High Performance Data Facility: Status and Plans

July 22, 2024

Jefferson Lab

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

The concepts in this talk are works in progress. Feedback is valued and greatly appreciated

We are eager to discuss possible alignments, opportunities, & gaps

We would like to talk further at your meetings, organizations, etc.

Innovation Through Partnership

Our mission: To enable and accelerate scientific discovery by delivering state-of-the-art data management infrastructure, capabilities, and tools

The HPDF Project team leverages the strengths and complementarity of both labs:

- Decades of experience with scientific missions and user communities
- A shared understanding of resilient, distributed infrastructure that supports the data life cycle
- A shared commitment to the IRI initiative and ASCR ecosystem

The HPDF will be a first-of-its-kind SC user facility:

- A distributed operations model will be essential to long-term success and required performance levels
- Project structure is integrated with JLab and LBNL staff

HPDF is a DOE 413 Project announced in Oct. 2023

HPDF is essential to handle the staggering amounts of data

The DOE envisions a revolutionary ecosystem – the Integrated Research Infrastructure – to deliver seamless, secure interoperability across National Laboratory facilities

The 2023 IRI Architecture Blueprint Activity identified three broad science patterns that demand research infrastructure interoperability:

- Time-sensitive patterns
- Data-integration-intensive patterns
- Long-term campaign patterns

HPDF will enable analysis, preservation, and accessibility of the staggering amounts of experimental data produced by SC facilities

Our mission: To enable and accelerate scientific discovery by delivering state-of-the-art data management infrastructure, capabilities, and tools

HPDF Will Address SC Priority IRI Science Patterns

Drivers	IRI Patterns
Supporting data curation, repositories, and archives	
Supporting data processing and analysis pipelines	
Data federation, sharing, and collaboration	
Real-time streaming and processing	

Time-Sensitive Data Integration-intensive Long-Term Campaign

These patterns are seen widely in the larger community, in other parts of DOE, and outside

HPDF will Support Data Lifecycle Management

Data science requires curated and annotated data that adheres to FAIR principles, and data reuse will be a metric for HPDF. Office of Scientific and Technical Information (OSTI) services will complement HPDF to provide full life cycle coverage.

