul style="list-style-type: none"> Readouts from the breakouts Discussion 	All	1 hr

Convergence/Iteration: Higher-level IRI focus and analytic questions

Topics for the breakouts

- Overarching IRI principles** (Facilitators: Debbie & Jini)
 Might also include consideration of common language/layer issues
- Blueprint compare/contrast** (Facilitators: Blueprint facilitators)
 Focus on the key distinctions among the patterns and among their respective driving requirements (and the commonalities as warranted)
- Ideas and options for IRI governance/steering** (Facilitator: Arjun)

Task for each

- A challenge exercise to **draft a 1-pager** that articulates the topic, drawing from all the preceding IRI ABA work and your experience
- Each should include an answer the question: **What is next?** (i.e., next steps)

Day 4 Convergence/Iteration Aug 4, 2022

Plenary - Welcome. Goals for today.	All	15 min
Breakouts: Refine topical 1-pagers from day 3 <ul style="list-style-type: none"> Overarching IRI principles Blueprint compare/contrast IRI governance/steering ideas/options Task <ul style="list-style-type: none"> Hear from new voices. Prioritize for importance and urgency. What is next? Early wins we could achieve in FY 2023? Note key opportunities for the research community. Compose a 1-2 sentence message at the top to a senior executive or lab director. 	Participants choose among the topics	1 hr
Plenary <ul style="list-style-type: none"> Readouts from the breakouts Summing up, open mic - further discussion Looking ahead. 	All	45 min
Facilitator and LG post-event meeting	LG & Facilitators	30 min

Thank you!

Q&A

Appendix G

IRI ABA BLUEPRINT: TIME-SENSITIVE PATTERNS

IRI ABA Blueprint: Time-Sensitive Patterns

Definition of This Pattern

Time-sensitive integrated research infrastructure (IRI) patterns comprise workflows that have time-critical/sensitive requirements (i.e., real time or near-real time), which can be motivated by various factors, such as:

- Decisions that cannot wait. These are policy decisions, based on simulations that are critical, for example, predicting wildfire spread, flooding severity, or the path of a hurricane.
- Experiment control. Without the right control, the observation will be missed or the instrument is not used well, and science is wasted.
- Virtual proximity. This means coupling unique research assets at multiple labs to function as if in the same room, to enable new research not otherwise possible when working independently.
- Loss of fidelity. Instruments produce data, and much of the time, they cannot all be captured, or getting all is very hard. This leads to two sub areas:
 - Data movement from the instrument to the High Performance Computing (HPC) center with limited buffering, is a problem. If data cannot be transported at the correct rate, it falls on the floor.
 - Edge computing. The data stream cannot be sent in entirety to the HPC center and needs processing in real time or near-real time for either data analysis and reduction or control.

These time-sensitive workflows involve integration across multiple facilities and resources, and are found in many science domains, such as beamline-based materials science, astronomy and astrophysics observational science, and experimental fusion science.

The key common factor is time criticality of the workflow to be able to perform the experiment or observation or otherwise accomplish the science goals.

1. Representative Cases for This IRI Pattern

The use cases highlighted here are grouped into three areas: experiment control, distributed systems administration, and data management.

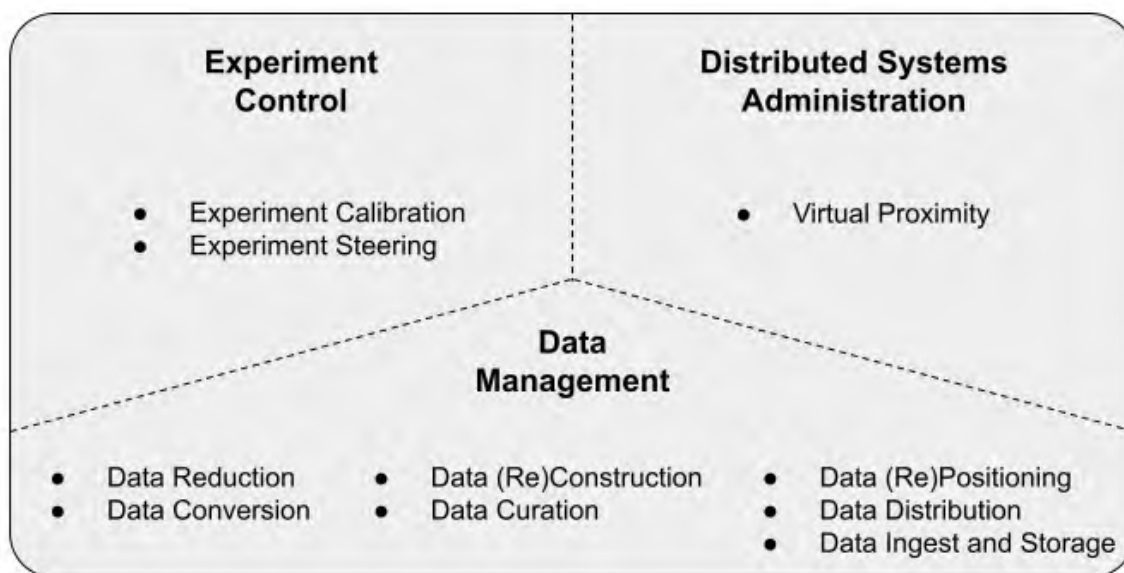


Figure 1. Time-sensitive representative cases

Experiment Control

Experiment Calibration

Experiment calibration requires the processing of data in real time to provide a feedback mechanism that prompts or directs a change in the "physical" hardware configuration of the experiment to get the optimal data quality/throughput.

The time sensitivities for experiment calibration are in the microseconds to minutes range and include activities such as:

- Changing focus, mirrors, or apertures in a synchrotron or free-electron laser beamlines. (APS, NSLS-II, LCLS) [*Time sensitivity: 200–500 msec*]
- Counteracting temporal drifts (temperature/wear). Infrared x-ray pulse characteristics—single pulse, double pulse, polarization and energy spectrum—serve as driver diagnostics to inform autonomous accelerator control, but the time to energy calibration in the diagnostic varies with the temperature variation in the hutch. Instrumentation is used to monitor the quality of the diagnostic, which is featurized using machine learning (ML) in Field Programmable Gate Arrays (FPGAs) (or specialized accelerators) at the source. When calibration is required, raw data from the experiment are streamed to specialized HPC at the ALCF for retraining, essentially a recalibration, with a turnaround time requirement of 15 minutes. The deformation of the floors in the experimental area due to temperature affects the Dynamic Reaction Microscope (DREAM) endstation in the time-resolved molecular optics instrument and the qRIXS

endstation. Realignment and recalibration of the instrument may be necessary in response to temperature drifts caused by diurnal cycles. The high-resolution monochromator for qRIXS may require recalibration with temperature and other sources of slow-motion drift. In anticipation of this, the instrument uses a diagnostic reflection from the mono exit slits, which requires an online photoemission measurement. We have observed drifts between the arrival time at the Arrival Time Monitor (ATM) and the interaction point (IP) that occur on the timescale of hours. As a result, we need to occasionally redefine the time mapping between ATM and IP. (LCLS) [*Time sensitivity: 15 mins*]

- Adjusting detector power supplies to keep experiment calibration constant. (CEBAF) [*Time sensitivity: 1–30 mins*]
- Pre-experiment calibration. To automatically steer an experiment, one must first configure the experiment: the beam must be steered into the hutch and focused onto the sample being probed, beam energies and profiles must be measured, timing information must be determined, detector responses and geometries need to be measured, and the position of each detector relative to the IP needs to be measured. There is an explicit calibration procedure that happens at the beginning of each experiment to set up a standard and known experiment configuration. Diagnostics are used to monitor for deviations from the optimal configuration. (LCLS) [*Time sensitivity: mins to hrs*]
- Real-time calibration. Nuclear physics detectors generate data at gigabytes per second for weeks to years of running time. Real-time monitoring of data quality requires conversion of the raw data from the detector into quantifiable physical values. This requires calibration of the detector during data taking. (CEBAF) [*Time sensitivity: keep up with data rate*]

Experiment Steering

Processing of data in real time to provide a feedback mechanism that prompts or directs an adjustment to the search pattern/dimension to get the optimal scientific result. This use case includes both “intra” and “inter” facility workflows. The former is confined within the well-defined boundary of an experiment, while the latter involves coordination across multiple experiments.

The time sensitivities for experiment steering are in the microseconds to minutes range and include activities such as:

- Protein Crystal Screening at the Coherent X-ray Instrument enables increased access to LCLS beamtime for biological structure determination by making use of short runs to screen the quality of different sample preparations or potentially collect a full data set under good conditions. Samples may be switched out every few hours. For each new sample, information about hit rate and quality of the sample preparation must be quickly evaluated. Determination of biological structure is a computationally challenging problem that requires rapid acquisition and characterization of data to ensure that results can be had within a few hours. During crystallography experiments, different sample

- preparations and sample delivery methods are often attempted, to determine which method and sample deliver the best result. (LCLS) [*Time sensitivity: 5–20 mins*]
- Adjust/prioritize follow-up resources for astrophysical surveys without human intervention. In astronomy it is common for a telescope to discover scores to thousands of events during a night with a wide-field survey telescope. But based on the discovery, the science might trigger a different set of telescopes to examine the same location in the sky. (Astrophysics) [*Time sensitivity: acceptable < 30 mins; objective < 10 mins*]
 - Highly automated beamlines could intelligently explore scientific problems. Beamlines equipped with high-throughput capabilities for fast measurements can explore different parameters, such as the effects of temperature, vapor pressure, and humidity. Real-time data analysis can be used to feed experimental decision-making. The ability to take measurements quickly is important, but not enough for revolutionary materials discovery because the material parameter spaces are very large and multidimensional. One experiment's parameter space may have multiple dimensions and tens of thousands of distinct points within that space to explore. Identifying the location in parameter space where information gain can be maximized can enable the targeted exploration of that region of interest. In the case of stochastic events, experiment monitoring can identify when a deformity or event of interest is beginning and then adjust the spatial or temporal scanning granularity to maximize the amount of information obtained. These sorts of adjustments would ideally be made as the experiment progresses. (LCLS) [*Time sensitivity: 200–500 msec*]
 - Coupling of simulations and development of “digital twins” to light source experiments to potentially unlock new materials science knowledge to, for example, better understand failure modes in materials, enable the synthesis of new materials, aid in the creation of purpose-built designer materials, and assist in additive manufacturing processes. This will also allow for a more efficient and optimum use of beamtime at the light sources. The coupling of simulations with light source experiments can be split up into three main areas in the experimental life cycle. First, before the experiment, simulations can be used to help prepare, plan, and determine if the experiment is even feasible. Second, during the allocated beam time, simulations can help guide and inform the strategy and guide the experiment. Finally, simulations can be used to aid in the data analysis to extract the maximum scientific information from the data. (LCLS) [*Time sensitivity: 200500 msec*]
 - Development of experimental machine control configurations based on results of more complex analysis than can be provided in real-time control systems. Especially relevant in pulsed nuclear fusion devices. These analysis tasks traditionally inform operators who configure the plasma control systems ahead of the next pulse. (NSTX-U, D-IIID) [*Time sensitivity: <15 mins*]

Distributed Systems Administration

Virtual Proximity

Virtual proximity is focused on coordinating the operation of physically separated research assets through the real-time exchange of command-and-control information. This is to address the need to use multiple distinct research assets to address the operation of complex systems, in a way that no one lab can do today.

The time sensitivities for virtual proximity are in the microseconds range and include activities such as:

- Emulating the behavior of large-scale complex systems via digital real-time simulations (DRTS) for distributed grid emulation. This latency dominated workflow, driven by the exchange of command-and-control information between unique laboratory research assets at disparate locations that need to operate as if they are residing in the same room, requires *virtual proximity* (i.e., near deterministic low latency) between distributed unique (dedicated) research assets and subject matter experts at multiple labs. In December 2021, NREL and PNNL collaborated on a demonstration project using ESnet OSCARS circuits to provide priority network performance and connected an OPAL-RT system at PNNL with an RTDS (Real-Time Digital Simulator) system at NREL. The goal was to emulate the power system of Cordova, AK, and demonstrate new advanced metering infrastructure (AMI) could ensure power to critical loads (e.g., a hospital and airport) in the event of system-level power disruptions. In this demonstration, NREL simulated a real-time digital twin of the Cordova microgrid with AMI controls, and PNNL simulated protection systems for load feeders using hardware relays. For this use case, low latency with very little jitter is crucial. (ESIF) [*Time sensitivity: 35–40 msecs*]

Data Management

Data Reduction

Real-time data reduction through compression, feature extraction, summarizing/aggregating information, or discarding of invalid data, to reduce downstream storage requirements, network bandwidth utilization, or computing bottlenecks.

The time sensitivities for data reduction are in the microseconds to hours range and include activities such as:

- Real-time filtering of science signals from noise and background. Due to image artifacts (systematic electronic noise, misalignment of images, etc.), ML/ artificial intelligence (AI) is employed in astrophysical surveys to separate noise/junk from the real signal. In addition, data from other surveys in the form of catalogs are used to further characterize

the true positives (real astrophysical objects/transients) to provide added scientific value. (Astrophysics) [*Time sensitivity: < 1 min*]

- Post processing of experimental data to reduce need for archival storage. While the original data are always saved, usually to tape, post processing in astrophysical surveys is done on subsections of the data with higher value, and kept on a spinning disk, to provide real-time analysis capabilities for the scientists. (Astrophysics) [*Time sensitivity: < 12 hrs*]
- The vast majority of instrument triggers on the Large Area Telescope (LAT) instrument of the Fermi Gamma-ray Space Telescope are background events caused by charged particles as well as earth albedo gamma-rays. To minimize the effects on the instrumental deadtime associated with reading out the LAT, an onboard filter (software algorithm) is used to eliminate a substantial number of background events without sacrificing the celestial gamma-ray events that are of interest with the goal that the resulting data that passes the filter can be transmitted to the ground within the available bandwidth. Instrument triggers occur at 1–10 kHz with about 1.5 kHz of events passed to the ground. Of these, 1–10 Hz are the desired celestial gamma-rays. (Fermi Gamma-Ray Space Telescope) [*Time sensitivity: 1ms*]
- The Fermi Gamma-Ray Space Telescope also has onboard science processing to provide rapid detection and localization of gamma-ray bursts (GRBs). The output of this processing can trigger an autonomous repointing of Fermi to keep the GRB within the LAT's field of view for observation of high-energy afterglows by other observatories. The onboard science processing consists of algorithms to select gamma-ray candidate events, reconstruct directions of gamma-ray candidate events, and search for and localize high-energy transients. Event selection is based on parameters previously calculated for the onboard filter. This algorithm needs to keep pace with the rate of incoming data. Once the data are on the ground, automated science processing pipeline performs reconstruction of classified events to facilitate the timely follow-up observations by other observatories. (Fermi Gamma-ray Space Telescope) [*Time sensitivity: < 1hr*]
- Removing long gaps (e.g., suppress below threshold signal) between useful data in parallel time-sequenced data streams. (CEBAF) [*Time sensitivity: 1ms to 1s*]
- Grouping time slices from all detectors and leveraging AI/ML feature extraction to look for "events" at data rates from 1 GByte/s to 1 TByte/s. (CEBAF) [*Time sensitivity: 1ms to 10s, criticality is matching the feature extraction rate to the data rate*]

Data Conversion

Real-time data conversion to facilitate additional downstream analysis, or comply with security requirements (e.g., real-time encryption).

The time sensitivities for data conversion are in the microseconds to seconds range and include activities such as:

- Converting data from formats optimized for data taking to formats optimized for archiving. One example of image formats for analysis is the one used in ptychography. An HDF5-based structure convention called Coherent X-ray Imaging (CXI) format [<https://www.cxidb.org/cxi.html>] is used for the essential data and metadata required by the analysis. Computationally intensive ptychography codes leverage the data format and metadata to do fast image reconstruction from diffraction data. (LCLS) [*Time sensitivity: several mins*]
- XTC2 is a self-describing scientific data format used in LCLS data analysis codes, such as psana, and contains data produced by the online data acquisition (DAQ) system. HDF5 is a general-purpose storage middleware used in many scientific applications and supported in many analytical tools. LCLS uses XTC2 for the raw data because it has performant access from C++/python, does not require serialization for network transport, and has the same format in-memory and “on the wire” (when data are transmitted over the networked or persisted to a file). HDF5 is often used for the first-pass reduced data, reduced in size by the user using experiment-specific knowledge. In the case that the LCLS real-time data reduction pipeline does not produce fully featured extracted data, it is necessary to process the data and produce HDF5 results files to assist in data analysis. Current HDF5 read/write performance for variable length (reduced) data is not sufficient to provide low-latency data access. Variable length data is slightly faster (5%) in writing and significantly slower (20%) for reading as compared with fixed length. HDF5 is slower in all cases than the binary files by 20–50%. (LCLS) [*Time sensitivity: < 10 mins*]
- LAT science data are compressed onboard and sent to the ground for processing. The processing pipeline is designed to allow parallel processing of events. About 300 Hz of downlinked on-orbit data can be processed by 100 cores within 1 to 2 hours, allowing processing to finish before the next downlink arrives. Reconstruction inflates the raw science data by a factor of 20. Reconstructed gamma-ray photon events are made available along with instrument response functions and high-level analysis tools to the Fermi Science Support Center for distribution to the community. (Fermi Gamma-Ray Space Telescope) [*Time sensitivity: < 1 hr*]

Data (Re)Construction

Real-time (re)construction of data (using a single or multiple sources (e.g., data merging/association)) to ensure data quality.

The time sensitivities for data (re)construction are in the microseconds to minutes range and include activities such as:

- Reconstructing data using both traditional analysis and AI/ML training in parallel, with the goal to have the AI/ML model replace the traditional analysis process (at the edge),

resulting in lower compute resource requirements. Such low-latency feedback also enables rapid higher-level decision-making, e.g., phenomena/feature detection, and therefore provides opportunities for experimental steering. Accelerated analysis or reconstruction of measurement data collected during X-ray experiments can be crucial depending on the scientific goals. For instance, detecting dynamically evolving features and adjusting experimental controls according to the features' states can provide opportunities for goal-oriented data acquisition while significantly decreasing the redundant data collection. However, such workflows require coupling the experiment with data analysis pipelines that can provide rapid feedback capabilities, close to the target temporal resolution of the experiments (or speed of features' state transitions). For data-intensive image analysis tasks, e.g., iterative reconstruction for micro/nanoCT or 2D/3D ptychography experiments, high fidelity/accuracy ML models can be used for rapid feedback. Some of the motivating experiments include observing microstructural changes during charging/discharging batteries, detecting chemical reactions during cement curing, crack detection and propagation in materials, and adjusting radiation exposure to biosamples to keep features intact. (APS) [*Time sensitivity: <1 min*]

- Calibrating data sets by taking in data in well-known conditions (i.e., calibration run), calculating the calibration, and then applying the calibration to the bulk dataset. With streaming data acquisition, the much larger dataset is self-calibrating (i.e., the calibration dataset is buried in the larger raw dataset). With enough real-time processing (i.e., < few mins) the calibration data can be filtered out of the data stream and a calibration calculated and applied to the physics data in time to be useful for experiment steering and monitoring. (CEBAF) [*Time sensitivity: < 10 mins*]
- Astrophysics often combines data from several pre-existing surveys, usually in the form of catalogs stored in large databases, with streaming survey data to best characterize real astrophysical sources and transients. (Astrophysics) [*Time sensitivity: several mins*]

Data Curation

The conditions under which data were collected, as well as their subsequent processing, are important factors in extracting a quality science result. Frequently the metadata that provide this information are synthesized and become decoupled from the raw data due to delays in processing. Prompt generation of and tagging with metadata is important to preserve data integrity. A well curated data set allows future researchers to establish the provenance of a published result. It also allows current researchers to track how data were collected and processed.

The time sensitivities for data curation are in the microseconds to minutes range and include activities such as:

- Synthesizing multimodel data streams (e.g., raw data, calibrations, conditions, detector geometry, even the processing code itself) to identify and preserve metadata for experiment reproducibility. Documenting the steps from experiment to result (e.g., results

-> analysis out -> reconstruction output -> raw data), as well as what “happened” to the data is important (e.g., given a raw data set, one should be able to determine if it has already been processed, what were the conditions, and where is the result of the processing.). The time-critical component is metadata added while the data are being taken and before they are archived. At the millisecond timescale, metadata are added documenting where and when the data originated, as well as any experiment conditions that change rapidly. At the timescale of seconds to minutes, additional metadata are added to document slower-changing variables. (CEBAF) [*Time sensitivity: msec to secs, between acquisition and storage*]

Data (Re)Positioning

Timely transfer of data (from instrument site local storage) to designated off-site storage to prevent oversubscription of local storage resources and resulting loss of data.

