High Performance Data Facility: Status and Plans

June 11, 2024

Jefferson Lab U.S. DEPARTMENT OF ENERGY | Office of Science BERKELEY LAB

Outline

Vision & Background – Amber Boehnlein
Design Overview – Graham Heyes, Shane Canon
Science Engagement & Partnership – Lavanya Ramakrishnan
Spokes – Lavanya Ramakrishnan

Innovation Through Partnership

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

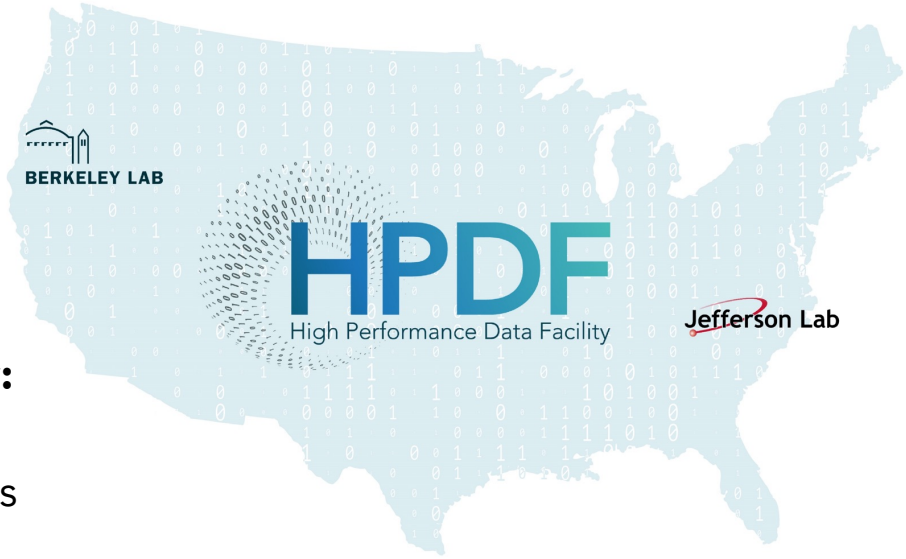

HPDF Governance & Execution

Meeting the Greatest Needs

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 in the ASCR Ecosystem

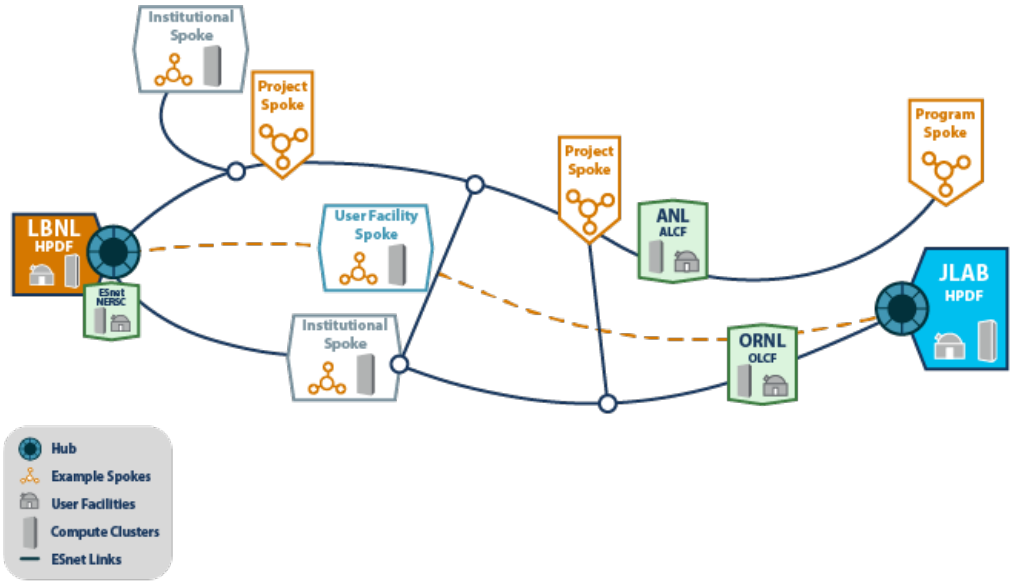

- Working with IRI and other ASCR facilities ensures a secure, high-performance mesh data fabric that enables data and workloads to flow freely
- The HPDF distributed infrastructure will be designed to maximize planned availability and resilience
- Partnering with Spoke sites will provide seamless data 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.