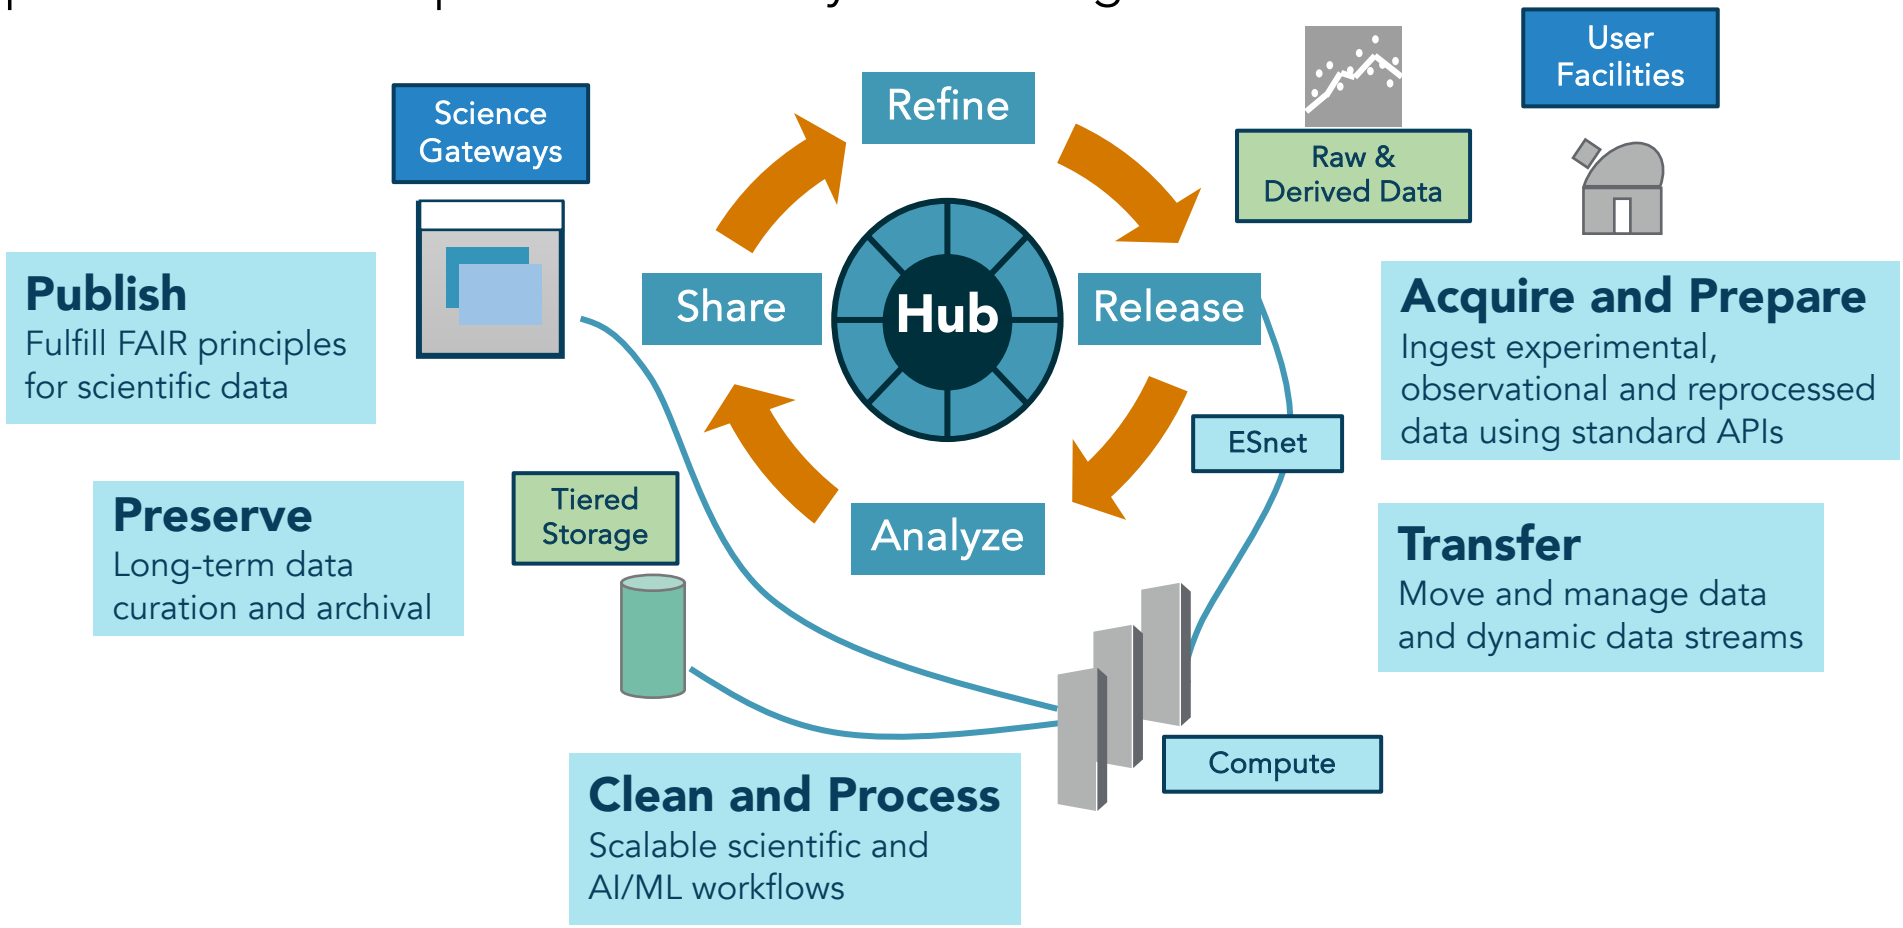

Key Facets of Data Requirements

Management – A dynamic and scalable data management infrastructure integrated with the DOE computing ecosystem

Capture – Dynamically allocatable data storage and edge computing at the point of generation

Staging – Dynamic placement of data in proximity to appropriate computing for reduction, analysis, and processing

Archiving – Extreme-scale distributed archiving and cataloging of data with FAIR principles

Processing – Resources for workflow and automation for processing and analyses of data at scale

Policy of data and providing collaborative environments around data are also critical