The time sensitivities for data (re)positioning are in the seconds to days range and include activities such as:

- Transferring of data from the experiment local storage (e.g., LHC Tier 0 (CERN)) to off-site storage/processing facilities (e.g., LHC Tier 1 - FNAL, BNL) to prevent the data in the local disk buffers from getting flushed to tape, and requiring data to be restaged back to the disk buffers if an external transfer is requested. (LHC) [*Time sensitivity: 15 days*]
- Timely data movement between fast-access storage layers and archival layers. (LCLS) [*Time sensitivity: 100–500 msec*]
- Providing multiple access methods for data storage where co-location is inefficient, such as an ITER-like access method to analyze tokamak data as they are produced by other devices. This would facilitate the development and testing of automated analysis routines ahead of the commissioning of other similar devices. (ITER) [*Time sensitivity: seconds for post shot analysis, msec for “real-time” control*]

Data Distribution

Real-time broadcasting of information or a subset (e.g., different streams/packets/events) to different consumers (e.g., pub/sub) for notification and/or follow-up action.

The time sensitivities for data distribution are in the seconds to mins range and include activities such as:

- Distributing of scattering or light source application data to different subscribers (e.g., data translation/conversion/cleaning, live statistics reporting, different visualizations and analyses of the experimental data, anomaly detection, data inspection, user alerting system, provenance capture mechanisms, buffering for data storage.) Data distribution occurs at different stages of the end-to-end pipeline, starting from DAQs to event identification, and ending at the consumption of events by one or more science

applications. Raw data distribution from DAQs to the first layer of compute is very different from a metadata + event stream pushed through a lower rate pub/sub pipeline to consumers of the experiment. (Light sources) [*Time sensitivity: 10–100 secs*]

- In astrophysical surveys after the raw data is processed and analyzed, vetted transients are streamed in real time to “brokers,” which run additional filters for consumption by humans. (Astrophysics) [*Time sensitivity: < 10 mins*]
- Support of remote collaboration sites and control rooms that could replicate the conditions found in the central control rooms. This would be applicable to a number of experimental operations facilities, to support scenarios such as real-time graphics or events that launch computational jobs. (D-IIID, ITER) [*Time sensitivity: < 1 sec*]
- Supernova Neutrino Early Warning System (SNEWS) (an NSF-funded project) aims to collect alerts from detectors capable of detecting supernova events, and then distribute those alerts to vested/interested parties in near-real time. (DUNE) [*Time sensitivity: 20 secs*]

Data Ingest and Storage

Real-time storage of data to mitigate against (real-time) data manipulation (e.g., reduction, conversion) pipeline failures to ensure that no data are lost. Analytics (at the edge) can be done in parallel to reduce sequential steps (e.g., store first and process later) and shorten time to results. The following detail the importance of real-time storage:

- Experiment schedules are set by beamline availability. Having a rigid allocation of a specific compute resource reduces the ability to be agile and move workflows to other resources “on the fly” while data are being taken.
- Experimenters are notoriously unable to accurately estimate computing needs. The ability to dynamically allocate or deallocate resources mitigates this.
- Real-time storage can support smaller experiments that might not be able to develop a complete real-time workflow.
- Data storage models frequently couple the physical aspects of how data are stored to the policies that determine how data are treated. An example is specific `/work`, `/cache`, `/volatile` posix paths where the name of the path determines both aspects.
- Data compression provides a direct benefit to storage, by increasing capacity as well as storage write speeds for the smaller data set. Implementing real-time compression that is transparent to the compute nodes makes for a very attractive use model.
 - All applications inherit the benefit of compression in the infrastructure layer, without any implementation overhead.

- Real time compression requires hardware like smartNICs. Selecting a single implementation in the infrastructure layer avoids any application specific requirements or coupling between the application and the hardware.
- Performance of the storage system can be updated and scaled based on best-in-class storage hardware, without affecting any application code.
- Real-time compression requires a data format and API that supports hardware real-time compression transparently.

The time sensitivities for data ingest and storage are in the microseconds to seconds range and include activities such as:

- Acquiring data from an instrument and positioning it for prompt processing in a way that does not impede flow of data from the instrument. Large instruments have data sources distributed around the hardware. Data are acquired from these sources in parallel and timestamped. Data must be transported to the physical location where prompt processing will take place without back pressure that would slow or destabilize data taking. (CEBAF) [*Time sensitivity: msec to few secs*]
- Rate matching between varying data rates from an instrument and varying ingest rates for a data center. (LCLS) [*Time sensitivity: 10 msec to 10 secs*]
 - When analyzing streaming data from LCLS, the goal is to keep up with the input rate and produce a result within seconds to minutes of stopping data acquisition. When data are streamed from LCLS to remote compute, such as NERSC, data actually pass through a number of hardware layers along the way, each of which can introduce latency: detector -> data reduction pipeline (point to point connection, but DRP can introduce latency depending on the nature of the data and the analysis) -> data passing DRP written to NVMe-based data cache -> transfer to spindle-based disk via ethernet -> DTN at SLAC border -> ESnet -> NERSC border -> burst buffer -> compute. At each stage where processing is done, we are prepared to buffer data if the processing takes longer to accomplish. We can also introduce controlled deadtime if the DRP cannot keep up. For crystallography, events are coming in every 10 us, calculating a veto takes of order 12 ms, peak finding takes of order 200 ms/event, and indexing about 15 seconds/event and structure determination.
 - Need to be able to launch jobs quickly, within a minute of starting acquisition, to return the results within a few minutes of finishing data taking.

2. What Are the Principal End-to-End Performance Factors, Technical and Operational Components, Interfaces, and User Experience Requirements?

This section will focus on the factors, components, interfaces, and requirements that cross-cut across the representative cases described previously. The subsections outlined here are grouped within the areas as highlighted in the DOE ASCR IRI whitepaper,¹ which are allocations, accounts, data, applications, scheduling, workflows, publications, archiving, policies and governance, and engagement and partnerships.

Allocations

Guaranteed/Assured Computing Resources When Needed

Availability of appropriate compute resources (e.g., HPC, cloud, edge compute, etc.) is critical in support of active experiments as they happen. Equally important to availability is elasticity of compute resources to scale dynamically with the need. Having dedicated compute resources can be costly, and using shared resources requires guarantee of availability. Examples of this include securing compute resources for traditional analysis and AI/ML (re)training for data inference while an experiment is running, and the ability to run psana (Photon Science ANALysis) which is a software package used to analyze data produced by LCLS on supercomputer architectures by changing the parallelization technology to allow scaling from hundreds of cores to hundreds of thousands.

- *Representative cases: Experiment Calibration, Experiment Steering, Data Reduction, Data Conversion, Data (Re)Construction, Data Curation*
- *IRI ABA practices:*
 - *Resource Co-Op*
 - *Scientific Data Lifecycle*

Accounts

Common/Interoperable Security Frameworks Across Facilities

Accessing resources across facilities (and in different administrative/security domains) can be nontrivial if each facility has a different security framework. A common (or interoperable) solution is needed that allows for easier user access and does not compromise the security policies of

¹ B. Brown et al., "Towards a Seamless Integration of Computing, Experimental, and Observational Science Facilities: A Blueprint to Accelerate Discovery," DOE ASCR IRI Task Force white paper, Mar 8, 2021, <<https://doi.org/10.2172/1863562>>.

each of the participating facilities. An example of this would be to deploy a FederatedID ecosystem across the DOE SC facilities.

- *Representative cases: Experiment Calibration, Experiment Steering, Data (Re)Positioning, Data Distribution, Data Ingest and Storage*
- *IRI ABA practices:*
 - *Cybersecurity and Federated Access*
 - *Resource Co-Op*
 - *User Experience*
 - *Portable/Scalable Solutions*
 - *Workflows Interfaces and Automation*

Data

Accessibility of Data

Uniformly implement Findable, Accessible, Interoperable, and Reusable (FAIR) data principles to accessing and sharing data, for example, at the five light sources or simulated data produced at local, university, or ASCR compute facilities, and AI/ML algorithms trained on a variety of data. Specifically, FAIR data principles provide uniform or standardized metadata collection, data provenance capture, accessibility for multiple users with appropriate access security, handling of multiple data sources and types, Digital Object Identifiers (DOIs) minting, and mechanisms to make the data discoverable and easy to find.

- *Representative cases: Experiment Calibration, Data Ingest and Storage*
- *IRI ABA practices:*
 - *Scientific Data Lifecycle*
 - *Workflows Interfaces and Automation*

Coordination of Multimodal Data Analysis

Multimodal data analysis from multiple sources (e.g., sample information, simulation, accelerator information, beamline information, etc.) requires coordinated data movement from multiple sources, and data translation functions (if needed).

- *Representative cases: Experiment Calibration, Experiment Steering*
- *IRI ABA practices:*
 - *Cybersecurity and Federated Access*
 - *Resource Co-Op*
 - *Scientific Data Lifecycle*
 - *Workflows Interfaces and Automation*

Time-Sensitive Data Movement

Time-critical movement of data between single/multiple source(s) and destination(s) requiring (scheduling) some level of network bandwidth guarantees (e.g., Quality of Service) and consumer/producer bandwidth/capacity guarantees to provide end-to-end quality of service to prevent data loss, and/or support latency/jitter sensitive applications. There are two types of data movement included here: file-based and memory-to-memory. File-based data movement is typically used when transferring data to remote compute resources (e.g., HPC, cloud, etc.), but memory-to-memory (streaming) data movement is more critical for time-sensitive workflows. Examples of this include streaming of experiment data to compute resources for quick turnaround analysis, broadcasting data to scientific users for follow-up action, and offloading local storage to prevent data loss from space exhaustion.

- *Representative cases: Experiment Calibration, Experiment Steering, Data Reduction, Data (Re)Positioning, Data Distribution*
- *IRI ABA practices:*
 - *Resource Co-Op*
 - *Scientific Data Lifecycle*
 - *Workflow Interfaces and Automation*

Applications

Easy to Use and Well-Supported Programming Constructs

Providing the right level of abstraction for API to bridge the “expert-gap” (e.g., descriptive vs prescriptive) is necessary in promoting ease of use. The choice of libraries and languages should be biased towards ones with good community support.

- *Representative cases: Experiment Calibration, Experiment Steering, Data Reduction, Data Conversion, Data (Re)Construction, Data Ingest and Storage*
- *IRI ABA practices:*
 - *Resource Co-Op*
 - *User Experience*
 - *Workflows Interfaces and Automation*

Common (Portable) Programming and Runtime Environments

Use of alternative computing resources if the initial compute site is down for maintenance. Similar or common services, programming and runtime environments, should exist on different compute resources/clusters (e.g., libraries, APIs, containerized workflows, common file formats, common implementations of data compression, common APIs for orchestration, common transport protocols over TCP/UDP, etc.). An example of this would be the migration of jobs from one HPC to another due to a scheduled/unscheduled outage.

- *Representative cases: Experiment Calibration, Experiment Steering*
- *IRI ABA practices:*
 - *Cybersecurity and Federated Access*
 - *Portable / Scalable Solutions*
 - *Resource Co-Op*
 - *User Experience*
 - *Workflows Interfaces and Automation*

Support for Rapid Prototyping/Coding

Ability to rapidly create new analysis and new workflows, test at small scale, and run at large scale is needed to adapt to rapidly changing experiments, with analysis code and workflows often tweaked and modified in response to new information emerging from experiment results, simulation, or other sources. Users need to be able to ask questions and write analysis to generate answers. While 90% of questions will be repetitive, that 10% that are new are triggered by data emerging from the experiment during beamtime. These are the questions that need an answer *right away*, and they may require significant computing and rapid coding.

- *Representative cases: Experiment Calibration*
- *IRI ABA practices:*
 - *Portable / Scalable Solutions*
 - *Resource Co-Op*
 - *Workflows Interfaces and Automation*

Scheduling

Orchestration and Verification of Resources Across Multiple Facilities

Coordination of multifacility resources may be necessary for real-time fast feedback analysis. For this to be done in an automated way, several functional capabilities are required. First is resource/service discovery: understanding what resources/services are available. Second is resource/service selection: determining the optimal set of resources/services to assemble. Third is resource/service request/negotiation: having the ability to request for the appropriate resources/services for the needed time window. Fourth is resource/service verification: verifying that the appropriate resources/services were allocated/committed. An example of this is the ability to control resource allocation within the exascale system to enable on-demand job scheduling, which may include allocation of different resources at different times, such as burst buffer allocation prior to compute node allocation.

- *Representative cases: Experiment Calibration, Experiment Steering, Data Reduction, Data (Re)Positioning, Data Distribution*
- *IRI ABA practices:*
 - *Resource Co-Op*

- *Workflows Interfaces and Automation*

Shared Use of Specialized Hardware

Specialized or purpose-built hardware may be required for feature extraction, data reduction, and data conversion. The location of specialized hardware may be remote, requiring timely transfer (e.g., streaming) of data to/from facilities, as well as scheduling of the specialized hardware (if it is shared).

- *Representative cases: Experiment Calibration, Experiment Steering, Data Reduction, Data Conversion, Data (Re)Construction, Data Ingest and Storage*
- *IRI ABA practices:*
 - *Portable/Scalable Solutions*
 - *Resource Co-Op*
 - *Workflow Interfaces and Automation*

Workflows

Automated/Programmatic Experiment Control

Automated control of the experiment is needed to remove human-in-the-loop delays if real-time calibration (or autonomous control) of the experiment is required.

- *Representative cases: Experiment Calibration*
- *IRI ABA practices:*
 - *Workflows Interfaces and Automation*

Dynamic Load Balancing of Resources

A mechanism to intelligently distribute data from running experiments to resources in a data center is needed. This system should have the ability to dynamically respond to events in the data center, for example, bottlenecks, overloaded resources, and nodes that need to be removed from service. As well, the load-balancing service should not be experiment or compute-infrastructure specific. Instead, the load-balancing service should provide the necessary hooks to have an experiment guide the load balancer as to which compute node to place a stream of data to, and an abstraction (API) from the compute nodes for backpressure or flow control. Additionally, the load-balancer service should support destination end-point addressing independent of the physical IP address or serial number of the compute node, using instead a mapping between a logical name or address of the node, to a physical address. This enables dynamic re-allocation of resources when needed. The load balancer should include the ability to replicate streams to multiple consumers when needed.

- *Representative cases: Data Ingest and Storage, Data Reduction, Data Conversion, Data (Re) Positioning, Data Distribution*

- *IRI ABA practices:*
 - *Resource Co-Op*
 - *User Experience*
 - *Workflows Interfaces and Automation*

Time-Critical Use and Management of AI/ML Models

Use of AI/ML models (to replace computationally intensive traditional analysis) for feature detection in experiment results. If AI/ML is running at the edge and needs to be responsive to changes in the experiment, the need for rapid streaming to retrain the model may be necessary. A mechanism for storing and tracking models and their provenance will become essential to avoid retraining from scratch, but from the “last, best” model that is available. The ability to identify when the model performance degrades is also essential.

Data generated by experiments, simulations, and digital twins, and by ML models derived from those data for use in digital twins, are used on multiple time and distance scales. For example, data from an in-situ experiment may need to be delivered quasi-instantaneously to an ML model trainer so that it can update the digital twin in time for it to be used to choose the next experiment. Here, the speed with which data relevant to a specific training scenario can be identified and delivered is of the essence. New data and trained ML models can also have value to other scientists, for example to design and steer subsequent experiments and to construct and update other ML models, and thus must also be efficiently accessible to those other parties. These characteristics require a FAIR Data and Model Service (fairDMS) to provide indexing, publication, enrichment, discovery, and access capabilities for both data and trained ML models for the DevOps of ML-based scientific applications. (<https://doi.org/10.48550/arXiv.2204.09805>)

- *Representative cases: Data Reduction, Data (Re)Construction*
- *IRI ABA practices:*
 - *Scientific Data Lifecycle*
 - *Workflow Interfaces and Automation*

Monitoring and/or Verification of Resource/Service Status and/or Performance

To understand how the workflow is performing, the selected resources/services must be able to provide feedback to the workflow so that it can take any appropriate actions accordingly. In addition, usage and access monitoring should be available for auditing purposes, e.g., generating usage reports correlating users and projects. This would imply that the facilities providing the resources/services have the appropriate measurement/monitoring infrastructure in place and are able to provide protected access to such information.

- *Representative cases: Experiment Calibration, Experiment Steering, Data Reduction, Data Conversion, Data Curation, Data (Re)Positioning, Data Distribution, Data Ingest and Storage*

- *IRI ABA practices:*
 - *Resource Co-Op*
 - *Scientific Data Lifecycle*
 - *Workflows Interfaces and Automation*
 - *Cybersecurity and Federated Access*
 - *User Experience*

Workflow Repeatability

Workflows should be repeatable; for example, most of the image reconstruction workflows perform the same computation steps on different experimental data. This also highlights the importance of operational reliability. Workflow blueprints/templates/services can be used to automate and generalize such repetitive operations.

- *Representative cases: Experiment Calibration, Experiment Steering, Data Reduction, Data Conversion, Data (Re)Construction, Data Ingest and Storage*
- *IRI ABA practices:*
 - *Workflows Interfaces and Automation*
 - *User Experience*

Archiving

Performant and Sufficient Storage

Storage resources (local or remote) need to have the predictable performance and appropriate capacity to serve as a data capacity for experiments. This prevents loss of data from pipeline or workflow failures and provides a mechanism to pace the data downstream and smooth out burstiness.

- *Representative cases: Experiment Calibration, Data Ingest and Storage*
- *IRI ABA practices:*
 - *Workflows Interfaces and Automation*

Engagement and Partnerships

User Engagement and Operational Support

Providing assistance on when and how to access the service will be important in its adoption. Well-defined processes to request support and enhancements are also needed if the software is part of a production workflow. Effective user engagement requires not only appropriate

documentation, but potentially workshops, and help desks. Long-term operational support will require funding and governance structures to support the software base.

- *Representative cases: Data Reduction, Data Conversion, Data (Re)Construction, Data Curation, Data (Re)Positioning, Data Distribution, Data Ingest and Storage*
- *IRI ABA practices:*
 - *User Experience*

Practice Group Contributions Specific to This Pattern

This subsection is an assembly of the practice group submissions specific to the time-sensitive pattern.

Cybersecurity and Federated Access Practice Group

We expect that there may be very different cyber implementations across these patterns, driven by unique needs. For example, memory-to-memory transfers across facilities require substantial changes to the existing security architectures of computing facilities but are needed for certain kinds of low-latency integration. Another consideration to time-sensitive workflows is the sensitivity to security incident containment, should, for example, the entire workflow be halted due to a detected security breach, or the jobs continue to run but restrict all external connectivity. However, the goal would be that the design principles and models developed here would be able to accommodate the diverse implementations across the patterns. Exploration of the specific controls needed for different architectures needs in-depth exploration. A primary example here is the diverse security requirements and processes for experiment data and control, where the former is closely aligned with informational technology (IT) practices, and the latter with operational technology (OT) principles (e.g., supervisory control and data acquisition (SCADA)).

User Experience Practice Group

In an ideal world where IRI is wildly successful, a human would be able to:

- Develop a trusted, reliable, and intuitive process that results in work accomplished within the needed time frame.
- Schedule work that requires multiple facilities in one place.
- Control multi-facility workflows from their experiment environment without individually managing separate facility logins/services/allocations.
- Run analysis concurrently with a large experiment.
- Develop software in an environment that reflects the production environment.
- Log into a single system and have easy access to all their data.
- Reduce or eliminate context switching to stay productive.

When time is of the essence, a user needs to be able to move quickly and maintain focus. Some of the outcomes relate to risk mitigation, and the others are to maintaining high engagement and productivity.

In an ideal world where IRI is wildly successful, a robot would be able to:

- Access to job status, success, and failure information via API.
- Submit jobs using well known APIs.
- Determine where a job will run the fastest.
- Move data from experiment to compute with persistent credentials.

The main difference between robots and humans is the need for programmatic interfaces that can determine the characteristics of the computing facility and availability of storage resources and move data quickly. Robots need autonomous access with long-lived credentials.

Workflows, Interfaces, and Automation Practice Group

The most important gap is a standard for interoperable workflows. The end points need to be identified, via a reference architecture, so that job/data services and resource descriptions can be standardized. It is important that monitoring should be considered as an integral part of the system design and implementation, and not as an after thought..

A common authentication/authorization standard may be an impossible dream, but by laying out a standard for communications via tokens each site can issue its own secure tokens. The standard should define the process for handling and exchanging tokens.

Portable/Scalable Solutions Practice Group

Many gaps are common across the patterns, including:

- Allocations and accounts that span resources are a precursor to enabling portability.
- Interfaces and tooling common across resources and stable over long periods of time such that users feel confident enough to invest in adoption.
- A well-defined governance structure that spans resources and facilities that help to define and oversee standards and common policies and maintain a common roadmap.
- Common frameworks that can easily be adopted within a pattern area so that, for example, bringing in a new instrument takes minimal effort.
- Expertise both within technology areas as well as domains that can assist communities in adapting their workflows to effectively use integrated infrastructure.
- Schedule and resource abstractions that allow users to express what their workflows require and then enable those to schedule across distributed resources.

Unique challenges in this area revolve around the immediacy of results.