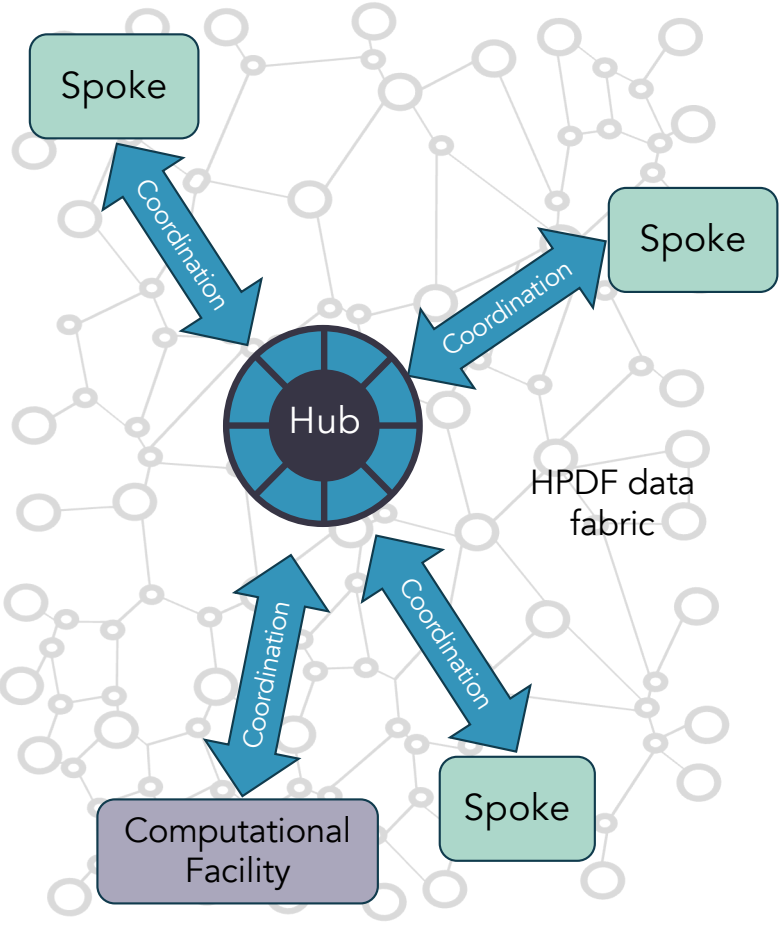

HPDF Hub Will Address Key Strategic Goals and Capability Gaps

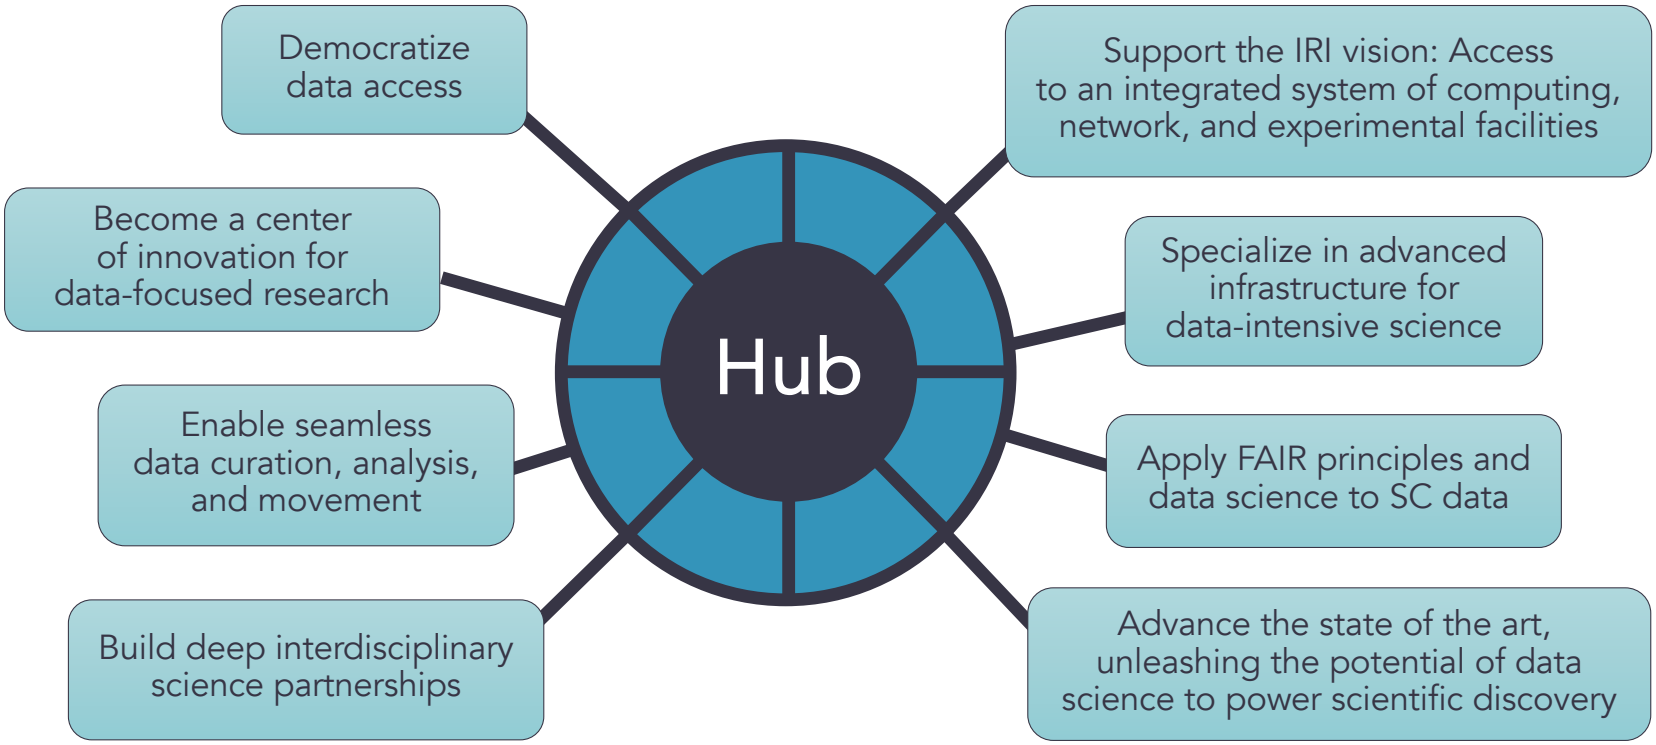

HPDF Will Set Standard for Data Life Cycle Management

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 services