HPDF in the ASCR Ecosystem

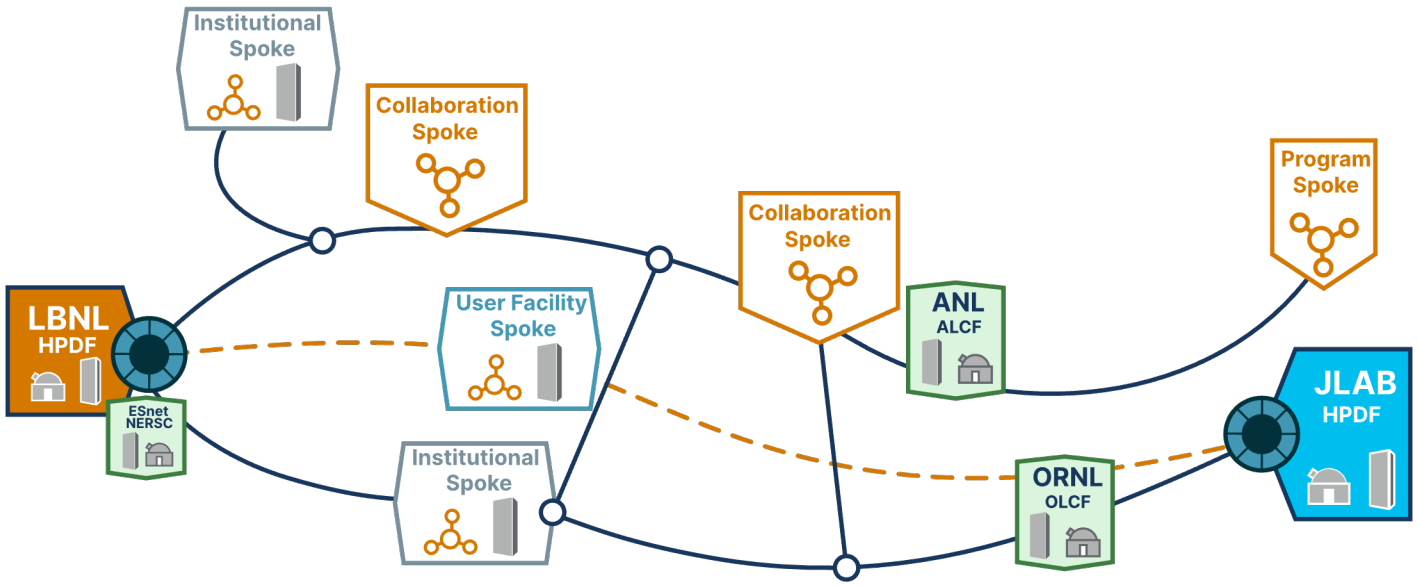

- Working with IRI and other ASCR facilities ensures a secure, high-performance mesh data fabric that enables data and workloads to flow freely
- The HPDF distributed infrastructure will be designed to maximize planned availability and resilience
- Partnering with Spoke sites will provide seamless data life cycle services to scientific users worldwide
- Pilot activities and partnerships will help refine the design as Hub software and hardware technology evolve and foster workforce development

HPDF — A Distributed Facility

Concept: HPDF is a distributed facility with a hub and spoke architecture.

- **Hub.** Data-centric infrastructure with high availability and performance, as well as geographically and operationally resilient active-active failover.
- **Spokes.** Distributed data-centric infrastructure to enhance HPDF access and support for science users and integrate distributed computing or storage resources.
- **Integration and Services.** Orchestration hardware, software, and services for data movement, storage and retrieval, and science workflow automation. These will use a mesh data fabric building on ESnet6 capabilities.