- Consequences of this are that resources must be available and potential latency effects must be considered. Having data potentially replicated at multiple resources could enable additional redundancy, but the time required for replication could restrict potential options.
- Standard queues often conflict with this requirement. This becomes even more challenging for workloads that may require both immediacy and scale.
- Frameworks or models that work for streaming data use cases, and not just bulk file movement models.

Resource Cooperations Practice Group

Facility cooperation requires a systematic ability to ask for reservations or resources in a timely manner. This requires automation of a service level agreement. Facilities that do not support time-sensitive patterns must adapt their queues and metrics to allow such requests. Resources must advertise their set of capabilities in a well-defined and commonly accepted vocabulary or interface. Through the steering structure, a facility may consider in advance the requests it may receive from other facilities to plan ahead. Finally, facilities should advertise failure models and objectives required for the facility to support time-sensitive patterns.

Facility cooperation also requires agreement between the facilities on standard interfaces and Software Development Kits so that users can build complex workflows without considering the idiosyncrasies of each facility. A steering group should facilitate this cooperation and standardization.

3. Where Are We Now? Current State of the Art in This Area.

- FPGA-based DAQs are ubiquitous today. This provides a starting point for custom protocols and interfaces to the compute backend that operate at 100 Gbps and beyond.
- Some experiments have incorporated real-time feature extraction, including the use of ML for feature identification. However, this is still the exception, rather than the norm.
- 10 Gbps DAQ systems are widespread, and some 100 Gbps DAQ systems are in production.
- PCIe gen4/5 based servers can stream 100 Gbps to SSD/DDR storage buffers, for software-based DAQ to compute data flows.
- ESnet6, now in production, can support dedicated optical circuits (i.e. L1), and user-driven (via APIs) dynamically provisioned (using the ESnet On-demand Secure Circuits and Advance Reservation System (OSCARS)) L2/L3VPNs with guaranteed bandwidth network connectivity (i.e., QoS, traffic engineering, etc.). Additionally, ESnet

has deployed FPGA devices throughout the ESnet6 footprint that can collect precision network telemetry information, providing unprecedented visibility at scale.

- LCLS-II runs ML at the edge, performs real-time data reduction, and streams data to NERSC using the Superfacility API to launch jobs remotely.
- Rubin Observatory, running the LSST, will discover thousands of transients every 30 seconds, with a requirement that these are vetted and distributed to several transient “brokers” before the next exposure is read out 30 seconds later. Current surveys, such as ZTF, vet candidates in 30 minutes to 1 hour and trigger follow-up shortly after that. This has allowed the discovery of several new types of short-lived transients as well as placed constraints on the physics of the progenitor systems due to early observations of several more well-known astrophysical transients. Pushing these limits to shorter time will open many new discoveries in this field.
- ML-based fast ptychographic reconstructions have been demonstrated by ANL and others. ANL is also collaborating with NVIDIA for more processing power at the edge.
- Tools for data management, analysis, and exploration (Cinema:Bandit and Cinema:Snap) as it becomes available have been developed by LANL and used for diamond anvil cell and shock compression experiments. Cinema:Snap uses a publish-subscribe approach.

4. Most Important/Urgent Gaps in Research, Technology, Resources, and Operations.

- Facility APIs, e.g., common libraries, common programming environments (e.g., containers) (Bjoern, Jana, Marcus)
- Efficient delivery of data directly to the appropriate compute resources(e.g., streaming into burst buffers, local compute may be insufficient and require remote compute)
 - Security concerns with direct ingest of data in HPC
- Services that enable easy interaction with facilities (federated facilities), e.g., light sources and supercomputers.
 - Common Authentication Infrastructures and Frameworks: access to compute resources, data transfers (e.g., Globus Auth, FederatedID)
 - Services: data transfers (Globus transfer, rsync) and remote function/procedure calls (funcX)
- Accessibility and sharing of data across facilities and user groups

- Common or interoperable security frameworks (e.g., AuthN) to enable cross facility APIs and data
- Streaming (i.e., memory-to-memory) of data to support time-sensitive workflows. Current infrastructures rarely provide memory-to-memory data movement capabilities. Domain scientists typically implement their solutions on workstation and/or commodity compute resources. These solutions, out of box, cannot efficiently use the resources or features provided by high-end compute resources or supercomputers, such as network or hardware topology for communication or heterogeneous resources (e.g., CPU+GPU, tensor cores and other HW-specific compute components). Optimized runtime systems, libraries, and tools are needed to rapidly develop new algorithms and efficiently use supercomputer-scale resources, e.g., an image processing runtime system for supercomputers.

5. Near-Term and Longer-Term Opportunities?

The near-term and longer-term opportunities would be to address the gaps mentioned previously.

Appendix H

IRI ABA BLUEPRINT: DATA INTEGRATION-INTENSIVE PATTERNS

IRI ABA Blueprint: Data Integration–Intensive Patterns

Definition of This Pattern

Data integration–intensive Integrated Research Infrastructure (IRI) patterns emphasize analysis of combined data from multiple sources, for instance, data from multiple sites, experiments, and/or simulations. Tracking metadata and provenance for reproducible science and interactive analysis of data possibly at scale are also key features.

1. Representative Cases for This IRI Pattern

Integration of Data from Simulations and Experiments / Observations

The integration of data from simulations and experiments/observations is needed to generate new insights and direct subsequent actions. Data from simulations are used to manipulate parameters that steer an experiment's settings or those of an observational device, and to compare real data to models as constraints. The results of experiments or observations are used to refine the parameter space of simulations and validate simulation results. A feedback loop, possibly in real time, combines both steps to accelerate and automate the discovery process. This pattern overlaps with time-sensitive and long-term campaign patterns.

Examples of data integration–intensive patterns that involve the integration of data from simulations and experiments/observations include:

- Large-scale simulations integrating experimental data from DOE and international nuclear physics facilities (High Energy Physics (HEP), Nuclear Physics (NP)).
- AI/ML incorporated into simulations to drive exploration of parameter space using codes like LAMMPS in the workflow (HEP, NP).
- Combining multimodal data from simulations and experiments for molecular-scale imaging workflows, including instrument, compute, and storage (Biological and Environmental Research (BER) Environmental Molecular Sciences Laboratory (EMSL) and the Joint Genome Institute (JGI)).
- DIII-D Fusion experiments coupled with simulations where near real-time analysis is required to determine/predict plasma conditions (Fusion Energy Sciences (FES)).

- Autonomous materials characterization workflows using data (hundreds of GBs per day) collected at light sources, neutron sources, and Nanoscale Science Research Centers (NSRCs) coupled with molecular dynamics (MD) simulations (Basic Energy Sciences (BES)).
- Coupled large-scale experiments and simulations involving multiphase, chemical reacting flows, such as computational fluid dynamics, for carbon capture and conversion technologies. This may be applied to direct air capture and point source capture, novel reactors that combine carbon capture and conversion to new chemicals, new biofuels applications (biomass, MSW) for H₂ production, novel sorbent-based oxygen separations systems, and process systems design and optimization (Fossil Energy and Carbon Management (FECM)).

Cross-Site Data-Driven Discovery

Data-driven discoveries, possibly ML-driven, relies heavily on new or existing data generated at different facilities/sites. This includes using similar, multimodal, or heterogeneous data already generated at different facilities, or running the same tool, e.g., simulation software, on different systems, or experimental/observational data originating at different sources, the results of which must be combined, processed, and analyzed. This pattern overlaps with the long-term campaign pattern.

Examples of data integration–intensive patterns that involve cross-site data-driven discovery include:

- Fusion workflows where large data on the order of hundreds of TBs are stored in custom SQL and noSQL databases, and data needs to be exposed for computing on local and institutional systems (FES).
- Analysis of multimodal materials data collected and stored across light sources, neutron sources, and NSRCs. Data must be moved to one or more common sites for data processing. The data must already be of a similar format or transformed into a common format for data-processing tools. Similarly, any metadata associated with the raw data must be available and organized in a similar fashion. Depending on the computing system, data-processing software may have to run on different underlying platforms and architectures. The orchestration of data movement across different systems should be coordinated so that the various access mechanisms for different sites are hidden from higher-level applications (BES).

2. What Are the Principal End-to-End Performance Factors, Technical and Operational Components, Interfaces, and User Experience Requirements?

Resource Co-operations

Data processing and storage are frequently needed across facilities. To ensure cooperating facilities support data integration-intensive patterns, each facility should set up shared and common mechanisms to expose information about the facility. These mechanisms, which require integration with facility operations, include data movement, scheduling and launching jobs, and monitoring. Across facilities, today's data transfers take place as bulk-file transfers, but tomorrow's integrated research infrastructure should offer a common method to access data as well as Findable, Accessible, Interoperable, and Reusable (FAIR) metadata resources. Mechanisms should be in place to enable access to large (many PBs) data sets remotely or to relocate computing tasks closer to the data. Streaming interfaces and file-based/memory-based interfaces will be required for integrating simulations with experimental data in real time. There is a need to acquire compute resources at scale with minimal wait. Above some request threshold, compute resource reservations may be tolerable. Allocation policies should be adapted/modified to better accommodate the use of distributed computing resources, and to enable cross-facility data sharing. All facility services should use common Application Programming Interfaces (APIs).

Cybersecurity and Federated Access

New common complex-wide federated access and cybersecurity models and systems must be developed and employed. Workflows and agents must be trusted at the same level as humans. Access must be persistent across all resources; logins should not be re-required as workflows access different resources. Policies must be modified/adopted to reflect the new landscape, and conversely technology must be able to accommodate various security models. Alignment of policies across sites, access monitoring, various permissions due to the Health Insurance Portability and Accountability Act (HIPAA), and international landscapes, should all be considered. Cybersecurity and federated access systems should use a common AP.

User Experience

In an ideal world where IRI is wildly successful, a human would be able to:

- Access storage resources without significant contention from other users.
- Easily move data throughout the storage hierarchy for analysis.
- Easily assemble all needed data sets in one location (transfer, reference, etc.).

- Avoid logging into multiple systems (because everything is integrated).
- Transfer files at a predictable speed.
- Shift seamlessly from one HPC site to another if a resource fails.
- Integrate and explore data interactively through Jupyter, RStudio, or the command line.

Data-intensive work has an emphasis on storage, I/O, and interactivity. The main barriers to productivity are distribution and facility policies that prioritize HPC. User training and documentation are paramount.

In an ideal world where IRI is wildly successful, a robot would be able to:

- Analyze data that has been transferred from other sites.
- Move data from experiment to compute with persistent credentials.
- Access resources using APIs.

The main difference between robots and humans is the need for programmatic interfaces that can determine the characteristics of the computing facility and availability of storage resources and move data quickly. Robots need autonomous access with long-lived credentials.

Workflows, Interfaces, and Automation

An open standard for workflows that allows interoperability and reusability, combined with a set of reference implementations, is required to facilitate this pattern. This standard should define a base abstraction API for workflows and automation and include monitoring capabilities.

Scientific Data Lifecycle

The scientific data lifecycle touches on all aspects of the scientific process, from conception of the scientific question, to theory/modeling/simulations, to testing with experimental and observational data, to data processing/reduction, to data analysis interpretation, to publication and sharing. Tools and services are required for data and metadata and provenance tracking. Provenance and metadata should include all scientifically relevant information, including information about the experiment/observation/simulation, as well as transformations applied to it. There should be common mechanisms for searching through data, metadata, and provenance. There should be common DOI minting capabilities. Data should be organized as FAIR data as soon as possible. Software should be treated the same way as data. A common (potentially distributed) facility for DOE data should be established.

Portable / Scalable Solutions

Software must run across multiple platforms and architectures, including edge, local, campus, cloud, and national. Tools to ease portability/scalability across these systems will be required to realize transparent IRI.

3. Where Are We Now? Current State of the Art in This Area.

Current use cases that take advantage of IRI (of those that exist) are primarily bespoke / ad hoc for a particular science problem, user community, or collaboration. General guidance and tools for realizing this pattern do not exist.

Resources and the ability to use resources are laboratory, facility, and sometimes instrument/system dependent. Each laboratory, facility, and sometimes instrument/system has its own way of exposing access and interacting; APIs are rare, and common APIs do not exist or are not used. Large-scale computing facilities do not accommodate streaming data well. There are, however, success stories, including the Open Science Grid (OSG), the European Open Science Cloud (EOSC), and earlier grid work that may point the way forward.

Federated access is nascent: technologies and tools exist but are not widely applied within our landscape. Federated access is site specific, and MFA mechanisms vary.

A plethora of workflow systems and tools exist. Robust and reliable file-based data movement and tracking tools exist (e.g., Globus and Rucio). The lack of standardization makes finding common ground difficult. Streaming-based workflows are a gap.

Access to remote data is largely by either straight file copy / data movement or remote file-system mount. Neither makes optimal use of network and storage resources. More efficient and scalable solutions based on advanced data access management, on-demand data movement, and cache technologies have seen some experimental uses but not been widely and reliably deployed.

Resources (allocation from DOE) must be managed separately at different sites. Each site has its own means of exposing access to resources; center APIs are rare.

4. Most Important/Urgent Gaps in Research, Technology, Resources, and Operations?

Cross-facility APIs need to be developed and maintained to bridge the gaps in interoperability required by this pattern. Cybersecurity and federated access models require new modes of access, including uniform authorization, MFA, and policies to realize this pattern. Common abstracted facility APIs and workflow and automation tools are needed for this pattern. A complex-wide data facility and searching and cataloging tools are required. Policies and mechanisms to allocate resources when needed and in a uniform manner are needed to further science represented by this pattern. There is a gap in cross-training of operations staff to have a working knowledge of computational topics and trends. For example, scientific staff often have little computing knowledge, and having legal, intellectual property, and tech transfer staff have a

working knowledge of AI basics along with definitions could help facilitate review/approval/governance of software and tooling for labs.

5. Near-Term and Longer-Term Opportunities?

Many early-win science opportunities exist in this area are derived from previous examples.

- Adopting common APIs for facilities.
- Establishing a standard set of metadata: high level to know what is in a data set and where it is, and lower-level metadata for a domain.
- Streaming data to/from compute and storage facilities.
- Supporting edge devices and edge computing paradigms. Software and infrastructure build out to take advantage of processing and computational compute resources at edge devices.
- Exploring decentralized data storage to build a complex-wide data storage capability for IRI.
- Clarifying data policies; evaluate frameworks in place that allow co-design, and users can shape/form what is required of IRI infrastructure.
- Exploring opportunities to catalog known data standards/ontologies—lessons learned from NASA and others using commercial clouds or open data platforms offered by commercial cloud providers; establish protocols with vendors.

Appendix I

IRI ABA BLUEPRINT: LONG-TERM CAMPAIGN PATTERNS

IRI ABA Blueprint: Long-Term Campaign Patterns

Definition of This Pattern

This class of Integrated Research Infrastructure (IRI) patterns is characterized by a need for sustained access to resources at scale over a longer time needed to accomplish a well-defined objective. Long-term robustness, reproducibility, and reliability are important to accomplish, which would potentially involve significant logistical planning. Examples include sustained simulation production and large data (re)processing for collaborative use.

We refine the above definition further so that a long-term research campaign for IRI purposes is one that conducts scientific research:

- **Lasting for 5–30 years.**
- **Using at least some resources that it does not own or control, e.g.:**
 - Instruments.
 - Computing resources.
 - Data resources.
 - Networking.
- **Relying on one or more user facilities for those resources to a varied research community.**

We will look at this through the lens of:

- **Research campaigns** that need resources to conduct research.
- **User facilities** charged with providing those resources, including computational capability, light sources, beam lines, atmospheric data, characterization of materials or biological systems, and more.

Long-term campaigns need to continuously produce valuable science while:

- They evolve (or do not!).
- The user facilities they rely on evolve (or do not!).

We expect that long-term campaigns will have all the challenges of the other two pattern groups, **plus** challenges associated with changes over the duration of the campaign in terms of at least:

- Programs
- People
- Priorities
- Policy
- Infrastructure

Over a 5- to 30-year duration, it is almost certain that the campaign will experience:

- Shifting priorities in funding agencies, facilities, and participating institutions.
- Turnover in workforce as people progress through their careers.

- Architectural changes in computing systems and networking.
- Multiple changes to all levels of software stack(s).
- Changes in quantity and quality of the evolving dataset.
- Changes in network technology and protocols.

The interface between campaigns and facilities becomes more complex if the campaign engages with and relies on multiple user facilities that are under separate administrative control. This will be compounded if the facilities adopt changes at different rates, as they are likely to do. The more user facilities involved, the more the campaign is likely to be affected by change over time. These changes incur costs to campaigns, especially if a campaign also finds it necessary to “move in” to a (new) facility or “move out” of (an existing) one. The IRI must be able to mitigate this kind of change for the campaigns and facilities.

1. Representative Cases for This IRI Pattern

These use cases are drawn from earlier IRI phases and supplemented with input from the Long-Term Campaigns pattern group. Note the use cases identified here are likely to be associated with other patterns (for instance, some science workflows are time sensitive *and* data-integration intensive *and* long term):

- Large-scale simulations integrating experimental data from many sources.
- Analysis of large data sets from observational/experimental science.
- Data streaming from experiment to High Performance Computing (HPC) facilities to be used to guide/target said experiment on a deadline or schedule.
- Large-scale data to be transferred, curated, and served long term.
- "Bursty" experiments that provide large quantities of data in hours/days with experiment downtime in between.
- Small data transfers from swarms of wireless devices and field sensors to user facilities.
- Reprocessing of legacy data sets to gain insight based on new theory or understanding of the instrument/experiment.
- “Moving in” to a facility, transferring in volumes of data, learning the supported architectures, workflows, and support mechanisms.
- “Moving out” of a facility, probably having to locate a final or temporary resting place for data, while gaining access to new resources if the campaign is to continue.

Long-term research campaigns will need many different resources, some of which will be subject to (possibly unwanted) change over the course of the campaign. The longer the campaign, the more these resources are likely to change, for example:

- Allocations:
 - Compute cycles.
 - Working space for data (terascale to exascale now).
 - Long- to indefinite-term storage space (petascale to exascale now).
- ID management/accounts.
- Data/archives/publishing support.
- Applications/scheduling/workflows.

- Engagement/partnerships.
- Data-transfer capability.
- User facility mission(s).
- Consistent, long-term underlying transport protocol support.

An archetypal use case is a decades-long physics project that would:

- Conduct a computational campaign to obtain theory predictions that inform the design of experimental facilities and interpret experimental data.
- Operate the machine to take data while conducting more computation to:
 - Process and analyze data produced by the experiment.
 - Calculate and retain calibration data to inform event reconstruction.
 - Produce simulated data of the same type as the machine for comparative purposes.
 - Perform more theoretical calculations.
- Upgrade the machine when possible to increase one or more of the:
 - Sample rate.
 - Resolving power.
 - Sensitivity.
- Obtain increased compute and/or data capacity needed by the upgraded machine.
- Eventually decommission the machine.
- Dispose of the data appropriately, whether archival, deletion, long-term curation, or other.

2. What Are the Principal End-to-End Performance Factors, Technical and Operational Components, Interfaces, and User Experience Requirements?

The principal measure of performance in the relationship between research campaigns and user facilities is **whether the facilities enable success of the research campaigns continuously over the campaigns' planned lifetimes.**

They must jointly manage change successfully over time.

What is likely to change in the user experience over the duration of a campaign?

- Compute hardware
- Compute software
 - Workflows
 - Interfaces
 - APIs
- Networking technologies (from interconnects to WAN)
 - Scalable, supportable protocol stacks
 - Migration to future networking protocols; maintenance of older protocols; protocol abstraction

- Instruments, in terms of:
 - Number
 - Sensors
 - Resolution/sampling
 - Data-collection/generation rates
- People, as long-term campaigns can span many generations of scientists moving in and out of the collaboration

The following sections focus on the factors, components, interfaces, and requirements that cross-cut across the representative cases described previously.

Resource Co-Operations

- Research campaigns do not have their resource allocations unilaterally disrupted by a facility's proposal/evaluation/allocation cycle.
- Campaigns continue to have the compute/data/network resources they need as instrument upgrades and changes to experimental techniques drive those needs higher.
- Experiments have the compute/data/network resources they need when they need them for experiment steering or for post-processing.
- User facilities deliver the uptime and availability requirements of the research campaigns while performing maintenance and upgrades.

Resource Co-operations Practice Group Contribution

These campaign patterns require allocations that straddle multiple accounting periods for a facility. Thus facilities whose metrics may consider particular time periods (e.g., an annual reporting interval) will need to agree on goals and metrics with other facilities that may have data-collection campaigns that span a different interval (e.g., multiple years). The two or more co-operating facilities will need to agree to support campaigns that operate across their resources over potentially new and unaligned time windows. This support will take the form of

- (i) reserving resources in a structured manner to allow for service levels that are acceptable to the end-to-end campaign,
- (ii) modifying agreements (to elevate or lower priorities) according to the needs of competing campaigns,
- (iii) creating a joint proposal acceptance mechanism, so that campaigns can request resources for longer (and, potentially, variable) durations of resources across the facilities,
- (iv) creating failure recovery mechanisms and agreements, and
- (v) deciding on joint metrics and mechanisms for credit for success. Once the multilateral steering structure is established as noted previously, the steps will enable the support of long-term campaign patterns.