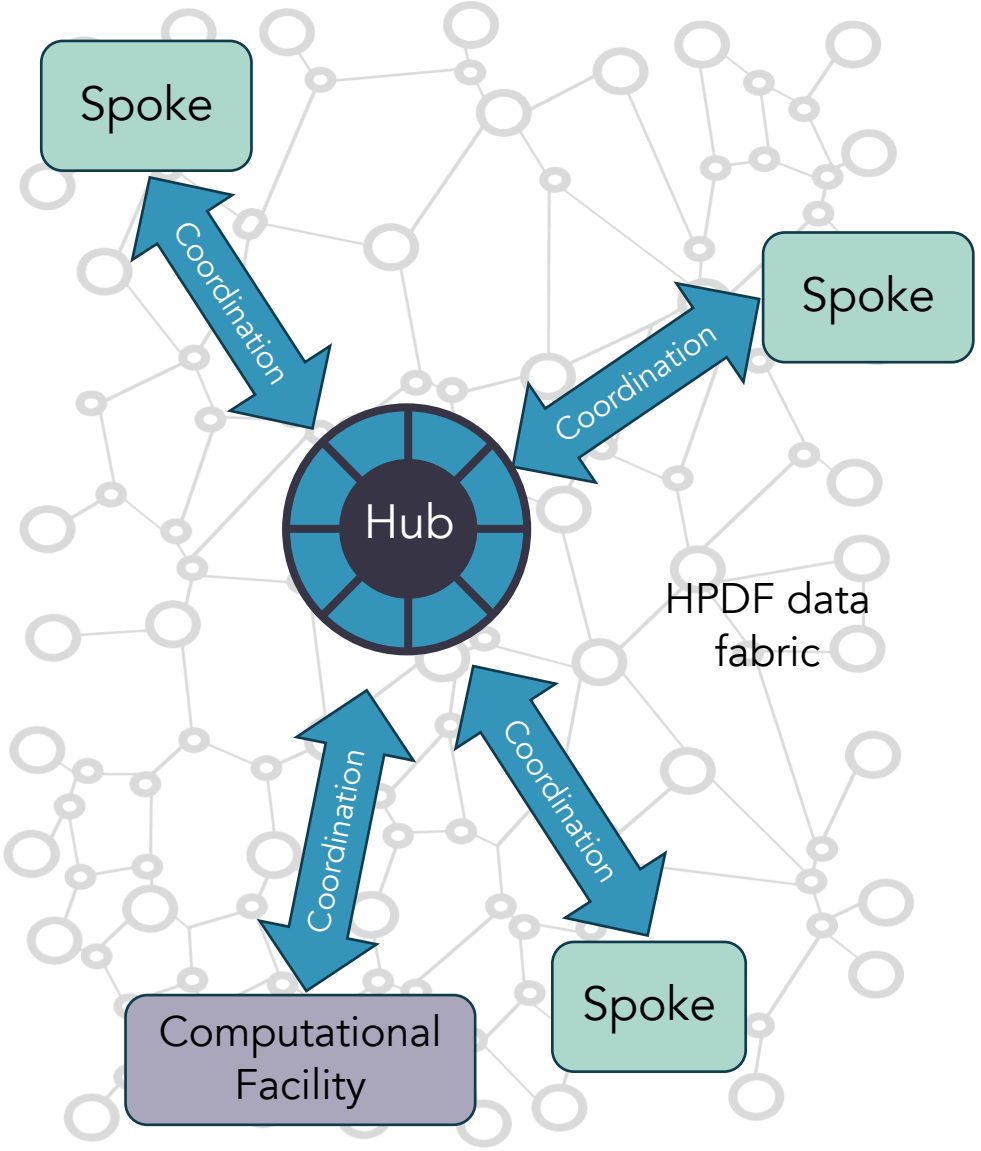

Technical Design — Core Capabilities

Hub Computing and Data Infrastructure

- High uptime
- Experiment-friendly availability
- Data-driven agility
- Support for new technologies
- Data storage, management, and interoperability
- Data preservation

Distributed Spoke Infrastructure

- User support
- Scientific application tailoring
- Hardware resources that mirror, supplement, or complement Hub resources
- Low-latency or high-bandwidth coupling of HPDF services to edge compute

Data-centric Orchestration of Hardware, Software, and Services

- High availability
- High-performance mesh data transport fabric
- Secure data paths
- Monitoring
- Orchestration

HPDF Architecture Stack

- Common APIs and data services to facilitate portability
- Distributed orchestration and execution layers
- Data transport, caching, communication, and monitoring built on ESnet6 capabilities
- Dynamic virtualized compute and storage ensuring portability between sites
- Cross-cutting components for security and monitoring
- Developed in partnership with IRI