Cybersecurity and Federated Access

- Conduct of research is not disrupted as security policy and practices evolve over time.
- Security fixes are applied in a timely manner and do not disrupt workflow.

- Users do not have to re-authenticate multiple times per day in multiple different ways and at multiple different user facilities in the distributed infrastructure.
- The security policy allows end-to-end performance testing across sites.
- Access to facilities is maintained for as long as data are retained rather than the term of a compute allocation; users who need access to data in the future might be different from those who created the data (or were on the original compute allocation).
- Long-running compute jobs, services, or other processes running on behalf of a campaign are decoupled from the identity and credentials of any individual user and can maintain their own credentials without human intervention.

Cybersecurity and Federated Access Practice Group

We expect that there may be very different cyber implementations across these patterns, driven by unique needs. For example, memory to memory transfers across facilities require substantial changes to the existing security architectures of computing facilities, but are needed for certain kinds of low-latency integration. However, the goal would be that the design principles and models developed here could accommodate the diverse implementations across the patterns. The specific controls needed for different architectures need in-depth exploration.

User Experience

- Users are not hindered by having to unnecessarily work at multiple user facilities, nor by being made to move unnecessarily from one facility to another.
- Computing architectural changes do not drive changes to software that campaigns do not expect and need to respond to.
- Computing architectures and research needs stay aligned and campaigns are not pressed to use unsuitable architectures.
- Users have a single place to go to have their needs met.
- Campaigns have the consulting and support services they need from the user facilities they work with.

User Experience Practice Group

In an ideal world where IRI is wildly successful, a **human** would be able to:

- Authenticate a session for the entire IRI in a single dialog.
- View all data associated with the campaign through a single portal, regardless of data location.
- Set data access policies, including release policy on all data they own from one system.
- Search and retrieve data from the campaign via associated metadata.
- Have confidence that the tools/processes will be available and well supported for the duration of the campaign.
- Access the needed resources for the full duration of the campaign.
- Access the data and metadata beyond the lifetime of the campaign.

User productivity is driven by data organization and management for long-term campaigns. Consistent access to resources during the life of the campaign makes it easier for users to reproduce analyses.

In an ideal world where IRI is wildly successful, a robot would be able to:

- Access systems with the same credentials for a long time.
- Leverage API-based access.

The main difference between robots and humans is the need for programmatic interfaces that the robot can use to discover the characteristics of the computing facility and availability of storage resources and move data quickly. Robots need autonomous access with appropriate authorization and long-lived (or renewable) credentials. Robotic access will be essential to automated experiment management, triggered when data are generated, when analysis results are produced, or when an intelligent agent suggests/imposes experiment changes.

Especially during long-term campaigns with generations of scientists moving in and out of projects, users will need to be trained in how to access systems and the historical data that reside on them. Humans can infer what they are looking at, but there is a lack of consistency in data formats and lack of descriptions of them across different domains. That is a hindrance to anyone who might want to use the data if they were not part of its creation. This is especially problematic due to the likelihood of missed opportunities for AI/ML training sets.

Workflows, Interfaces, and Automation

- Workflows, interfaces, and automation undergo planned change over 5- to 30-year time frames without disrupting the conduct of research.
- As software stacks become deeper and more complex, workflows become simpler, and troubleshooting does not become more complex.
- Workflows do not needlessly cross facility boundaries, or seamlessly cross facility boundaries when needed.
- Infrastructure that supports workflows is managed intentionally so that quick fixes (especially security fixes) do not create technical debt.
- Campaigns can select and analyze outputs across a user's/project's workflow pieces via workflow interface (API-based or web-based).

Workflows, Interfaces, and Automation Practice Group

The most important gap is a standard for interoperable workflows. As technologies (both hardware and software) evolve over the 5- to 30-year time frame, a long-term, stable standard for interfaces will enable workflows to adapt to these changes with minimal disruption. The end points need to be identified, via a reference architecture, so that job/data services and resource descriptions can be standardized.

Monitoring should be considered a priority. Long-term campaigns often have simulation or analysis campaigns that last for months or years, which need to run and respond to resource

availability without human intervention. The interface standard needs to provide resource status and performance information, in a format that a long-term campaign workflow can react to.

Documentation should also be considered from the beginning, for example by planning for a self-documenting workflow.

A common authentication/authorization standard is required. By laying out a standard for communications via tokens, each site can issue its own secure tokens. The standard should define the process for handling and exchanging tokens.

Scientific Data Life Cycle

- Data storage costs are successfully managed as technology and storage requirements evolve over time.
- Data are retained and curated appropriately while data generation and retention rates increase over time.
- The lifetime requirements of datasets are specified to user facilities up front.
- Data and metadata are documented and described sufficiently to support reproducibility of scientific results and re-analysis or use as training sets.

Portable / Scalable Solutions

- Campaigns can successfully migrate to new computing architectures, enabled by common interfaces in use at different facilities.
- Campaigns have access to the computing architectures (or abstraction layers representing those architectures) best suited to their needs.
- Campaigns continue to have the compute/data/network resources they need as instrument upgrades and changes to experimental techniques drive those needs higher.
- Research campaigns have the funding and personnel available to adapt to technology, architecture, and scale changes over time.
- Networking technology changes continue to bring performance improvements without disrupting the conduct of research.
- The mix of wide-area networking requirements (especially related to performance) is well-characterized and projected to future needs.
- Applications are reused across facilities as an aid to reducing software development and portability burdens.

Portable/Scalable Solutions Practice Group Contribution

Many gaps are common across the patterns, including:

- Allocations and accounts that span resources are a precursor to enabling portability.
- Interfaces and tooling common across resources and stable over long periods of time such that users feel confident enough to invest in adoption.
- A well-defined governance structure that spans resources and facilities, which will help to define and oversee standards, common policies, and maintain a common roadmap.

- Common frameworks that can easily be adopted within a pattern area so that, for example, bringing in a new instrument takes minimal effort.
- Expertise both within technology areas as well as domains that can assist communities in adapting their workflows to effectively use integrated infrastructure.
- Schedule and resource abstractions that allow users to express what their workflows require and then enable those to schedule across distributed resources.

Since these campaigns can last years or even decades, longer-term allocations are critical. As projects plan out campaigns, they need to know that access to resources can be counted on over the lifetime of the project.

The interfaces need to be not only common across resources but stable over extended periods of time. Reproducibility is also a challenge, as various resources will likely have to evolve over the lifetime especially for very long campaigns. So balancing the need for systems and services to be updated and evolve over time with the need for reproducibility is a challenge.

3. Where Are We Now? State of the Art in This Area.

In scientific domains where (some) data are kept “forever,” burden is placed on the campaigns and user facilities to have aligned expectations about what data are kept. Not only must the data be stored, but time and resource cost must be incurred if the data are to be pulled out of long-term storage for further analysis.

Wide-area networking providers (ESNet and other GNA-G partners) can dynamically provision bandwidth for traffic, including deadline-driven data transfers, but this is not often taken advantage of.

Physics and computational chemistry are two domains that have a long history and established use patterns of computational and data resources. Biology-driven computing is still nascent and is likely to drive very different use cases and more computational diversity.

Observational and experimental sciences will continue to push the boundaries of available computational and data resources as they strive to increase the quality and quantity of the data they produce.

It has been suggested that zettascale computing/data will come in ~2035 and that data will be the center of gravity for those architectures. Compute would be “in” the data rather than the current practice of moving data to and from the compute.

4. Most Important/Urgent Gaps in Research, in Technology, in Resources, in Operations?

- Dearth of permanent data storage resources that span all DOE Office of Science domains
- Data policy implementations vary from facility to facility and campaign to campaign
- There is no capability to curate large observational data sets that could be re-analyzed as new theory emerges
- Resources (especially compute time) are not guaranteed for the lifetime of a campaign
- The capacity of observational and experimental science to generate data exceeds the compute/data/network capacity to make use of all of the data
- There is no capability to handle multiple streams of small data transfers from a swarm of wireless devices
- It is difficult to test performance end-to-end because of site security policies
- There is no funding for a dedicated workforce to port/develop/update software and maintain data over the entire life cycle of research campaigns
- Software development/update/porting activities are vital, yet frequently underplanned and underfunded
- The validation of software stacks across systems and user facilities requires time and effort
- Program offices do not make it a priority to cooperatively align their resources for campaigns
- There are no standards for various IRI areas to enable long-term engagement and encourage projects to buy in

5. Near-Term and Longer-Term Opportunities?

- Establish permanent data storage to be used by all program offices that endures beyond any research campaign and is decoupled from the campaign lifetime. Ensure it has adequate computation for reprocessing (this may be a good commercial cloud use case).
- Align allocation cycles and durations with the life cycles of longer-term campaigns to assure continuity of resources over long time frames.
- Ensure that computational and data processing and network capacity is commensurate with the data production rates of observational and experimental facilities across the program office and remains so over time
- Provide guaranteed funding for software porting/development/update activities and for maintaining data over the entire lifecycle of the campaign.
- Establish robust practices to maintain continuity of the workforce: have robust workforce training, transfer authorities and permissions, provide data descriptions (or self-describing data) to maintain capability as individuals move on in their careers

- Ensure that data are accessible in terms of being described well or being self-describing. Create documentation, common language, and lexicon for scientific data and metadata in the IRI.
- Create a generalized DOE facility interface (or API if you will) with layers of abstraction to allow for different resources (compute resources, schedulers, storage systems, network protocols, and so on) at the facilities while presenting a common interface to all users. Monitoring will be highly important to provide a window into current activity.
- Establish a holistic IRI approach including coherent *policies* and *practices* on top of the technical underpinnings (compute and instrument infrastructure, APIs, abstraction layers, security policy, workflows, etc.).

Appendix J

IRI ABA DESIGN PHASE: PRACTICE GROUP ON RESOURCE CO-OPERATIONS

IRI ABA Design Phase: Practice Group on Resource Co-Operations

Members

Arjun Shankar (Facilitator), shankarm@ornl.gov
Johannes Blaschke (Facilitator), jpbblaschke@lbl.gov
David Matin, dem@alcf.anl.gov
Xi Yang, xiyang@es.net
Sarp Oral, oralhs@ornl.gov
Tom Lehman, tlehman@es.net
Lee Ann McCue, leeann.mccue@pnnl.gov
Taylor Childers, jchilders@anl.gov
Scott Atchley, atchleyes@ornl.gov

Description

<i>Resource Co-Operations</i>	<i>Allocations/provisioning of multiple heterogeneous resources must be aligned in time and planned in advance to enable integrated workflows. IRI requires new levels of cooperation, collaboration, co-scheduling, and joint planning across facilities and across DOE programs.</i>
-------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

This practice area aims to describe the essential requirements and approaches to create an integrated research infrastructure across DOE by thinking of facilities and their capabilities “co-operating” and cooperating to offer DOE facilities as a whole to scientists and end users. The responses here discuss what it will take to enable this overarching objective, including mutually beneficial multilateral steering bodies to service establishment modes to set up cooperating services for urgent science, data integration science, and long-term campaigns.

1. Vision/Goal Statement

Establish a DOE-wide virtual infrastructure for resource use with a set of common operational policies towards improving scientific progress for the three major scientific workflow patterns (time-sensitive patterns, data-integration-intensive patterns, and long-term campaign patterns), without impeding independent mission requirements for the participating DOE user facilities.

Scientific users should have access with a common or federated cross-facility account/identity authentication and authorization mechanism to the shared resources. Users also should access the resources (i.e., experimental, compute, network, and data) through an interface allowing them to compose and schedule workloads and scale them up and down as needed. The interface should (1) be intuitive enough that the users easily grasp what shared resources are expected to be available at any given time and what common policies govern the said shared

resources; (2) be simple enough to shield the users from different resource characteristics, independent policies, and operational procedures/practices of the participating DOE user facilities; (3) be flexible enough for the users who want to interact with it at a lower level to increase efficiency or performance of their scientific workflow patterns.

There is a natural tension between local control and global scheduling. Mechanisms should be developed so users can negotiate with facilities for access to the needed resources. This should be a multi-way structured communication between the facilities and the user, so that the commitments from each are clear. A similar mechanism will also be needed to resolve concurrent competing requests from users (e.g., two experiments requesting more than 50% of a given resource at the same time). This could be approached by explicitly defining policy and mechanisms developed to support the objective of co-operating facilities. Shared governance would facilitate communication and the construction of coordination rules and engagement structures.

The policies and operational practices of DOE user facilities have some commonalities but at the same time differ widely (e.g., maximum wall times for queues, capacity management/data purge policies, set priorities for different scales of jobs). To steer the community towards establishing a common set of policies across the participating DOE user facilities, a high-level advisory organization needs to be established. This advisory organization should have members from the DOE, program managers, and DOE user facilities, as well as representatives of the scientific end users, domains, and the workflows that will benefit from the shared resources. This advisory organization should help define the common set of policies (https://www.worldscientific.com/doi/10.1142/9789811204579_0018) that will be supported by all participating DOE user facilities and also the interfaces to access them. There also needs to be a low-level technical organization established for the execution, communication, and shared responsibility of the operational aspects of the (e.g., system management, security) shared resources. This low-level technical organization should have operational representatives from the participating DOE user facilities, and mechanisms (potentially with a third neutral or DOE-established entity) for tool building and continuing support.

2. What Is the Current State of the Art, Including State of the Community, Related to IRI?

Experimental and observational science (EOS) teams and interinstitutional networking projects are leaders in this space. Noteworthy is that the science communities (such as photon sources, biochemistry and bioinformatics, telescopes, high-energy particle physics, and fusion energy science) have done a lot of the heavy lifting themselves. Examples are the Energy Sciences Network (ESNet), Conseil Européen pour la Recherche Nucléaire (CERN), and National Energy Research Scientific Computing Center (NERSC)'s superfacility collaborations (doi.org/10.48550/arXiv.2206.11992). The state-of-the-art model that emerges from these collaborations is (a) persistent storage used to orchestrate the workflow beyond the lifetime of individual jobs; (b) persistent services to support workflows and control capable of monitoring

multiple sites and requesting resources; (c) a portable software deployment; (d) processes capable of requesting resources; and (e) a reliable high-speed network between facilities.

A common design pattern that we observe ([10.1109/BigData52589.2021.9671421](https://doi.org/10.1109/BigData52589.2021.9671421)) is that (a) and part of (b) are hosted on microservice platforms. When hosted at the edge of the High-Performance Computing (HPC) center's network, these microservice platforms are capable of bridging the gap between HPC resources (such as job submission, high-speed networks, and file systems) and EOS data sources. Furthermore, microservice platforms enable relatively high up-time (by providing workload schedulers, such as Kubernetes, which restart failed or failing services, and in some cases by drawing on redundant networking and power to microservice clusters), ensuring that orchestration services between institutions remain available. Software portability (c) is frequently accomplished using a container runtime (e.g., Shifter, Singularity, and Podman). As there is little standardization among HPC centers, implementation details, especially around (b) and (d), vary drastically. The CERN collaboration's Worldwide Large Hadron Collider (LHC) Computing Grid (WLCG) is a pioneer in developing a distributed workflow engine. Recently, HPC centers have begun to adopt more general-purpose workflow platforms, with Argonne Leadership Computing Facility (ALCF)'s Balsam workflow platform being an example of a more general-purpose cross-facility workflow platform. In general, however, each science community has adopted its own tailored workflow engine (e.g., [10.1109/UrgentHPC54802.2021.00011](https://doi.org/10.1109/UrgentHPC54802.2021.00011)) Wherever possible, Application Programming Interfaces (APIs) are used to programmatically monitor workflow steps across facilities, and request resources. Most of these workflows use the high-speed network via bulk file transfers with tools like CernVM-File System (CVMFS), XRootD, and Globus being prevalent. Some workflows copy data directly to memory (via tools like ADIOS2); however, these examples are relatively rare. A common concern that has arisen from the use of edge services and a broader adoption of APIs is that of security. For example, an edge service can expose file systems previously only accessible from within a secured network, and APIs commonly require long-lived access tokens, which can be stolen.

In addition to the EOS design patterns discussed previously, which usually serve to tackle a narrow scientific problem, ESnet is a multifacility organization that supports all DOE Office of Science programs. ESnet, therefore, is a leader in orchestrating resources across many organizations. ESnet has a focused mission to provide state-of-the-art networking to support scientific research. Maintaining a leadership-class facility in this area requires ongoing network systems development and periodic equipment upgrades. The fundamental building block is providing fully managed high-speed connectivity between all the key scientific facilities in the DOE Office of Science research portfolio. This includes close coordination with connecting sites to provide the level of capacity, service level guarantees, and network state visibility as needed for their science objectives. The current state of the art includes continuous upgrade planning to incorporate the latest network speeds and technology, as well as building network management and monitoring systems tailored to the needs of the scientific process.

One key aspect of advanced Research and Education (R&E) networks today is a focus on the aggregate use by all the science workflows. The network monitoring and management systems are highly tuned to evaluate the network facility in the context of this aggregate use. In addition

to this aspect, ESnet has long been a leader in providing services that can be tailored to individual science workflows, sites, or even specific data flows. These types of services are enabled by providing API-driven services, where users can reserve deterministic services and interact with the network. ESnet is also highly active in research and development into these more workflow-specific services. This will likely be an important component of future science workflows, which increasingly depend on the real-time access and use of distributed sets of resources. ESnet On-Demand Secure Circuits and Advance Reservation System (OSCARS) and Software-Defined Network for End-to-end Networked Science at Exascale (SENSE) systems are examples of workflow-tailored services with multi-resource orchestration.

ESnet provides a full set of network services, which includes high-performance IP-based routed connectivity between sites, extensive network monitoring with portal access, and advanced workflow-specific services. These advanced services can be tailored to individual science workflows, sites, or even specific data flows. These are enabled by providing API-driven services, where users can reserve deterministic services and guaranteed bandwidth, and interact with the network. ESnet connects the 28 DOE Office of Science facilities and some non-DOE sites where there is a science collaboration with a DOE facility. ESnet has a goal of 99.9% uptime. There are no formal service-level agreements between ESnet and DOE sites.

3. What Are the Most Important/Urgent Gaps in Research, Technology, Resources, Operations, Policy, and Environmental Constraints for This Practice Area?

The primary gap in setting up resource co-operations across facilities is the absence of established protocols and policies to set up end-to-end campaigns. The resource profiles at each participating facility are seldom aligned in a way that a science campaign can be overlaid seamlessly across the facility. Thus, the process of co-operation requires first an understanding of capabilities across facilities, an agreement for how capabilities will be aligned, and an establishment protocol of resource use (across time and capacity) to enable end-to-end campaigns. In addition, during failures in a campaign, or when priorities change for a particular facility, the different facilities need a unified governance model to ensure that campaigns can be jointly reprioritized in a manner agreeable to the scientists as well as the facility objectives. We use the term “steering” in lieu of governance to strive for decentralized efficiencies in each facility and to also imply advisory roles. To effect this structure of co-operations, we expect facilities to join and be advised by bilateral and multilateral steering and/or groups.

Specific challenges include a lack of a common service establishment protocol (as seen, e.g., in networking systems) to create a cooperating set of facilities, and a methodology (including a vocabulary and ontology) to align resources availability (in time and size) across facilities. The rules for steering bodies to oversee objectives, constraints, and metrics for individual and collective facilities also need to be defined. Choosing appropriate messaging is an important

way that all stakeholders will feel heard. For example, whether to describe this steering group as an advisory group vs a governance organization will send a very different message to members. The former sends the message that this body has little authority to adjudicate, whereas the latter invokes a picture of “top-down” control. We therefore recommend a middle ground: a steering group, indicating a largely democratic organization with authority to resolve conflicts.

IRI Pattern-Specific Requirements

For the specific patterns in the IRI, we identify the following specific requirements.

Time-Sensitive Patterns

Facility co-operation requires a systematic ability to “ask for reservations or resources in a timely manner.” This requires automation of a service-level agreement. Facilities that do not support time-sensitive patterns must adapt their queues and metrics to allow such requests. Resources must advertise their set of capabilities in a well-defined and commonly accepted vocabulary or interface. Through the steering structure, a facility may consider in advance the requests they may receive from other facilities to plan ahead. Finally, they should advertise failure models and objectives required for the facility to support time-sensitive patterns.

Data Integration–Intensive Patterns

Data integration is an immediate requirement for resource co-operation. Data processing and storage are frequently needed across facilities. To ensure co-operating facilities support data integration–intensive patterns, each facility should set up mechanisms to exchange (or alternatively, access) data and metadata information. Across facilities, today’s data transfers take place as bulk file transfers, but tomorrow’s integrated research infrastructure should offer a common method to access data as well as metadata resources. State of the practice Globus and Science DMZ structures can enable data movement at first, but a plan for scalable repositories and FAIR (Findable Accessible Interoperable Reusable) interfaces to data will be required.

Data provenance and repositories with associated computing and analysis resources will be needed to exchange facility data for campaigns that move from one facility to another.