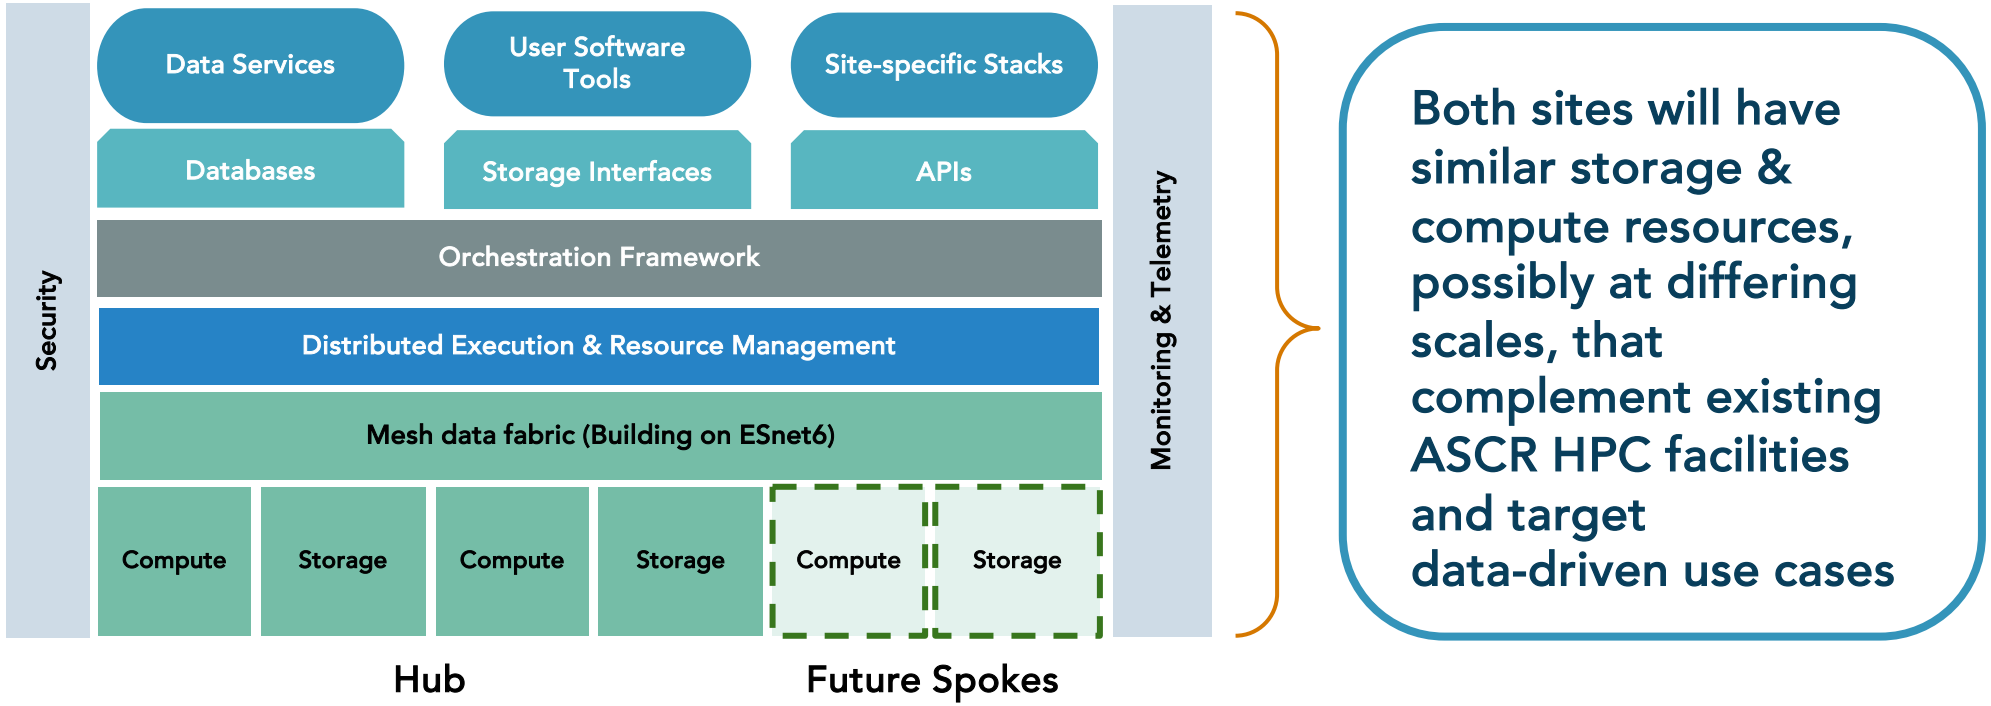

Both sites will have similar storage & compute resources, possibly at differing scales, that complement existing ASCR HPC facilities and target data-driven use cases

The HPDF Hub: Unique Hardware Capabilities

- Combines high availability, flexibility, and support of time-critical workflows
- Composable storage will be configured to limit the need to modify existing code
- A local archive will be available along with a federated data catalog of data archived elsewhere
- The data processing design is based on the concept of "standard units," hardware elements following well-defined architectures targeting specific use cases
  - Batch jobs, AI/ML intensive, streaming, real-time, and dynamic reconfiguration
  - A mix of CPU/GPU flavors to run existing optimized code
- The Hub will incorporate a range of standard units in a mix that meets the science needs yet can evolve over time
- This is not a one-size-fits-all approach; it allows tailoring to needs and lowers the barrier to HPDF use

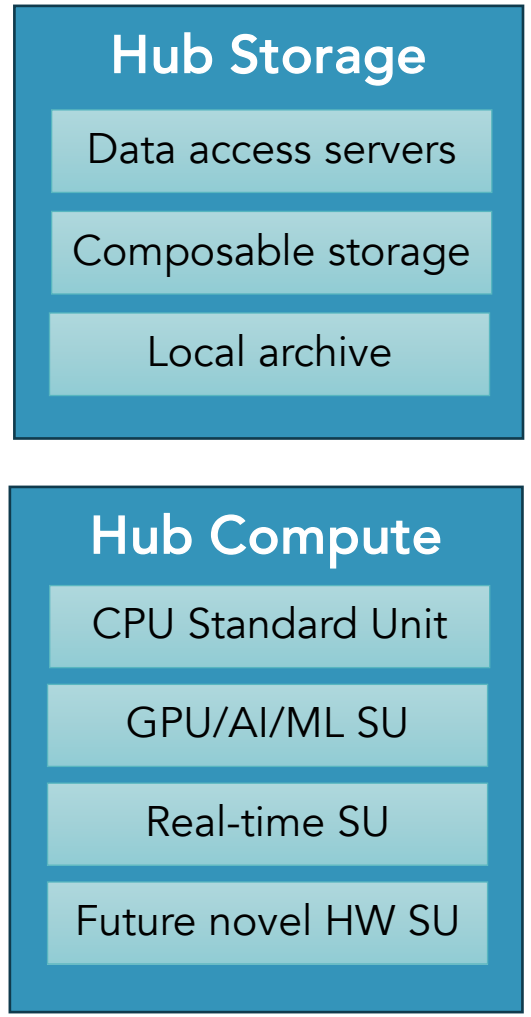

Software & Data Services Strategy: Innovation & Stewardship