Long-Term Campaign Patterns

These campaign patterns require allocations that straddle multiple accounting periods for a facility. Thus, facilities whose metrics may consider particular time periods (e.g., an annual reporting interval) will need to agree on goals and metrics with other facilities that may have data-collection campaigns that span a different interval (e.g., multiple years). The two or more co-operating facilities will need to agree to support campaigns that operate across their resources over potentially new and unaligned time windows. This support will take the form of (i) reserving resources in a structured manner to allow for service levels that are acceptable to the

end-to-end campaign, (ii) prioritization agreements (to elevate or lower) priorities according to needs of competing campaigns, (iii) creating a joint proposal acceptance mechanism so that campaigns can request for resources for longer (and, potentially, variable) durations of resources across the facilities, (iv) creating failure recovery mechanisms and agreements, and (v) deciding on joint metrics and mechanisms for credit for success. Once the multilateral steering structure is established as noted previously, the steps will enable the support of long-term campaign patterns.

Supporting change in the nature of campaigns will be an important part of the shared governance structure in enabling long-term campaigns. Facilities and the science needs will evolve, and the governance structures must have a process to explicitly entertain and incorporate rule modifications to address change. There will need to be continuing requirement reviews, and changes in the landscape should be officially communicated to the governance bodies and DOE and partner agencies.

4. How Does This Practice Area Interact With and How Can It Advance the End-User Experience? (Refer to the Sprint 1 Synthesis)

The scientific processes today rely on access to a heterogeneous set of resources across a distributed set of facilities. This aspect is expected to increase, as science workflows become more complex, facilities add value by developing leadership expertise in specific areas, and networks evolve to further enable this integrated distributed facility paradigm. From a distributed cyberinfrastructure perspective, this will require a new class of services providing access to this infrastructure to users in the context of the science work upon which they are focused. That is, the science users should not need to develop cyberinfrastructure system expertise.

Orchestration of distributed heterogeneous networked resources, in a manner which is presented to users in a simple and science-objective meaningful context, will be an important part of these future systems. This will require continued development of networked services tailored to a full science workflow lifecycle's requirements. Design and architecture activities starting from a user and workflow perspective, and evolving to co-design across heterogeneous resources, interconnect models, and service interfaces will be needed. A unified facility interface and negotiation mechanism will facilitate the simplification of users' experience.

In addition, best practices need to be defined with respect to how domain science workflow developers can communicate their requirements to the developers of this next-generation distributed cyberinfrastructure. This may include a common ontology for how compute, storage, instruments, and networks systems are described, and services defined. Development of a common method to describe requirements and services would facilitate the development of API-based services. It may be better to coordinate on the common resource and service description concepts, and then let individual resource owners innovate on the specific services.

Key guidelines may be defined, such as:

- All services must be accessible via an API that is well documented.
- Distributed facility integration should follow a logical progression of:
 1. Policy definition
 2. Service definition
 3. API definition
 4. Tool definition

Common definitions and implementations for all these areas are encouraged but should not be enforced. This is because allowing facility-specific innovation is very important.

A balanced compromise may be that there is a base set of features associated with items 1 and 2, which are expected to be in place across multiple related facilities. Individual facilities are encouraged to innovate beyond those in a manner which best uses their systems. Items 3 and 4 can be facility specific, as long as they are well documented and accessible for orchestration middleware developers. In this practice area, the above guidelines may evolve into a co-design architecture for both common and domain science-specific orchestrators and workflow tools to grow organically, adapt to innovations in the distributed facilities, and develop nimble services that suit best for the end users.

5. Near-Term and Longer-Term Opportunities? Summarize Key Issues and Recommendations for Future Stewardship and Engagement, Including Potential Impact and What the Implications Are for Future Practice.

Historically, it has been difficult for DOE user facilities to meet their required metrics while operating in a cooperative and coherent manner. The formation of a community organization would better motivate a model of co-operating multiple facilities that ties these independent entities together, defines common mission goals, and provides a forum to discuss, compromise, and make decisions that advance the community.

Forming a steering group (SG) that bridges political/institutional divides in a democratic way is a clear short-term goal that would facilitate many long-term returns. There are many good examples of scientific research collaborations (High Energy Physics (HEP) collaborations, ESNet, WLCG, HPSS, etc.) where decisions are taken through representative methods that ensure the community's needs are being met and addressed, while moving forward to advance common goals. These examples can be used as a template for an inter-laboratory SG that would discuss and take decisions toward co-operating facilities. It would provide a place for all stakeholders (domain scientists, facility operators, DOE management, etc.) to be represented.

Through this SG, long-term benefits can be achieved. As an example, domain scientists have long requested unified environments at DOE HPC sites. Currently, each site has a very different computing environment due to system and distributor differences. In addition, there currently is little incentive for facilities to self-organize around common solutions with the current purely competitive environment. With an SG in place, democratic decision-making can be taken to adopt community solutions and build community support around common tools. The community could rally around common software for scheduling, data transfer, and web portals for scientists, where stakeholders can put forth suggested solutions, discuss them, and settle on the best option.

Each facility would describe its value-added capabilities and publish, document, and disseminate them. This would enable an end-to-end orchestration service to pull out capabilities and offer a baseline common-denominator service which may be enhanced with value-added services.

Appendix K

IRI ABA DESIGN PHASE: PRACTICE GROUP ON CYBERSECURITY AND FEDERATED ACCESS

IRI ABA Design Phase: Practice Group on Cybersecurity and Federated Access

Participants:

[Rachana Ananthkrishnan](#), [Carl Bai](#), [Amber Boehnlein](#), [Nicholas Buraglio](#), [Adam Carlyle](#), [Evan Felix](#), [Gina Fisk](#), [Tom Harper](#), [Carol Hawk](#), [Damian Hazen](#), [Raj Kettimuthu](#), [Shawn Kwang](#), [Eric Lancon](#), [Jeff Neel](#), [Michael Skwarek](#), [Adam Slagell](#), [Cory Snavely](#), [Adam Stone](#)

1. Vision/Goal Statement: What Does the Ideal Future Look Like for This Practice Area Focused on Enabling Integrated Research Infrastructure (IRI) and Science/Workflow Patterns Under Consideration (Time-Sensitive Patterns, Data Integration–Intensive Patterns, and Long-Term Campaign Patterns)?

Vision: Embrace the opportunity to accelerate the pace of innovation through novel secure design patterns and architectures to support open science-integrated architecture for seamless scientific collaboration

Goal: Create a business life-cycle framework that positions cybersecurity to enable and drive scientific productivity. This will facilitate scientists and service providers to develop low-friction data movement mechanisms and the ability to move workloads between facilities.

Goal: Develop, in partnership with the IRI community, well-documented and long-lived interfaces to all resources (storage, compute, artificial intelligence (AI), etc.) that use secure methods to authenticate and authorize use.

Goal: Establish a cadence and community of practice for the continued maturity of our shared vision and strategy for future cybersecurity principles, service offerings, and policies that align with our IRI partners.

2. What Is the Current State of the Art, Including State of the Community, Related to IRI?

Overall, the cyber posture of the labs follows National Institute of Standards & Technology (NIST) 800-53 and is risk-based, with risks accepted locally by the site offices. The labs have tremendous experience supporting large and small collaborations across institutions, but the kind of seamless integration imagined by IRI creates new needs, which may challenge existing cyber controls.

Cybersecurity and the approach to contractual requirements is typically DOE-lab specific and can be program-specific within a given DOE lab. Cybersecurity has often been focused on the perimeters of both sites and facilities, but interoperability across institutions challenges this paradigm.

Collaborative projects often present significant complications to the current site-specific model. Policy incompatibilities on collaborative tools hamper collaborative activities. Identity and access management present challenges where collaborators may exist outside of the institution and/or outside of the agency.

3. What Are the Most Important/Urgent Gaps in Research, Technology, Resources, Operations, Policy, and Environmental Constraints for This Practice Area?

Consider urgencies and gaps that cross all the IRI science/workflow patterns.

Cross-Cutting Urgencies:

Urgent: Federated Identity management is a pressing issue. Policy issues include the implementation of DOE O 142.3B, Unclassified Foreign National Access Program and differences in Identity Assurance Levels/Authenticator Assurance Levels (IAL/AAL) requirements at different labs/facilities. Technical implementations remain under discussion.

Important: Establishing design principles and architectural patterns for a new cyber paradigm to support IRI in a flexible and extensible way.

Important: Development of an operational security coordination model for IRI to address additional shared risk and to respond to incidents.

Important: There is a gap in understanding the totality of stakeholders to better manage and forecast strategic requirements, associated risks, and operational impacts: It may be insufficient to address issues only within the Office of Science labs/sites. This will enable a more strategic

and flexible approach in integrating, for example, university-based instruments and for partnerships with other agencies.

Pattern-Specific Considerations:

We expect that there may be very different cyber implementations across these patterns, driven by unique needs. For example, memory-to-memory transfers across facilities require substantial changes to the existing security architectures of computing facilities but are needed for certain kinds of low-latency integration. However, the goal would be that the design principles and models developed here would be able to accommodate the diverse implementations across the patterns. Exploration of the specific controls needed for different architectures needs in-depth exploration.

The patterns groups explicitly called out the following requirements:

- The conduct of research is not disrupted as security policy and practices evolve over time.
- Users do not have to re-authenticate multiple times per day in multiple different ways and at multiple different user facilities in the distributed infrastructure.
- Security policy allows end-to-end performance testing across sites.
- Machine-to-machine authentication:
 - A new model where devices, facilities, workflows, or agents are trusted at the same level as humans should be explored. This will be required to enable the automation of devices and experiments.

Time sensitive patterns group identified two additional considerations:

- Some consideration needs to be given to an abstraction layer between operational technology (OT) and information technology (IT) systems.
- Containment mechanisms that could affect time-sensitive workloads should be put in place, with the understanding that in the case of malicious actors, the entire workflow (which could cross multiple resources) has to be fully contained.

4. How Does This Practice Area Interact With and How Can It Advance the End-User Experience? (Refer to the Sprint 1 synthesis)

Firewall rules and lack of Federated Identity/Access are referenced repeatedly as barriers to the end-user experience. Due to factors beyond the IRI, there will be the need for changes within the cyber postures of the labs, which typically were architected ~20 years ago. Adding in IRI as

a design element for forward-going implementations could mitigate or remove the existing barriers.

5. Near-Term and Longer-Term Opportunities? Summarize Key Issues and Recommendations for Future Stewardship and Engagement, Including Potential Impact and What the Implications Are for Future Practice.

Explore barriers to IRI (e.g., Federated Identity, consistent trust domains, perimeter security) by creating forums to surface policy/risk-management barriers with lab, SC, and DOE leadership.

Explore/leverage the identity and data pillars of the move to Zero Trust Architecture (ZTA). Sustained activity is required to transition systems from existing perimeter-based cyber models to operationalized interoperable ZTA models.

Establish a working group (similar to Distributed Computing and Data Ecosystem (DCDE)) that looks at the issues of federated storage, perhaps with one or two target user communities.

Evaluate reciprocity of foreign national verification.

Appendix L

IRI ABA DESIGN PHASE: PRACTICE GROUP ON USER EXPERIENCE

IRI ABA Design Phase: Practice Group on User Experience

Contributors:

Eli Dart, Doug Benjamin, Christopher B. Fuson, Georg Rath, Jay Bardham, Wayne Joubert, Markus Diefenthaler, Peter Ercius, Qin Wu, Qingteng Zhang, Todd Munson, Vivek Thampy

User Experience Practice

Scientific computing systems should be a joy to use, meaning they should enhance a person's scientific productivity. The DOE Office of Science (SC) deploys bleeding-edge supercomputing resources that make higher-quality simulations and large-scale data analysis possible. However, these resources can be incredibly challenging to use.

High-Performance Computing (HPC) is now available in the cloud, and the community's perspective is shifting from general acceptance of computing facilities that are hard to use to demanding the resources better meet their needs. Scientists will use the resources that will help them achieve their goals in the least painful way possible.

Interdisciplinary work requires revamping our processes to incorporate perspectives from experts in different fields. Building trust at these intersections is a challenge and opportunity. User-centered design and design thinking create space for curiosity, exploration, and collaborative problem solving: critical elements of a high-functioning, diverse team.

Definitions

Usability can be described as the capacity of a system to provide a condition for its users to perform the tasks *reliably, effectively, and efficiently* while enjoying the experience.

The **Scope of User Experience for Integrated Research Infrastructure (IRI)** is the interaction of humans or robots with the computing and data infrastructure for the purposes of advancing the mission of the DOE SC.

Vision Statement

Suppose we were successful in advocating for all the things that make IRI highly usable. How will we know that the solutions we have put in place work?

Vision Statements

- Create a robust data and computing infrastructure that is a joy to use.

- Ensure the process for systems design and execution puts the user first.
- Achieve broad adoption of user-centered design for discovery and usability.

Key User Groups

We can potentially develop some more detailed personas that align to each of the different patterns, but for now we thought it important to call out two different user types.

Humans

- Scientists, developers, analysts: anyone who seeks to execute work on the IRI system.

Robots

- Software, systems, services, etc., that interact with the system in an automated way.

State-of-the-Art

Existing work in this area within the DOE complex has been considered “user experience” and is captured here.

Soliciting input for problem identification:

- ESnet requirements reviews
- DOE workshops
- Community-driven consensus problems over multiple years
- Surveys

Representation of user needs at facilities:

- User services/consulting groups
- Committees comprised of users

Gaps

DOE lacks a consistent, rigorous process for including the end users. The following are practices that are considered state-of-the-art when building systems or services for users/customers.

- User-centered design software practices: actively engaging users throughout the process
 - Verify that you are solving a real problem for your community.
 - Generate prototypes and conduct usability tests.
 - Use the feedback to refine the prototype.
 - Develop versions of the product based on the prototype and continue testing.
- Usability testing
 - Checking that what was built works across all intended user groups
- Product management

- A key member of a team that advocates for the customer/end user
- IDEA
 - Intentional inclusion of diverse perspectives through surveys, interviews, focus groups
- Clear messaging regarding the scope for a particular solution
 - Who is the intended end user?
 - What problem is being addressed?
 - What problems are not being addressed?

Pangeo reference

Magellan project learnings

Desired Outcomes

The following subsections describe, by type of pattern, what a human could do and what a robot could do if IRI is wildly successful

Time-Sensitive Pattern

A human could:

- Develop a trusted, reliable, and intuitive process that results in work accomplished within the needed time frame.
- Schedule work that requires multiple facilities in one place.
- Control multi-facility workflows from their experiment environment without individually managing separate facility logins/services/allocations.
- Run analysis concurrently with a large experiment.
- Develop software in an environment that reflects the production environment.
- Log into a single system and have easy access to all their data.
- Reduce or eliminate context switching to stay productive.

When time is of the essence, a user needs to be able to move quickly and maintain focus. Some of the outcomes relate to risk mitigation, and the others are to maintaining high engagement and productivity.

Data Intensive—Integration Pattern

A human could:

- Access storage resources without significant contention from other users.
- Easily move data throughout the storage hierarchy for analysis.
- Easily assemble all needed datasets in one location (transfer, reference, etc.).
- Avoid logging into multiple systems (because everything is integrated).
- Transfer files at a predictable speed.

- Shift seamlessly from one HPC site to another if a resource fails.
- Integrate and explore data interactively through Jupyter, RStudio, or the command line.

Data-intensive work has an emphasis on storage, I/O, and interactivity. The main barriers to productivity are distribution and facility policies that prioritize HPC.

Long-Term Campaign Pattern

A **human** could:

- View all data associated with the campaign through a single portal, regardless of data location.
- Set data access policies on all data they own from one system.
- Search and retrieve data from the campaign.
- Have confidence that the tools/processes will be available and well supported for the duration of the campaign.
- Access the needed resources for the full duration of the campaign.

User productivity is driven by data organization and management for long-term campaigns. Consistent access to resources during the life of the campaign makes it easier for users to reproduce analyses.

Time-Sensitive Pattern

A **robot** would have:

:

- Access to job status, success, failure information via API.
- API-based job submission.
- The ability to determine where a job will run the fastest.
- Fast movement of data from experiment to compute with persistent credentials.

Data Intensive—Integration Pattern

A **robot** would have:

- The ability to analyze data transferred from other sites.
- Fast movement of data from experiment to compute with persistent credentials.
- API-based access.

Long-Term Campaign Pattern

A **robot** would have:

- The ability to access systems with the same credentials for a long time.
- API-based access.

The main difference between robots and humans is the need for programmatic interfaces that can determine the characteristics of the computing facility and availability of storage resources, and move data quickly. Robots need autonomous access with long-lived credentials.

Opportunities

It is clear from the discussions, interviews, and synthesis of all IRI inputs that a core set of barriers can and should be addressed. As solutions are being considered and prototyped, users should be engaged to evaluate the efficacy of a solution and provide feedback. The feedback should be incorporated and the cycle should repeat. The cycles should be fast for the prototyping phase, and then production versions of the solution can be released.

For software infrastructure solutions, i.e., workflow managers, there should be enough resourcing to ensure that the solution remains viable over the long term. This is important, as no user wants to invest effort in leveraging a solution only to have it disappear in the future. We recommend that a rigorous effort be undertaken to recommend existing solutions with robust user communities prior to developing new software.

The IRI architects should engage in a concerted effort to communicate the specific problems that will be addressed first.

Lightning talk points

- Vision statements
 - Create a robust data and computing infrastructure that is a joy to use.
 - Ensure the process for systems design and execution puts the user first.
 - Achieve broad adoption of user-centered design for discovery and usability.
- It is clear from the discussions, interviews, and synthesis of all IRI inputs that a core set of barriers can and should be addressed.
 - Users need capabilities that reliably perform/ behave as they should.
 - Easy-to-integrate and easy-to-use are important aspects of the user experience.
 - IRI needs to grow capabilities in features, sophistication, number, and scope over time as success begets success.
- Two important constituencies
 - Humans
 - Scientists, developers, analysts: anyone that seeks to execute work on the IRI system.
 - Robots
 - Software, systems, services, etc., that interact with the system in an automated way.
- For software infrastructure solutions, i.e., workflow managers, there should be enough resourcing to ensure that the solution remains viable over the long term.

- No user wants to invest effort in leveraging a solution only to have it disappear in the future.
- We recommend that a rigorous effort be undertaken to recommend existing solutions with robust user communities prior to developing new software.

Appendix M

IRI ABA DESIGN PHASE: PRACTICE GROUP ON WORKFLOWS, INTERFACES, AND AUTOMATION

IRI ABA Design Phase: Practice Group on Workflows, Interfaces, and Automation

Team members

Deborah Bard <djbard@lbl.gov>, John MacAuley <macauley@es.net>, Bin Hu <bhu@lanl.gov>, Christian Engelmann <engelmann@ornl.gov>, Ben J Mintz <mintz@ornl.gov>, Rafael Ferreira da Silva <silvarf@ornl.gov>, Sterling Smith <smithsp@fusion.gat.com>, Stephen Chan <sychan@lbl.gov>, Beckman, Pete <beckman@anl.gov>, Michael Kirby <kirbsox@gmail.com>, Malachi Schram <schram@ilab.org>, William (Bill) E. Allcock <allcock@anl.gov>, Pete Beckman <beckman@mcs.anl.gov>, Jha, Shantenu <shantenu@bnl.gov>, Mehta, Apurva <mehta@slac.stanford.edu>, Nicola Ferrier <nferrier@anl.gov>

Scope:

Workflows, Interfaces, and Automation: Assembling system components to support Integrated Research Infrastructure (IRI) science cases systematically in the form of end-to-end pipelines. Users should be able to manage these overlays and middlewares effectively across facilities.

1. Vision/Goal Statement: **What Does the Ideal Future Look Like for This Practice Area?**

The ideal future for workflow interfaces and automation is an open standard for workflows that allows interoperability and reusability, combined with a set of reference implementations. We should not commit to a specific technology stack as there is no one-size-fits-all in this space. Instead, the standard should empower users to develop their own workflow. The standard should incorporate the following:

- It should scale in terms of the number of instruments/sensors in the workflow, as well as the number of users.
- It should enable monitoring of each component of the workflow, built into the standard rather than an afterthought.
- It should standardize application programming interfaces (API) / messages, data schema, behaviors, and error (exception) handling.

- It should introduce minimal additional latency in the workflow (time-sensitive patterns).
- It should provide access to, and information about (e.g., last change), external data sources in a way that can allow automated queries and secure data verification (data-intensive patterns).
- It should leverage other standards such as making data Findable, Accessible, Interoperable, Reusable (FAIR) / Public Reusable Research (PURE) / Zero Trust rather than re-inventing them.

In an ideal world, this would be accompanied by unified universal authentication and authorization. More realistically, we need a framework that allows sites to have the required level of access control, with the onus on users to provide the needed authentication. [Note: Cybersecurity and Federated Access]

2. What Is the Current State of the Art, Including State of the Community, Related to IRI?

Currently there exist many fragmented, isolated APIs that cannot communicate with each other. Some are REpresentational State Transfer (REST) APIs, some are pub/sub.