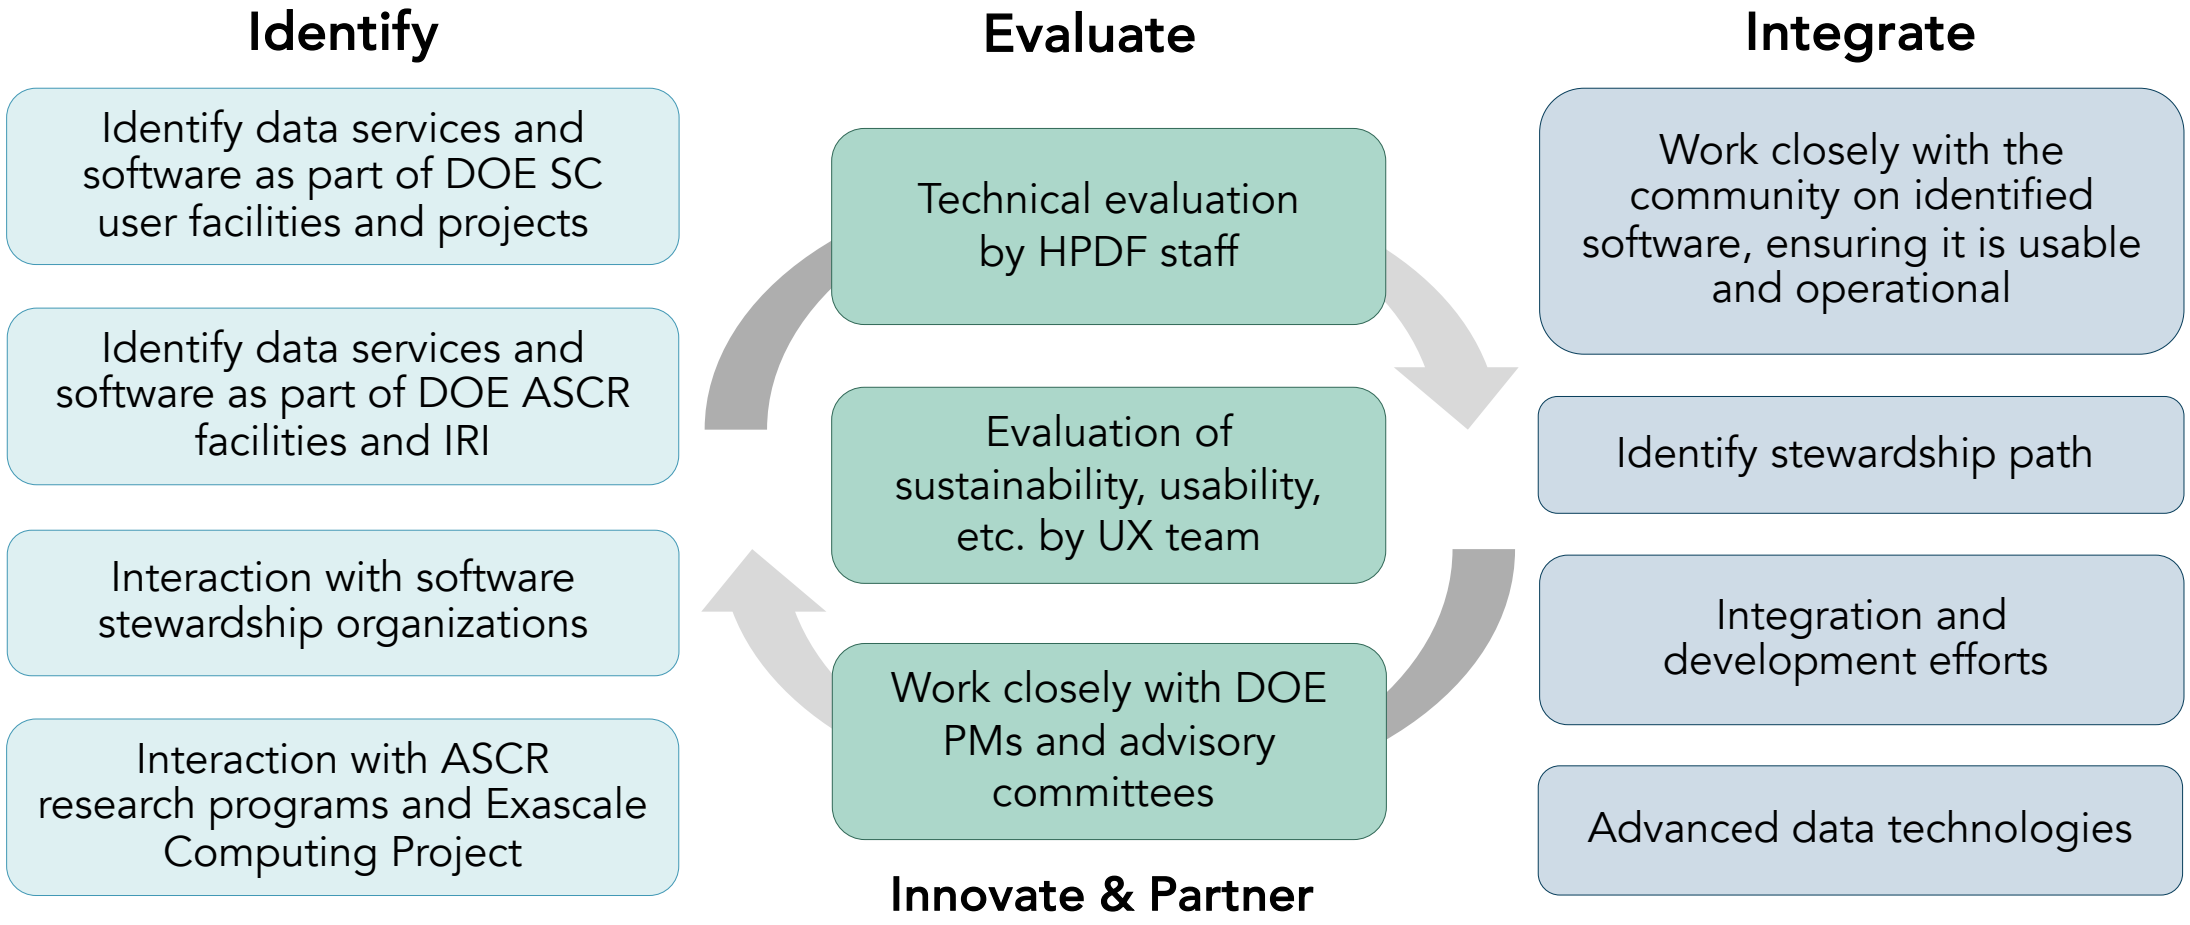

Coordination and collaboration with the governance for IRI and ASCR scientific software