- Examples include Globus, S3 for data movement; Software-defined network for End-to-end Networked Science at Exascale (SENSE) , Open Grid Forum (OGF) Network Service Interface (NSI) Slurm, K8s for resource orchestration; Superfacility API and Workflows Community Initiative for component interfaces.

There is no common authentication framework, but if the user can collect the relevant tokens, it is possible to work across sites. [Note: Cybersecurity and Federated Access]

Many good monitoring tools exist within a stack, but not across the whole workflow. Resilience in distributed systems for complex workflows is also an unsolved problem, e.g., a computing resource going away means a workflow falls apart.

- Examples in this space are the Google Cloud stack, K8s, and the Large Hadron Collider (LHC) Grid (Worldwide LHC Computing Grid (WLCG)).

Specific Example:

Open Grid Forum development of the NSI specification. Members from the National Research and Education Network (NREN) across the world were working to develop a set of standards documents relating to network services. Weekly group meetings to discuss open issues, track progress, etc., and quarterly face-to-face meetings (two to three days) to make focused progress, review proposals on specific solutions for different problems, etc., were held. Twice a year, we coordinated demonstrations of the standards so far at specific conference events. We forced software teams in multiple organizations to focus on implementations, discovering issues

with the standards and different implementations. Each standard demanded many hours of work and required many hours to complete.

3. What Are the Most Important/Urgent Gaps in Research, Technology, Resources, Operations, Policy, and Environmental Constraints for This Practice Area?

The most important gap is a standard for interoperable workflows. The end points need to be identified, via a reference architecture, so that job/data services and resource descriptions can be standardized. Resource/infrastructure monitoring should be considered at the outset and given the same priority as other areas of implementation.

A common authentication/authorization standard is desired but may be an impossible dream. By laying out a standard for communications via tokens, each site can issue its own secure tokens, thereby facilitating secure communications. The desired standard should define the process for handling and exchanging tokens. [Note: Cybersecurity and Federated Access]

4. How Does This Practice Area Interact With, and How Can It Advance, the End-User Experience? (Refer to the Sprint 1 Synthesis)

Users want to work across facilities but do not know how to do this. We want to ensure this is made easy for them via technology and policy. Currently there are huge start-up barriers for scientists, so providing both out-of-the-box technology (e.g., for routine tasks) and the ability to refine/customize a workflow will improve productivity.

We will need to carefully consider what metrics will show success and how those metrics come into the end-user experience.

We will need a carefully considered training and documentation plan, for both users and admins.

5. Near-Term and Longer-Term Opportunities? Summarize Key Issues and Recommendations for Future Stewardship and Engagement, Including Potential Impact and the Implications for Future Practice.

Short Term:

- Identify communities that we want to build and leverage. The challenge is how to handle the community of communities!
- Build understanding and agreement on best practices and develop that understanding across facilities. This will help get buy-in from facilities and help change policies (which is slow)
- Will need bottom-up support from end users, to tell facilities this is where we need to go.
- Will also need top-down support/mandate from funding agencies.
- An agile approach to developing a standard, that incorporates user feedback/testing from the beginning, will help make short-term progress and engage users. We must not workshop/white paper this to death.
- Need a management plan along with an update cycle for the plan. And a path on how to keep this plan alive/funded/successful.

We will need a management plan, which considers both short-term milestones and long-term goals. There needs to be a corresponding long-term funding commitment for users, developers, and facilities.

- To start with, we will develop standards/specs/ref implementations (w/ connection to users), then later phases will ensure facilities support them. Showing these different phases from the beginning will be important to the success of IRI.
- Facilities also need to be able to plan for this and have multiyear procurement efforts/staffing that need to be considered.

Dependencies From Other Groups

[IRI ABA Design Phase - Charge for Long-Term Campaigns](#)

[IRI ABA Time-Sensitive Patterns Report](#)

[IRI ABA Data-Integration-Intensive Patterns Blueprint](#)

[IRI ABA PRACTICE Groups Charge: Portable/Scalable Solutions](#)

Appendix N

IRI ABA DESIGN PHASE: PRACTICE GROUP ON SCIENTIFIC DATA LIFE CYCLE

IRI ABA Design Phase: Practice Group on Scientific Data Life Cycle

Participants:

[Franck Cappello](#), [Vincent Garonne](#), [Michael Hu](#), [Eric Lancon](#), [Kurt Maier](#), [Giri Palanisamy](#), [Amedeo Perazzo](#), [Lavanya Ramakrishnan](#), [Sreeranjani "Jini" Ramprakash](#), [Kelly Rose](#), [Simon Roux](#), [David Schissel](#), [Jonathan Taylor](#), [Rama Vasudevan](#), [Andrew Wiedlea](#),

Description of This Area:

Users need to manage their data across facilities from creation (along with metadata), staging, movement, storage, dissemination, curation, archiving, publishing, etc. Technologists need to understand the requirements across different communities to develop solutions appropriate for an integrated research infrastructure (IRI). This pattern group is focused on this aspect of the scientific data life cycle.

1. Vision/Goal Statement: What Does the Ideal Future Look Like?

The ideal future will have an established and well-articulated system design pattern that maximizes the return on scientific investment by capturing, cataloging, curating, and sharing data, including rich provenance information to provide a Findable, Accessible, Interoperable, Reusable (FAIR) based data pipeline from acquisition to publication and future use with end user-focused interfaces.

Background considerations for the vision:

- We recognize that several federal and organizational data compliance requirements exist and should be considered in the design patterns.
- Provenance should include rich forms of metadata including but not limited to ownership and citation.
- A system design pattern is analogous to a set of specifications that various entities (facilities, institutions, centers, etc.) can take to design and implement solutions that will result in seamless integration with other research infrastructures.

2. What Is the Current State of the Art, Including the State of the Community, Related to IRI?

Definition: Communities here are the various scientific domains/program offices that are represented by the group members.

Common across different domains:

- Data curation is very domain specific and tough to generalize and get right.
- There are differing data management capabilities and policies based on the home base/home program of the project/user.
- Data catalogs are diverse and distributed.
- Expertise in data management varies considerably between various communities and facilities.
- Long-term curation policies, when they exist, are different across facilities.
- FAIR principles are implemented sporadically at best (except where federal acts have mandated them).
- Short-term data needs (real-time analysis) versus long-term data needs (reuse and reproducibility) vary and can be complicated to navigate.

Neutron sciences:

- Facilities act as custodians of the data, with few to no policies defined.
- Facilities also provide data management and access to centralized data services.
- Simple machine scrapeable metadata are automatically collected and stored.
- Data processing frameworks are developed with Continuous Integration (CI) and versioned digital artifacts are Digital Object Identifier (DOI) citable.
- Data processing workflows and parameters are stored within the Hierarchical Data Format version 5 (HDF5) data container along with the software version used for the processing.
- It is recognized that the time to market (beam time to publication) is rate limited by data interpretation and associated data services.

High-Energy Physics and Nuclear Physics:

- Centralized large storage facilities (hundreds of PB) supported by operation programs exist.
- Data is curated for decades.
- Data is largely not public with the notable exception of some domains related to astronomy.

Office of Fossil Energy Carbon Management:

- Energy Data eXchange (EDX) includes services for software and data, including:
 - o DOI services for publication of data products (including software, datasets, databases, tools, application programming interfaces (API), and other R&D data products).
 - o APIs to facilitate use of data via EDX for user needs.

- Git-repo federation services to support software development privately, and then maturation and publication of those products via EDX but through the Git's as appropriate.
- EDX presently manages data products (of all types) from Fossil Energy and Carbon Management (FECM) R&D spanning materials science, computational, subsurface, infrastructure, energy justice (EJ) / social justice (SJ), geospatial, systems analysis (life cycle and techno-economic), experimental/lab generated, field, and more.

EERE:

- Some programs use National Renewable Energy Laboratory (NREL)'s OpenEI platform, which is a public-facing data management and curation capability in its present state.

Materials Science:

- Most data are not available, and reproducibility is poor to nonexistent in most cases.
- Simulations/theory datasets are usually disconnected from corresponding experimental validation.

Geospatial Information Science:

- More mature community with the first full-time geospatial information officer (GIO) at DOE.
- DOE has a geospatial project management advisory group, led by the GIO, and an R&D community DOE geospatial users group (GUG) that meets biannually and offers input and support related to DOE geospatial R&D and IRI-related activities.
- The [2018 Geospatial Data Act](#) requires all federal agencies to manage and report their geospatial data products, in compliance with the requirements of the act in alignment with Federal Geographic Data Committee (FGDC) guidance.
- Most DOE users and facilities use Esri Arc resources for their geospatial infrastructure needs, but there is a lack of federation of these capabilities.

BES:

- Some facilities offer a different experience even across different beamlines within the same facility.
- Some unification efforts are driven through pilot programs. The most successful cases of shared ideas and methods have derived from the facilities spontaneously adopting tools developed at another facility.

Subsurface:

- Data is very challenging to manage and integrate.
- The tools and capabilities for managing, virtualizing, and conducting computational elements in this domain are often costly, commercial in nature, or restricted-access codes and platforms generated by various DOE national laboratories.

ASCR:

- Compute facilities are currently working with several DOE user facilities (Advanced Photon Source (APS), Advanced Light Source (ALS), Stanford Linear Accelerator Center

- (SLAC), etc.) on connecting, transporting, and processing data across these facilities through complex workflows and other software (such as Balsam).
- Globus plays a big role in data movement and management in this community.
 - The efforts related to data management at scale are mostly at the compute facility level combining efforts with individual domain communities such as light sources, climate community, etc.

Nanoscale Science Research Centers (NSRCs):

- Centers have local solutions to their data storage and data workflows issues.
- There are no intentional connections between the different facilities beyond ad-hoc transfers (e.g., using Globus).
- Each facility produces different types of data without common agreement on formats, data structures, schemas, etc.
- There is a need to accommodate both short-term storage (for enabling real-time or near real-time analysis, or model training) coupled with longer-term storage for reuse and reproducibility.

Magnetic Fusion Energy:

- Contained within an experimental user facility, magnetic fusion data (including metadata) is well maintained and curated.
- Long-term centralized data storage is considered within the mission of the facility.
- Access to this data (via a unified API) is available to anyone who joins the collaboration and agrees to the usage policy.
- Providing external open access to select data sets has not yet occurred, but there is active work in this area examining how this will be done and what policies must be in place to support this capability.
- For the theory/simulation community, there is no common data management system.
- Unlike the experimental data, there is no external open access to select datasets.

Genomics:

- The community is relatively mature at this point, and there are a handful of well-identified databases through which all the data are typically shared and/or retrieved.
- “Data ownership” plays a significant role in this community.

3. What Are the Most Important/Urgent Gaps in Research, Technology, Resources, Operations, Policy, and Environmental Constraints for This Practice Area?

Top level urgent gaps / issues:

- Not being able to articulate the value of implementing data management principles (such as FAIR) when compared to the return on that investment is preventing the creation of a cohesive solution:
 - o Concentrated costs (current data producers or curators must take on the bulk of resources needed to implement) but distributed benefits (vague, abstract future state where this may be useful to someone else).
- At the DOE agency level, there is no clarity on the data requirements <https://edx.netl.doe.gov/reference-shelf/orders-initiatives-and-policies/>.
- Lack of clarity of “ownership” of the data being generated for curation or stewardship.
- Lack of federation and creation of data silos, extending to solution silos.
- If data has been moved around so they are persistent, this puts the creators and potentially any users into navigating site-specific cyber rules.
- Data management when researchers leave (e.g., student leaves and principal investigator [PI] wants their files) can be solved with group management tools, but in practice lots of research is in individually owned files. Similar issues arise when a project terminates.
- So much experimental data is wasted. Often these data are used as part of a paper and then abandoned and never found again or used for future research.

4. How Does This Practice Area Interact With and How Can It Advance the End-User Experience? (Refer to the Sprint 1 Synthesis.)

We start with a visual representation of the elements of the scientific data life cycle.



Figure: Example data life cycle components from Atmospheric Radiation Measurement (ARM) observational facility.

Data, from creation to curation, to long-term storage and use, are critical to the scientific user. High-quality, easy to use software that provides users with ready access and relevant metadata at every stage of the life of data is critical to turn the experimental investment (national facilities, research funding, scientists' and students' time, and efforts) into science.

Multidisciplinary scientific endeavors rely more and more on data integration and cross referencing. The availability of curated and previously used data from across science disciplines will encourage more systematic research and further metastudies.

From real time to future reuse, from small footprint data in memory to large petabyte archival, the needs of the user community and the experience users seek is diverse and sophisticated.

Every user community could benefit from the adoption of tools that aim at facilitating citing, accessing, and reusing publicly available scientific research datasets produced by its researchers and assigning persistent identifiers to these datasets. More specifically, these identifiers would offer the following benefits:

- Enable researchers to discover, access, and reuse data to verify the original experiment or produce new results with the latest methods.
- Facilitate linkages among documents or published articles, their underlying datasets, and other related research objects.
- Make datasets that have been announced and registered become discoverable through search tools, increasing the opportunity for discovery of additional data, specialized interfaces, toolkits for data analysis, etc.

- Make data easy to cite in a standardized way, making datasets more easily citable so that data creators, contributors, data centers, and others can receive proper attribution.

Facility users ask: Where is my data, how do I access them, and what do they mean? Solutions to address those questions should be considered a high priority.

5. Near-Term and Longer-Term Opportunities?

Near-Term Opportunities:

- Allocable storage resources, with a need for different classes of storage, to support the needs of researchers
- Workforce development at facilities through education on data management principles and data analysis, including the use of AI/ML
- Gathering the most complete information about data usage for scientists and data stewards
- AI-based upgrades for improving data management and facilitating data interoperability
- Improving the operational efficiency of data life cycle components by adopting community-developed standards
- Establishing prototypes of a data and supporting computing infrastructure for next-generation data science, including global-scale search capacities
- DOE's Office of FECM has recently implemented Funding Opportunity Announcement (FOA) / Field Work Proposal (FWP) data requirements for all its funded R&D. This guides performers on expectations for product preservation/delivery and also helps FECM programs be more compliant with federal and executive requirements and laws. That FOA language was vetted by DOE legal and could be leveraged by other offices or modified for their use.
- Developing DOE-wide incentives for the implementation and use of FAIR principles

Longer-Term Opportunities:

- Establishing an open DOE data catalog that enables researchers to search and find DOE-funded data located across the DOE complex
- Identifying needs and sites for archive center(s) to preserve DOE-funded data long term
- Establishing data and supporting computing infrastructure for next-generation data science, including global-scale search capacities

- Establishing and supporting long-term multi-sites archive center(s) for DOE-funded data
- Creating an opportunity to partner with the DOE OCIO GIO about geospatial-related IRI needs and build upon their efforts for enterprise license agreements, compliance with GDA 2018 requirements, and more
- Enabling multidisciplinary repositories, along with APIs, applying FAIR principles
- Setting up data management as a service—offered at each facility to the user community—where user teams could potentially specify one person in their group to be the one to perform domain-specific data-related tasks like specific metadata creation and curation

Workforce Opportunities:

The group recognizes the value of research software engineers to the work required in the practice area and suggests DOE encourage young scientists to transition into research software engineers and data management specialists and provide career opportunities for them at national facilities.

We also recognize the workforce development challenges associated with this type of role change for individuals in an ecosystem that places publishing over all else, and encourage DOE to evaluate ways to recognize the contributions of the people doing this critical infrastructure work.

There is a need for an enduring statistical sciences capability to support the creation of ontology/data schema and to support the design of data structures to support the analytic demands that may be placed upon data sets and inferencing.

Appendix O

IRI ABA DESIGN PHASE: PRACTICE GROUP ON PORTABLE/ SCALABLE SOLUTIONS

IRI ABA Design Phase: Practice Group on Portable/Scalable Solutions

Participants:

Antonino Miceli <amiceli@anl.gov>, Feiyi Wang <fwang2@ornl.gov>, Jerome Lauret <jlauret@bnl.gov>, Ezra Kissel <kissel@es.net>, Chris Knight <knightc@anl.gov>, Matthew Cherukara <mcherukara@anl.gov>, Verónica Melesse Vergara <vergaravg@ornl.gov>, Paolo Calafiura <pcalafiura@lbl.gov>, Jason Zurawski <zurawski@es.net>

Scope

Portable/scalable solutions: Users and technologists need to move/translate their efforts across heterogeneous facilities (be portable) as well as go from smaller to larger resources (be scalable).

1. Vision/Goal Statement: What Does the Ideal Future Look Like?

We envision an integrated research environment where users can seamlessly move and scale all or portions of their workflows across resources. This portability can be to achieve better reliability and exploit unique capabilities or scale that may accelerate parts of a workflow. This vision not only helps to accelerate research and boost productivity but serves to unlock new opportunities for discovery.

2. What Is the Current State of the Art, Including State of the Community, Related to IRI?

Several related efforts or existing platforms can serve as an inspiration for this vision:

- The Open Science Grid is in many ways a precursor to the IRI, since it tackled similar themes.
- The commercial cloud space through its API-centric model, standardized interfaces, infrastructure as code, and emphasis on scalable, distributed infrastructure. This also includes tools and frameworks, such as Kubernetes, and extended ecosystems and tools like Terraform.

- The High Energy Physics projects with well-established computing grids that address the challenges holistically and deeply integrated with their scientific process.
- Extreme Science and Engineering Discovery Environment (XSEDE) and Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS), which have also tackled some of these challenges and could provide many lessons learned and best practices.
- Foundational technologies like Globus, containers, and service orchestration platforms will likely play an important role.

3. What Are the Most Important/Urgent Gaps in Research, Technology, Resources, Operations, Policy, and Environmental Constraints for This Practice Area?

Many gaps are common across the patterns, including:

- Allocations and accounts that span resources, a precursor to enabling portability.
- Common interfaces and tooling across resources and stable over long periods of time such that users feel confident enough to invest in adoption.
- A well-defined governance structure that spans resources and facilities that help to define and oversee standards and common policies and maintain a common roadmap.
- Common frameworks that can easily be adopted within a pattern area so that, for example, bringing in a new instrument takes minimal effort.
- Expertise both within technology areas as well as domains that can assist communities in adapting their workflows to effectively use integrated infrastructure.
- Schedule and resource abstractions that allow users to express what their workflows require and then enable the users to schedule across distributed resources.

Time-Sensitive Patterns

Unique challenges in this area revolve around the immediacy of results.

- Consequences of this are that resources must be available and potential latency effects must be considered. Having data potentially replicated at multiple resources could enable additional redundancy, but the time required for replication could bound potential options.

- Standard queues often conflict with this requirement. This becomes even more challenging for workloads that may require both immediacy and scale.
- Frameworks or models that work for streaming data use cases, not just bulk file movement models, are needed.

Data Integration–Intensive Patterns

The unique challenge here is having the relevant data easily accessible so that resources can easily be applied to analysis. This particular pattern may not be as latency sensitive, but ease of access/availability when needed is important.

Other challenges are making sure the data sets are Findable, Accessible, Interoperable, and Reusable (FAIR) and stored in common formats so they can be analyzed correctly. This goes beyond portability and scalability but is a precursor.

Portability here is less about reliability or urgency but may be to leverage unique capabilities that are necessary for some types of analysis (e.g., specific architecture or specialized hardware).

Long-Term Campaign Patterns

Since these campaigns can last years or even decades, longer-term allocations are critical. As projects plan out campaigns, they need to know that access to resources can be counted on over the lifetime of the project.

The interfaces need to be not only common across resources but stable over extended periods of time. Reproducibility is also a challenge, as various resources will likely have to evolve over the lifetime, especially for very long campaigns. So balancing the need for systems and services to be updated and evolve over time with the need for reproducibility is a challenge.

4. How Does This Practice Area Interact With and How Can It Advance the End-User Experience?

The impact of achieving this vision is evident, since it goes directly to boosting productivity and eventually opening the possibility of new modes of discovery. While it is already possible today for projects to distribute parts of their workloads across distributed resources, and many projects are already doing this, it often requires very advanced expertise, significant investment, and constant maintenance. This means smaller projects can often not replicate these approaches, and achieving the next level of portability and scalability is not feasible because of the effort that would be required. By having a highly integrated distributed infrastructure that supports portable and scalable workflows, projects of all sizes can exploit the capabilities and start to tackle even larger challenges.

5. Near-Term and Longer-Term Opportunities?

Near Term

- Establish a governance group that would start to define common interfaces and standards, allocation models, authentication and access policies, roadmaps, and processes for maintaining these.
- Establish a cross-site organization that has a range of expertise in critical areas.
- Create an inventory of capabilities and services that can help identify potential common interfaces but also areas of expertise.
- Define a set of pilots with demonstration goals with each pattern being potentially represented by more than one domain.
- Launch a series of hack-a-thons to support the pilots and assist projects in integrating interfaces or tools into their workflow. These should work towards demonstrations.
- Review and report out on any lessons learned from Open Science Grid (OSG), XSEDE/ACCESS, and other related efforts.

Longer Term

- Provides an opportunity to redefine the ecosystem and landscape as we work across sites.
- Expand on the organization and governance model to extend to new communities and potentially bridge across agencies.
- Establish workforce development opportunities, including postdoc fellowships.
- Host competitions to help spur innovation in the area.
- Fund R&D to greatly improve on the current state of the art, with demonstrations of new capability.
- Address challenges around software sustainability, especially for critical software and services that underlie the integrated infrastructure.

Cross-Cutting Themes

- Governance
- Access to experts and training/workforce development

- Strong dependency on many of the other practice areas. Portability requires addressing challenges in authentication and access, resource co-operation, automation and APIs, and data life-cycle management. And much of the vision starts to align with issues around usability.

Appendix P

IRI ABA DESIGN PHASE: FOCUS AREA ON OVERARCHING IRI PRINCIPLES

IRI ABA Design Phase: Focus Area on Overarching IRI principles

What Is the Integrated Research Infrastructure (IRI)?

IRI is an overarching organization connecting sites, hardware, people, and knowledge, working together to advance science.

Executive Message:

IRI enables the integration and federation of DOE investments in unique world-class facilities to work collectively to advance scientific discovery beyond what is possible today.

IRI is designed to enable science cases that cross multiple facilities (both within and across physical sites). It encompasses all DOE Office of Science (SC) programs, as well as the broader scientific community outside of SC. IRI prioritizes user-focused design, with users being an integral part of the design and development process.

IRI Overarching Principals:

- **IRI is designed to enable science cases that cross multiple facilities** (both within and across physical sites).
- **IRI is a process that encompasses all SC programs and the community outside of SC.**
 - All offices need to be involved in IRI all the way through to keep this engagement active beyond the Architecture Blueprint Activity (ABA).
 - The IRI is a “science enabler,” not a purely computing project.
 - Facilities and labs need to have buy-in from the beginning. IRI needs to cross the boundaries between labs and facilities.
 - Consider that many facilities serve scientific communities outside SC, e.g., university-based instruments.
- **IRI prioritizes user-focused design.**
 - Documentation, training, and ease-of-use will be built in from the beginning, not as an afterthought. Both reference docs and reference implementations will be provided.
 - IRI will iterate with users early and often in the development process.
 - Facilities will have their own implementations of IRI, but need to ensure those implementations are user friendly and consistent across sites.
- **Principle of nonuniformity: facilities and labs have different missions, constraints, and governance structures. IRI will respect and facilitate that.**
 - IRI can help defuse this with its own funding / program model, rewarding work that works in multiple IRI facilities.
 - IRI will not define implementation (to avoid being too proscriptive). We will work towards common standards and reference architectures/implementations, but sites will have the flexibility to implement according to their needs.

- IRI will support facilities to opt-in and implement IRI infrastructure, including funding for the work this places on facilities.
- **Principle of interoperability: IRI will enable scientists to move between the facilities, respecting the needs of both the scientists and the labs/facilities.**
 - This includes a strategy for authentication across sites, for both users and machines.
- **IRI will be persistent: persistence will be designed into all infrastructure plans**
 - Infrastructure requires continuous upkeep and updating. These require long-term funding and planning.
- **IRI will continuously evolve, reacting to changing needs of scientists and technologies.**
 - Documentation will be updated along with any changes to IRI standards or infrastructure.
- **IRI will use an agile/iterative development process that will iterate early and often with users to represent their requirements.**
 - We will not white-paper this to death.
 - We suggest adoption of agile principles of early and frequent interactions between developers/designers and end users as well as delivery of early prototypes for evaluation.
 - Whenever possible, the science users should be part of the sprint team.
 - We suggest setting up a requirements gathering process early in the effort to understand in depth the needs of the experimental facilities and how to meet their requirements. Software designs and products coming from the IRI will be derived from a mixture of end-user long-established techniques and methodologies on the one hand and what the IRI can offer to add value to their efforts on the other.
 - IRI will regularly re-evaluate its processes.
- **The IRI framework will appear simple to the end user, to reduce the barrier to getting started and overcome institutional inertia.**
 - Upsides of using IRI will need to be obvious, not just in lofty terms of advancing science, but also because it makes scientists' lives and science easier.
- **IRI will measure progress and success in each area.**
 - Articulating measurable benefits will help get buy-in.
 - IRI should consider metrics, e.g., happiness of scientists using the systems.
 - Metrics can be used to determine where to develop further, and whether to cut parts that do not work.
 - IRI should perform continuous introspection of the system and decisions.
- **IRI will not reinvent the wheel.** We will adopt existing standards/principles such as Findability, Accessibility, Interoperability, and Reuse (FAIR), and will leverage lessons learned from the community, e.g., funding agencies, industry, and existing projects.
- **IRI will be inclusive** and will create an equitable framework for the entire community.

Next Steps:

- Immediate action: **need a plan for continuing communication across facilities.**
- **Define the process by which labs will contribute** and also how they can bring IRI back to their own facilities.
- Identify **long-term business goals** and where they conflict (e.g., flexibility with a standard). This will directly affect the architecture planning.
- **Gather lessons learned from existing projects**—what went well, what did not—and collect in a repository for reference in the planning process.
- **Need buy-in from agencies, as well as from scientists.** A near-term action could be to frame what is needed to get buy-in, and how this would translate into decision-making in DOE headquarters.
 - Consider whether a top-down requirement to support IRI is useful/possible (e.g., build on the requirement for a data management plan).
- **Consider/review the funding model.** Facilities have long-term funding for long-term projects, making joint cross-facility collaborations difficult in the short term.
 - Write down all the funding constraints and needs that IRI needs to consider/include.
- Define a **set of end points** for IRI to help us define what we want to connect and help shape what the standard needs to do.
 - In computing, instruments, and networks, frame the general space of end points (and their maturity level), which can be used to spec out the IRI framework.
 - Specific use cases can be used to validate the standard/reference architecture as it develops, including existing projects/frameworks that could be used as IRI early adopters.
- Define a **compatibility matrix**, a base-level set of things needed to operate within IRI (e.g., a minimal standard for maximal compatibility).
 - Identify areas that this matrix should address: compute / storage / user access / archival / data distribution, etc.
- Need to figure out **where to standardize** and what is obvious/easy to start with.
 - In an agile environment, standards can grow through incremental design and implementation (e.g., Google’s internal convention of programming in Go, using protobufs across their software, etc.). We in the DOE have rallied around Jupyter Notebooks. It isn’t a standard, but it is a good idea broadly.
 - The objective is to make things easier for the user, and initial standardization must deliver this with a balance between making it easier to implement something and making it easier to use.
- **Continue the requirement-gathering** process to refine what we’ve heard in the patterns (i.e., take a deeper dive). Focus on doing the most good up front: 80% is good enough to start with.

Research opportunities:

- Guidance on the best approach for problem solving for user-centered design and human-computer interaction, grounded in research.
- IRI could be seen as an overarching umbrella organization, akin to ECP/OSG.

- ASCR research is already funding work relevant to IRI, e.g., data management and distributed management systems.
- AI/ML.
- Grid computing was a large area of investment in the past for ASCR.
- Transition from research into operations is a specific effort that needs support and funding (e.g. the DOE-funded software sustainability effort).
 - Will need to test across multiple sites, and so will need to develop a framework to do this.
- Translational research. Meet applications where they are today and translate them to IRI-compatible production. Design solutions for future applications based on where they are today, rather than future fresh/new apps. Not building things from the ground up: use existing tools.
 - Being able to scale up will distinguish apps/tools to focus on. IRI requires scale.

Appendix Q

IRI ABA DESIGN PHASE: FOCUS AREA ON GOVERNANCE/ STEERING APPROACHES

IRI ABA Design Phase: Focus Area on Governance/Steering Approaches

Executive Message:

Addressing national and societal grand challenges and unlocking new opportunities around energy, science, and technology for US competitiveness will require highly coordinated, collaborative research and integrating capabilities across our world-leading facilities, which currently operate largely independently. We can achieve this vision if the facilities, projects, and science communities have the right incentives, governance, and operating structure to enable them to deliver an integrated research platform, accelerating time to discovery and time to innovation.

Governance/Steering Approach:

- Primary coordination: governance
 - Policy: a place to define community policies, point of view, priorities. Need to consider short-, medium-, and long-term objectives and priorities.
Extract policy areas in the principles document to apply to the charter of this working group (WG).
 - Technical:
 - WG: standards/interoperability: identify the need and guide the development of community standards (including international efforts).
 - WG: evaluation: evaluation/maturity assessment body
 - WG: cyber
Extract technical areas in the principles document to apply to the charter of these working groups.
- Features of the governance structure:
 - Multi-tiered
 - Executive board
 - Higher level, a “board” of stakeholders (steering, “council”)
 - Planning and view into technical/backlog
 - Technical board/committees
 - Policy and technical working groups would explore opportunities and piece together outcomes to recommend future campaigns and strategic priorities
 - Stakeholder representation (agencies, labs/facilities, science domains, technology providers, program managers)
 - Open meetings (e.g., like the ASCR Science Advisory Committee (ASCAC))
 - Time-limited roles, rotating roles
 - External reviews

- Agile engagement as the science ecosystem evolves. Allows evolution of representation and priorities (e.g., a proposal to add new stakeholders).
- Clear accountable structures and expectations. Support autonomy but ensure minimal requirements.
- Rethink support model for a user of IRI (it's an "IRI user" and not just a facility user).
 - Should not forestall direct access between users and facilities.
- Facilities can evolve their metrics with DOE/PM guidance to support IRI alignment.
- Establish a communication/user engagement office.
- Establish structures for funding and sustainability, overcoming the hurdle of self-interest weighing down the interests of the whole.

Next Steps:

- Establish the IRI charter for each of the governance bodies (guiding principles, include cycles/cadence, releases). Helps adopters know how to change to become part of IRI.
 - Determine the points of contact for this activity and build out the initial (virtual) organization levels and people.
- Priority for the governance group should list early and long-term priorities to start clarifying activity: offer small/clear victories to increase buy-in and adoption and prevent errors in thinking this is "everything for everyone."
 - Evaluate and estimate required funding to deliver a minimum level of operation and / or integration.
 - Develop a multi-annual roadmap for activities and funding on a year-by-year basis.
 - Develop a communications and stakeholder engagement plan (including outreach to users and technical people, sponsors, etc., and a website).
- Technical activities
 - Set up a WG on measures of success and metrics for IRI.
 - Set up a WG to identify successes and lessons learned from ESnet/WLCG-OSG/IETF/EOSC/MPI-OpenMP approaches.
 - Set up a WG to investigate a maturity model that will help understand how ready each facility is for the IRI, and what each facility needs to do next.
 - Define outcomes for all of the above.
 - Prototype Findable, Accessible, Interoperable and Reusable (FAIR) multidisciplinary repositories.¹
 - Prototype federated/sustained login solutions.

Research Opportunities:

- Specifications, reference implementations for interoperable/standard systems
- Programming language hooks/improvements to support IRI campaigns
- Use of AI/ML methods for optimizing data placement and discovery
- Data storage/management/processing in an IRI; at scale/distributed

¹ Explore building on similar efforts such as <https://dataverse.org/> and DOE programs that have specifically requested FAIR repositories. Examples are <https://microbiomedata.org/> and <https://workflowhub.eu/>.

- Levels of portability (perfection is hard, but pragmatic steps)
- Co-scheduling distributed and federated resources
- Automation of experiments and workflows: edge-to-exascale
- Cyber-security questions (including federated/sustained logins)
- Human centered design questions (monitoring, control, etc.; human-computing interaction) for research communities working with an IRI; key user experience questions
- How to abstract data locality to users while maintaining performance?
- Reproducibility/provenance/FAIR and meta-data R&D
- Better co-design of facilities' data and data-services and computing (edge-to-exa)

Appendix R

IRI ABA DESIGN PHASE: FOCUS AREA ON COMPARING AND CONTRASTING THE PATTERN BLUEPRINTS

IRI ABA Design Phase: Cross-Cutting Focus Area on Comparing and Contrasting the Pattern Blueprints

Executive Message:

The DOE Integrated Research Infrastructure (IRI) enables the seamlessly integrated use of multiple scientific user facilities to allow scientists to perform research that otherwise could not be done and is a key enabling capability for solving the grand challenges facing the nation. IRI will advance scientific discovery and impact by providing the integrated capabilities for scientists to use advanced, high-powered instruments for cutting-edge research. It provides a common framework for scientists to schedule experiments at IRI facilities, take data, analyze them, and preserve and share the scientific results, better connecting people to science through computing.

What's Unique About Each Pattern?

Time sensitive patterns: These patterns include workflows that have time critical/sensitive requirements for real-time or near real-time results, for example, experiment steering, near real-time event detection, and deadline scheduling to keep up with data production.

- *Real-time access to resources:* Workflows have a time constraint for how science and measurements are performed.
- *Resource co-operations:* Real-time scheduling and synchronization across facilities is required to provide guaranteed networking, compute, and instrument resources.
- *Cybersecurity and federated access:* Federated ID is necessary to enable these workflows, with a level of service that meets time constraints and enables automation.
- *User experience:* Users need easy processes and timely execution to reserve, launch, tweak, and monitor resources.
- *Workflows, Interfaces, and Automation:* Support for rapid prototyping/coding.
- *Portable / Scalable Solutions:* Performant software tuned for architecture should be able to deliver results within *real-time* requirements; as a result, this software may not be able to "run anywhere," depending on available architecture at facilities.

Data integration-intensive patterns: These patterns include the analysis of combined data from multiple sources that may include data from multiple sites, experiments, and/or simulations. They require tracking metadata and provenance for reproducible science and interactive analysis of large-scale data.

- *Resource co-operations:* Workflows need to acquire many possibly scarce resources at one time.
- *Cybersecurity and federated access:* Common complex-wide federated ID system where workflows and agents are trusted at the same level as humans.
- *Workflows, interfaces, and automation:* Common abstraction Application Programming Interfaces (APIs) are needed for workflows and automation across complexes.
- *Scientific data lifecycle:* A key requirement is data services that track all aspects of the scientific process, from conception of the scientific question, to theory/modeling/simulations, to testing with experimental and observational data, to data

- processing/reduction, to data analysis interpretation, to publication and sharing. This also includes shared metadata and provenance tracking, including software provenance.
- *Portable / scalable solutions*: Software needs to run across different various resources.

Long-term campaign patterns: These patterns are characterized by sustained access to resources at scale over a longer time period, to accomplish a well-defined objective. Also important are long-term robustness, reproducibility, and reliability, which would potentially involve significant logistical planning. Examples include sustained simulation production and large data (re)processing for collaborative use.

- *Resource co-operations*: Approaches to managing IRI resources are needed that span the program offices and support research over decades, including allocations and curated data.
- *Cybersecurity and federated access*: Cybersecurity should be viewed from a strategic, long-term perspective rather than dealing tactically with disparate responses. An intentional cybersecurity infrastructure is needed that supports all the DOE Office of Science facilities and removes barriers between them.
- *Workflows, interfaces, and automation*: A long-term approach to managing software stacks is required, including workflows and everything under them.
- *Scientific data lifecycle*: Reproducibility, usability, and reuse of data with appropriate retention and curation should be supported.
- *Portable / scalable solutions*: A key requirement is support for the ongoing porting/developing/updating of software and workflows over life cycles of research campaigns.

What's Common Across the Patterns?

- **Resource co-operations**: Schedule and allocate instrument, computing, networking, and storage resources across the complex in a common manner. Establish common API(s) for resource allocation, launching jobs, monitoring job status, and accessing data, and streaming interfaces to large-scale computing resources.
- **Cybersecurity and federated access**: New complex-wide federated access tools and policies, cybersecurity and federated access as a service via APIs, strategic, cross-IRI view of cyber security policies and controls.
- **User experience**: Focus on bridging domain science and computational capabilities with easy-to-use, reliable, and repeatable tools.
- **Workflows, interfaces, and automation**: Interoperable workflow solutions; common base APIs and standards upon which domain or applications specific workflows may be created.
- **Scientific data life cycle**: Sustainable data repositories, data curation and search tools, and practices to help adopt Findable, Accessible, Interoperable, and Reproducible (FAIR) data practices and clear and understandable data policies and expectations across the complex. Software should be included in the curation process.
- **Portable / scalable solutions**: Allocations and accounts that span resources. Easy portability across sites and resources, and scalability for performance.

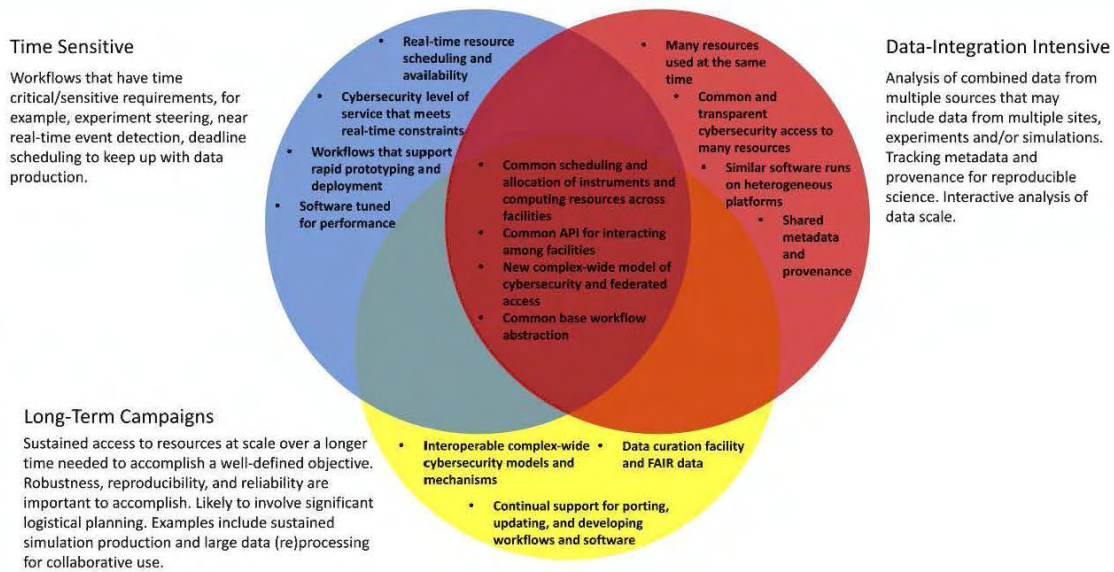
What's Next?

IRI should identify an early win use case for each pattern, for example.

- **Time sensitive:** real-time computing for experiment steering and automation at experimental facilities, such as the BES light sources.
- **Data-integration intensive:** combining data from multiple large-scale simulations for scientific discovery, such as at HEP and NP facilities, or across the NSRCs.
- **Long-term campaigns:** aligning facility resource commitments with campaign requirements, e.g., facility resource allocations, schedules, and upgrades with campaign schedules.

IRI should begin efforts to realize the following in support of these early win use cases:

- Implement capabilities (both technical and policy) for on-demand use of large-scale computing facilities for time-sensitive science.
- Develop common APIs for resource co-operations across computing (compute, storage, and network resources) and other facilities. (Research opportunity)
- Research and develop a common base workflow abstraction layer and reference implementations. (Research opportunity)
- Explore and realize new cybersecurity and federated access models that accommodate complex-wide IRI utilization. (Research opportunity)
- Evaluate and adopt/modify policies for allocations and resource utilization across facilities.
- Establish a data curation facility, and future plans for data curation services, for DOE SC.
- Build a community of practice around patterns and use cases for IRI across DOE SC.
- Create a governance mechanism to guide priorities for IRI.