Preliminary Results: Scalable Data Management Infrastructure Mapped to IRI Patterns

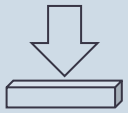

Data Capture & Storage

- Data replication & tiering
- FAIR data support, curate data with metadata
- Streaming data core & edge services co-dev with ESnet

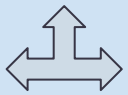

Data Management & Staging

- Techniques for data filtering, data scheduling, parallel stream processing
- Replication
- APIs for schedulers

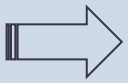

Programmable APIs

- APIs to services across entire data lifecycle
- Access through web-based APIs & Python/C++
- Interface to SF-API

Data Life Cycle Services

- Robust & reliable distributed data management layer
- Data analysis tools/services: user feedback, vendor/OSS engagement

Data Repository & Archiving

- Publication QA/QC pipeline, search tools, AI/ML dataset tagging
- Long-term storing, archiving, access, & discovery through web interface, DOIs & APIs

Data Analysis & AI

- Connect data to clusters/ clouds/HPC
- Integrated AI platform with uniform APIs
- Re-use & reproduce previous results

Time-Sensitive Data Integration-intensive Long-Term Campaign

User Experience Engagement: Core to Our Strategy & Plan

User research gives us a process to verify/validate our "intuition about what the user needs" (hypothesis) and convert it into action

Deep partnership model to serve user needs, mature data stewardship across SC, and develop a workforce

User support and data stewards will provide critical help to HPDF users to leverage resources effectively and efficiently, allowing us to address computation and data needs early

"Tale of caution; you get one shot at making things accessible. Simple, easy (Google-esque) interfaces. You get one shot at it, or you'll lose your PIs. Interface simplicity."

– Dan Jacobsen, ORNL/UT

HPDF Governance & Execution

HPDF Engagement modes

Users

Use HPDF hardware, software, and services to support their data lifecycle needs.

- Store and retrieve data
- Acquire, prepare, preserve, and publish
- Process/analyze data
- Leverage IRI capabilities to seamlessly use ASCR facilities

Partners

Partners provide software/services that contribute to the HPDF ecosystem.

Spokes

Long-term partnerships that support the data lifecycle needs for specified communities

A mutual commitment to align the teams, resources, and policies users depend on for maximum productivity

Spokes provide first level of service for the communities they serve and provide enhanced user experience and customizations

Co-evolve with Hub infrastructure as computing and data technology and solutions advance

Add value by providing community-specific software and services, user support, and possibly hardware resources mirroring, supplementing, or complementing Hub resources

Key Spoke Activities: High-level View

Spoke Activities

Collaborative Design: Work between Hub, Spoke, & focal communities on resources such as hardware, software, data policy, and data lifecycle needs.

Systems Operations: Responsible for the management and support of different types of operational tasks & resources (hardware, software, etc.).

Partnership Modalities: Spoke staff will contribute to HPDF partnership and governance model as part of systems and community operations activities.

Data Policy in Practice: Spoke staff will ensure appropriate collaborative data policies are put into practice in concert with a community.

Support: Spoke staff will offer front line user support that leverages their deep subject matter expertise in the scientific domains, instrumentation, or data products at hand.

Training & Outreach: Leading training & outreach activities for community specific workflows, specific HPDF resources, and outreach to the Spoke's community, or general advocacy of Open Science.

Summary

Next Steps

- Working toward CD-1: Conceptual technical design and scope and alternative analyses
  - Includes design of Hub and initial Spokes
- Community outreach
  - ✓ 6-way Light Sources meeting (Jan, in-person)
  - ✓ IRI Management Council (April, virtual)
  - ✓ Community Webinar (June, virtual)
  - ✓ FES PI Meetings (June, in-person)
  - ✓ HPDF/IRI workshop (July, in-person)
  - Monterey Data Conference (August, in-person)
  - Community Webinar (Q4 FY24, virtual)
  - Small-group interviews with groups identified through initial HPDF workshop (summer/fall, virtual)
  - Supercomputing '24 (November, in-person)

User & Community Engagement

Science engagements & partnerships are critical to success of HPDF

We are in the early conceptual design and community engagement (i.e., listening) phase.

We are eager to continue learning about your community's needs, challenges, science drivers, and so on.

Pointers to tools, reports, and so on are always helpful resources for our team.

Q&A

https://hpdf.science

Share thoughts & questions or request to be added to our mailing list via our form.
Answers will be provided via the website within a few weeks.