Appendix S

COMPENDIUM OF SC AND NATIONAL REPORTS RELEVANT TO AN INTEGRATED RESEARCH INFRASTRUCTURE/ECOSYSTEMS APPROACH, AS OF FEBRUARY, 2022

**Compendium of SC and National Reports
Relevant to an Integrated Research Infrastructure/Ecosystems Approach**

(AC reports, workshop reports, requirements reviews, other agency and Office of Science and
Technology Policy (OSTP) reports ca. 2015–2021) as of February 2022

Program	Report and link	Notes
ASCR	<p>ASCAC/Hey 2020. <i>Opportunities and Challenges from Artificial Intelligence and Machine Learning for the Advancement of Science, Technology, and the Office of Science Missions</i></p> <p>https://www.osti.gov/services/purl/1734848/</p>	<p>“(Finding A) The growing convergence of AI, Data, and HPC provides a once in a generation opportunity to profoundly accelerate scientific discovery, create synergies across scientific areas, and improve international competitiveness. PDF page 19. Science and computing are now in an era of post-Moore’s Law silicon technologies and there is an urgent need for a sea-change in the productive use of increasingly complex/heterogeneous systems, and in the seamless integration of data and computing resources. There are also major challenges in the management, reduction, visualization, provenance, and curation of the scientific Big Data generated at scale by DOE’s most advanced facilities.”</p> <p>PDF page 19. “The combination of ML, high performance computing (HPC), and advanced data acquisition and handling will uncover a range of opportunities for breakthrough science – allowing the analysis of huge datasets, the exploration of enormously complex parameter spaces and the discovery of extremely subtle effects, leading to unforeseeable discoveries that will benefit the nation and, ultimately, the world.”</p>
ASCR	<p>2020 ASCAC ECP Transition Report https://science.osti.gov/-/media/ascr/ascac/pdf/meetings/202004/Transition_Report_202004-ASCAC.pdf</p>	<p>For instance, from recommendation A.2: “[T]he challenges and needs of non-ASCR user facilities can be quite different from ASCR’s current portfolio: data acquisition rates and needs for persistent storage, are increasing exponentially, even as there is increasing interest in performing significant amounts of computation on the data, e.g., for AI.”</p>
ASCR	<p><i>AI for Science: Report on the Department of Energy (DOE) Town Halls on Artificial Intelligence (AI) for Science, 2020</i> https://www.osti.gov/biblio/1604756</p>	<p>Large focus on integration needs, e.g.:</p> <p>“International leadership in AI over the coming decade will hinge on an integrated set of programs across four interdependent areas—new applications, software infrastructure, foundations, and hardware tools and technologies, feeding into and informed concurrently by DOE’s scientific instrument facilities and by DOE’s</p>

		<p>leadership class computing infrastructure.”</p> <p>“A set of integrated new AI workflow frameworks and exemplar applications will be needed to evaluate emerging AI architectures from edge SoCs to HPC data centers.”</p>
ASCR + other SC programs	ASCR ESnet requirements review reports	<p>Contain a plethora of priority use cases most of which are directly relevant to IRI/ecosystems:</p> <ul style="list-style-type: none"> • 2021 <i>HEP-ESnet Network Requirements Review Report</i>, https://escholarship.org/uc/item/78j3c9v4, or, https://science.osti.gov/-/media/hep/pdf/Reports/2021/2020-HEP-ESnet-Network-Requirements-Review-Report.pdf • 2020 <i>Nuclear Physics Network Requirements Review: One-Year Update</i>, https://escholarship.org/uc/item/4sf7n3pc • 2019 <i>NP Network Requirements Review Report</i>, https://www.es.net/assets/Uploads/20200505-NP.pdf • 2015 <i>BER Network Requirements Review Report</i>, https://www.es.net/assets/Hester/RequirementsReviews/BER-Net-Req-Review-2015-Final-Report.pdf • 2014 <i>BES Network Requirements Review Report</i>, https://www.es.net/assets/Hester/RequirementsReviews/BER-Net-Req-Review-2015-Final-Report.pdf • 2014 <i>FES Network Requirements Review Report</i>, https://www.es.net/assets/pubs_presos/FES-Net-Req-Review-2014-Final-Report.pdf
ASCR	<p><i>Report of the DOE Workshop on Management, Analysis, and Visualization of Experimental and Observational Data – The Convergence of Data and Computing</i>, 2015</p> <p>https://escholarship.org/uc/item/3vf1w91z, or https://www.osti.gov/biblio/1525145/</p>	<p>Report contains use cases/case studies from the programs and findings and recommendations.</p> <p>Some findings:</p> <ul style="list-style-type: none"> • "Specifically, the science use cases reveal a trend towards the convergence of data and computing: data- and compute-centric needs and opportunities are increasingly intertwined, interrelated, and symbiotic. " • "Meeting the challenges of the explosion of data from EOS projects requires computational platforms, networking, and storage of greater capacity and lower latency, along with software infrastructure suited to their needs. However, existing HPC platforms and their software tools are designed and provisioned for high-concurrency HPC workloads, single-project data products, and comparatively simpler data needs. The result is a significant gap between the needs of EOS projects and the current state of the art in computational and software

		<p>capabilities and resources.”</p> <ul style="list-style-type: none"> • “EOS projects increasingly rely on low-latency, fast turnaround resource response to meet datacentric needs.” • “Collaboration and sharing of data, tools, and methodologies are central to modern EOS projects, yet there is insufficient infrastructure to facilitate such interactions.” <ul style="list-style-type: none"> • “EOS projects are impeded due to significant “data lifecycle” needs that are largely unmet.” [Data Lifecycle refers to all stages of data collection, movement, processing, analysis, management, curation, and sharing] • “The highly specialized nature of skills and expertise in the data sciences and their application to EOS problems raises concerns about workforce training, development and retention.” <p>Some recommendations:</p> <ul style="list-style-type: none"> • Rec 1c: “Cultivate multidisciplinary teams and programs that focus on software solutions to data centric challenges that are broadly applicable beyond a single EOS project.” • Rec 2: “Evolve HPC computational facilities to include focus on the needs of the EOS community.” • Rec 3a: “Develop a systematic, end-to-end understanding of time-critical EOS needs that includes the appropriate metrics and that takes into account human-in-the-loop scenarios.”
ASCR	<p><i>The Future of Scientific Workflows</i>, DOE-NNSA workshop report, 2015 https://science.osti.gov/-/media/ascr/pdf/programdocuments/docs/workflows_final_report.pdf</p>	<p>Workshop objectives are focused on emerging integrative use cases for High-Performance Computing (HPC) and “distributed-area instruments and computing (DAIC)”. Organizers and participants are the key players in science workflows across DOE and National Science Foundation (NSF) space.</p>

ASCR	<p><i>The Future of Scientific Workflows</i>, Deelman et al. 2017, International J of HPC, DOI: 10.1177/1094342017704893.</p>	<p>Excerpts:</p> <p>The community identified four main research areas for future workflow development:</p> <ul style="list-style-type: none"> • Design of <u>task coupling</u> and data movement between workflow tasks. Scalable and robust control and data flow and the need for efficient and portable migration of heterogeneous data models across tasks. • <u>Programming and usability</u>. programming models, design patterns, the user interface, task communication, and portability. • <u>Monitoring</u>: anomaly detection, gracefully recovery from errors. • <u>Validation of results</u>: reproducibility, provenance capture. <p>“It is important to understand and classify various workflows and workflow needs through user studies. Identifying common patterns for next-generation in situ and distributed workflows is needed to address programmability and usability concerns.”</p>
------	-------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

ASCR (Facilities)	<p><i>DOE HPC Operational Review (HPCOR) 2014: Enabling Data-Driven Scientific Discovery at HPC Facilities</i> https://www.osti.gov/services/purl/1163236</p>	<p>“On June 18-19, 2014 representatives from six DOE HPC centers met in Oakland, CA at the DOE High Performance Operational Review (HPCOR) to discuss how they can best provide facilities and services to enable large-scale data driven scientific discovery at the DOE national laboratories.”</p>
ASCR and other SC	<p><i>Exascale Requirements Review Crosscut Report</i> https://exascaleage.org/crosscut-report/</p>	<p>Contains the links to the individual program-specific reports.</p>

<p>ASCR and other SC</p>	<p>2019. <i>Background and Roadmap for a Distributed Computing and Data Ecosystem</i> https://doi.org/10.2172/1528707.</p>	<p>“The National Laboratory Research Computing Group (NLRCG) and the Advanced Scientific Computing Research (ASCR) program office jointly established the Future Laboratory Computing Working Group (FLC-WG) with a charter to identify the benefits and obstacles in creating and operating a DOE/SC wide federated Distributed Computing and Data Ecosystem (DCDE).”</p> <p>Pg 8. “We envision the creation of a DOE Office of Science (SC) wide federated Distributed Computing and Data Ecosystem (DCDE) which comprises tools, capabilities, services and governance policies to enable researchers to seamlessly use a large variety of resources (i.e., scientific instruments, local clusters, large facilities, storage, enabling systems software, and networks) end-to-end across laboratories within the DOE environment.”</p>
<p>SC</p>	<p>2020. Office of Science User Facilities. Lessons Learned from the COVID Era and Visions for the Future https://www.osti.gov/biblio/1785683/</p>	<p>Pg 23. “The greatest obstacle to effective virtual communities involves the ability to share information seamlessly and securely among geographically dispersed participants.”</p> <p>Pg 24. “The ability to find, access, and reuse data stored in pools of geographically and logically distinct storage resources is critical to ensuring a high level of scientific productivity for staff and users. ... To support these activities, user facilities will require data management systems that integrate all</p>

		<p>data, enabling researchers to use a common set of tools to work across the broadest range of applications. Artificial distinctions between data structures made based on their origin lead to redundant efforts and impede scientific progress. Support is required for all needed data types and structures.”</p> <p>“Facilities may also be able to deploy a distributed network of connected and interoperable computing resources that enable all scales of computing, data exploration, and analysis. Seamlessly connecting a user with data and computing enables more uniform and egalitarian data exploration and analysis capabilities.”</p> <p>“With collaboration among all its user facilities, DOE SC is in a position to facilitate all aspects of the data life- cycle across its facility complex, including simulations, experiment design, data generated at scientific instruments, data analysis, and data archiving for future use. Data management tools that provide transparent data movement among these facilities would enable users to log in from anywhere to focus on the science.”</p>
BER	<p><i>Breaking the Bottleneck of Genomes: Understanding Gene Function Across Taxa Workshop</i>, 2019 https://www.osti.gov/servlets/purl/1616527/</p>	<p>Pg 21. “Discovering new gene functions and accurately transferring these annotations across taxa are both experimental and computational challenges. ... Consequently, advances in computational tools are urgently required to automate the inference of gene function from diverse data and interactive databases that maintain and propagate accurate gene annotations across taxa.”</p> <p>“One opportunity for discovering new gene functions and rapidly increasing the quality of genome annotations is a proper computational infrastructure, with community coordination and appropriate experimental data (see Fig. 3.2, this page). This platform could integrate seamlessly with (or be a part of) existing U.S. Department of Energy (DOE) computational resources, including the Systems Biology Knowledgebase (KBase), Joint Genome Institute (JGI), Environmental Molecular Sciences Laboratory (EMSL), and National Energy Research Scientific Computing Center (NERSC), as well as the National Center for Biotechnology Information (NCBI) supported by the National Institutes of Health, Protein Data Bank (PDB) managed by the Research Collaboratory for Structural Bioinformatics, and the UniProt database.”</p>

BER	<p><i>Genome Engineering for Materials Synthesis Workshop Report, 2019.</i></p> <p>https://genomicscience.energy.gov/wp-content/uploads/2021/09/GEMS_Report_2019.pdf</p>	<p>Discusses emerging computational tools and application of ML for various aspects of synthesis research and development; does not discuss the details of compute/data infrastructure needs per se.</p>
BER	<p><i>Atmospheric Radiation Measurement (ARM) User Facility ARM Mobile Facility Workshop Report, 2019.</i></p> <p>https://science.osti.gov/-/media/ber/pdf/community_resources/2019/ARM_Mobile_Facility_Workshop_Report.pdf/</p>	<p>Page iv (Pdf page 8). “Dedicated site-focused modeling activities, like Large-Eddy Simulation (LES) ARM Symbiotic Simulation and Observation (LASSO), should be used to bridge observations with efforts to improve larger-scale [earth system models] ESMs. “</p> <p>Page 31. PDF 41—image caption—“Improving model parameterizations based upon the understanding of processes gained through AMF observations is a high priority. There is potential to pair AMF deployments with model studies performed by DOE’s Energy Exascale Earth System Model (E3SM). “</p> <p>Page 31 text: “As a means to enhance the effectiveness of using AMF data by modelers, one could automate the production of large-scale forcing from the analyses or re-analyses of these models so that they are available as early as possible during the campaign. These analyses could be of higher quality if the AMF observations, particularly for the radiosondes, could be ingested into these analysis models in real time.”</p>
BES	<p>Brochure. <i>BES Producing and Managing Large Scientific Data with Artificial Intelligence and Machine Learning—Enabling transformative advances at BES Scientific User Facilities, 2020</i></p> <p>https://science.osti.gov/-/media/bes/pdf/reports/2020/AI-ML_Report.pdf</p>	<p>PDF page 2</p> <p>“Autonomous control of experimental systems promises to open the study of problems previously considered impossible. Automating the entire experimental workflow— instrument setup and tuning, sample selection and synthesis, measurement, data analysis and model-driven data interpretation, and follow-up experimental decision-making—will bring about revolutionary efficiencies and research outcomes. “</p> <p>PDF page 3</p> <p>“AI/ML-based methods are needed to efficiently search large, complex parameter spaces in real time and to predict the health and failure of instruments that operate at high-power sources and the experiments that are run on those instruments. Such capabilities could dramatically reduce facility tuning time and downtime, improve facility performance, and maximize the productivity of the BES scientific user facilities.”</p>

BES	<p>Full Report BES AIML Roundtable on Producing and Managing Large Scientific Data with Artificial Intelligence and Machine Learning, 2020 https://science.osti.gov/-/media/bes/pdf/reports/2020/AI-ML_Report.pdf/</p>	<p>Page 34 (PDF page 44) “AI/ML models are fundamentally linked to the datasets on which they are trained, and data infrastructure needs are ubiquitous in AI/ML workflows. ... nearly every topic covered during the roundtable will face challenges relating to data workflows for training, testing, and deployment of models. AI/ML methods will eventually support rapid data processing at HPC facilities to enable quasi-real-time feedback on experiments and observations. These advances are fundamental to the PROs identified in this BES AI/ML roundtable. “</p> <p>PDF page 45. “Future computing environments that can address these challenges will likely be heterogeneous, consisting of GPU accelerators, possibly in conjunction with FPGAs, application-specific integrated circuits (ASICs), and emerging hardware custom designed for deep learning workloads.”</p>
BES	<p>AIML companion doc: <i>Facilities’ Current Status and Projections for Producing and Managing Large Scientific Data with Artificial Intelligence and Machine Learning</i> https://science.osti.gov/-/media/bes/pdf/reports/2020/AI-ML_Companion_Document.pdf</p>	<p>Page 30 (PDF 38) “There are ongoing efforts to integrate capabilities across the SUFs via AI/ML, networking, and advanced math. Coordination and execution are highly collaboration-based and mostly fall under guidance from the Energy Sciences Network (ESnet) and the CAMERA project.”</p> <p>“Current AI/ML methods are most effective in the regime of supervised learning, for which access to training datasets is a critical requirement. BES, ESnet, and NERSC facilities currently lack standardized tools to capture, label, and share such datasets broadly within their respective user communities or the wider research community.”</p>
BES	<p><i>Basic Energy Sciences Exascale Requirements Review</i>, 2017 https://science.osti.gov/-/media/bes/pdf/reports/2017/BES-EXA_rpt.pdf/</p>	<p>“BES and the ASCR facilities are experiencing a pressing need to mature their capabilities in data science. Improvements and new capabilities at BES facilities are creating challenges that the community is not prepared to address. These include unprecedented growth in data volume, complexity, and access requirements; the need for curation of the massive amounts of data that are retained; and integration of diverse datasets from different experiments to enable new scientific conclusions. Efficient and effective use of BES facilities requires real-time access to ASCR HPC facility-class resources to support streaming analysis and visualization to guide experimental decisions.”</p>

FES	<p>FESAC, 2020. <i>Powering the Future: Fusion & Plasmas</i></p> <p>https://science.osti.gov/-/media/fes/fesac/pdf/2020/202012/FESAC_Report_2020_Powering_the_Future.pdf/</p>	<p>PDF page 12. “A vital part of the program is the continued development of validated models at a range of complexities and experimental fidelities, along with the predictive integrated modeling capabilities that utilize them. Creating such models will require continued close partnership between FES and ASCR to fully leverage US investments in high performance computing, including coming exascale machines.”</p>
FES	<p>FESAC 2018. <i>Transformative Enabling Capabilities for Efficient Advance Toward Fusion Energy</i></p> <p>https://science.osti.gov/-/media/fes/fesac/pdf/2018/TEC_Report_1Feb20181.pdf</p>	<p>Discusses needs for exascale and HPC, “integrated data analysis,” simulation.</p>
FES	<p>2019 <i>Advancing Fusion with Machine Learning Research Need Workshop</i> cosponsored by FES and ASCR</p> <p>https://science.osti.gov/-/media/fes/pdf/workshop-reports/FES_ASCR_Machine_Learning_Brochure.pdf</p> <p>https://science.osti.gov/-/media/fes/pdf/workshop-reports/FES_ASCR_Machine_Learning_Report.pdf</p>	<p>PDF page 43 (of 66) discusses issues relevant to IRI-ABA under Priority Research Opportunity #7: Fusion Data Machine Learning Platform</p>
HEP	<p>2018 <i>HEP Portfolio Review: Report of the LHC Subpanel</i></p> <p>https://science.osti.gov/-/media/hep/hepap/pdf/Reports/HEP_Portfolio_Review-Report_LHC_Subpanel.pdf</p>	<p>Discusses future computing needs and use of ASCR facilities</p>
HEP	<p><i>Computing in High Energy Physics. Report from the Topical Panel Meeting on Computing and Simulations in High Energy Physics</i>, 2014.</p> <p>https://science.osti.gov/-/media/hep/pdf/files/BannerPDFs/Computing_Meeting_Report_final.pdf/</p>	<p>Workshop in 2013 included a DOE ASCR session, and discussed computing, new strategies in data, software, etc. “Evolution of data archiving, data-intensive computing, and storage will drive new computational strategies”</p>

NP	<p>NASEM. <i>An Assessment of U.S. Based Electron-Ion Collider Science</i>, September 2018 https://science.osti.gov/-/media/np/pdf/NASAnAssessmentofUSBasedElectronIonColliderScience.pdf</p>	<p>Has section on Electron Ion Collider (EIC) and Advanced Scientific Computing; large-scale simulation of lattice Quantum Chrono-Dynamics (QCD).</p> <p>Page 112 (PDF page 127). "An EIC will be among the first facilities to come online in the era of exa-scale computing, an era that will see unprecedented integration of computing in the collider and experiments. These developments, combined with continued advances in machine learning and other areas, will open up opportunities for truly new approaches to nuclear physics experiments and analyses of scale, perhaps removing altogether the current distinction between acquiring the data from the instruments and their subsequent analysis."</p>
Other agencies: NASA	<p><i>NASA SMD's Strategy for Data Management and Computing for Groundbreaking Science 2019–2024</i></p> <p>https://science.nasa.gov/science-red/s3fs-public/atoms/files/Knezek%20SDM%20WG%20Strategy%20Update%20to%20APAC%20March2020.pdf</p>	<p>Guiding principles and plan for NASA SMD.</p>
Other agencies: NSF	<p><i>Transforming Science Through Cyberinfrastructure: NSF's Blueprint for a National Cyberinfrastructure Ecosystem for Science and Engineering in the 21st Century</i> https://www.nsf.gov/cise/oac/vision/blueprint-2019/</p>	<p>Consists of five vision/blueprint documents (thus far) for a more holistic approach to cyberinfrastructure for science:</p> <ul style="list-style-type: none"> • OAC Vision & Blueprint: Overview and Computational Ecosystem (As of April 2019) • OAC Vision & Blueprint: Coordination Services (As of November 2019) • OAC Vision & Blueprint: International Research & Education Network Connections (As of November 2019) • OAC Vision & Blueprint: Data & Software Cyberinfrastructure • OAC Vision & Blueprint: Cyberinfrastructure Learning & Workforce Development <p>The website also cites several relevant RFIs (with posted RFI submissions) and other relevant reports.</p>

Other bodies: AAAC	<i>Report of the Astronomy and Astrophysics Advisory Committee</i> , March 15, 2021 https://www.nsf.gov/mps/ast/aaac/reports/annual/aaac_2021_report.pdf	Includes findings and recommendations related to NASA DOE-NSF coordination/collaboration in the areas of computing, data sharing, software, workforce, artificial intelligence / machine learning (AI/ML), cyberinfrastructure writ large.
Other agencies: OSTP	<i>National Strategic Computing Initiative Update: Pioneering the Future of Computing</i> , 2019 https://www.nitrd.gov/publications/National-Strategic-Computing-Initiative-Update-2019.pdf	Guiding interagency doc. (Goal 1 Enabling the Future of Computing Objective: Pioneer new frontiers of digital and nondigital computation to address the scientific and technological challenges and opportunities of the 21st century): “....At the same time, application workflows are evolving with new requirements that necessitate the integration of heterogeneous platforms, including those within a given architecture as well as network-centric and edge computing.”
Other agencies: OSTP	National Strategic Computing Reserve, 2021 https://www.whitehouse.gov/wp-content/uploads/2021/10/National-Strategic-Computing-Reserve-Blueprint-Oct2021.pdf?utm_medium=email&utm_source=govdelivery	Proposes cross-agency partnering on integrated/federated computing reserve approach built on experience with the COVID-19 HPC Consortium.
Other agencies: OSTP	National AI Research Resource https://www.ai.gov/nairtf/	Website provides links to the meetings and presentations. Much discussion of how to approach a national resource via some form of federated or integrated existing and new research infrastructure and data sources.
Other agencies: OSTP	<i>National Strategic Overview for Research and Development Infrastructure</i> , 2021 https://www.whitehouse.gov/wp-content/uploads/2021/10/NSTC-NSO-RDI-REV_FINAL-10-2021.pdf	The RDI report is notable for specifically identifying both research cyberinfrastructure and knowledge infrastructure as key categories on equal footing with experimental and observational infrastructure within a broader more-inclusive definition of IRI. In particular, computing and cyberinfrastructure are cited repeatedly throughout the report, recognizing the increasing prominence of these needs across science.