will complement HPDF to provide full life cycle coverage.

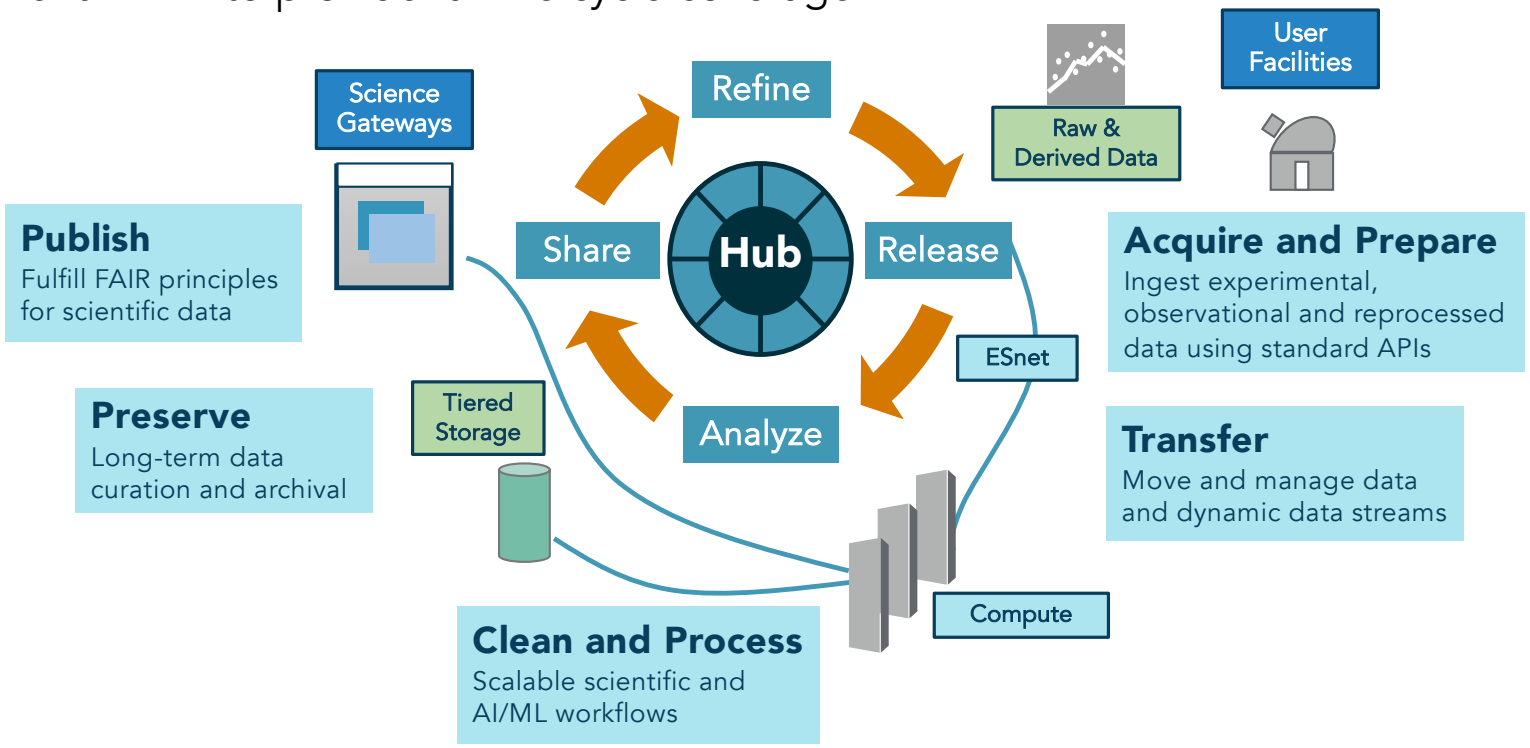

HPDF Will Address SC Priority IRI Science Patterns

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

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

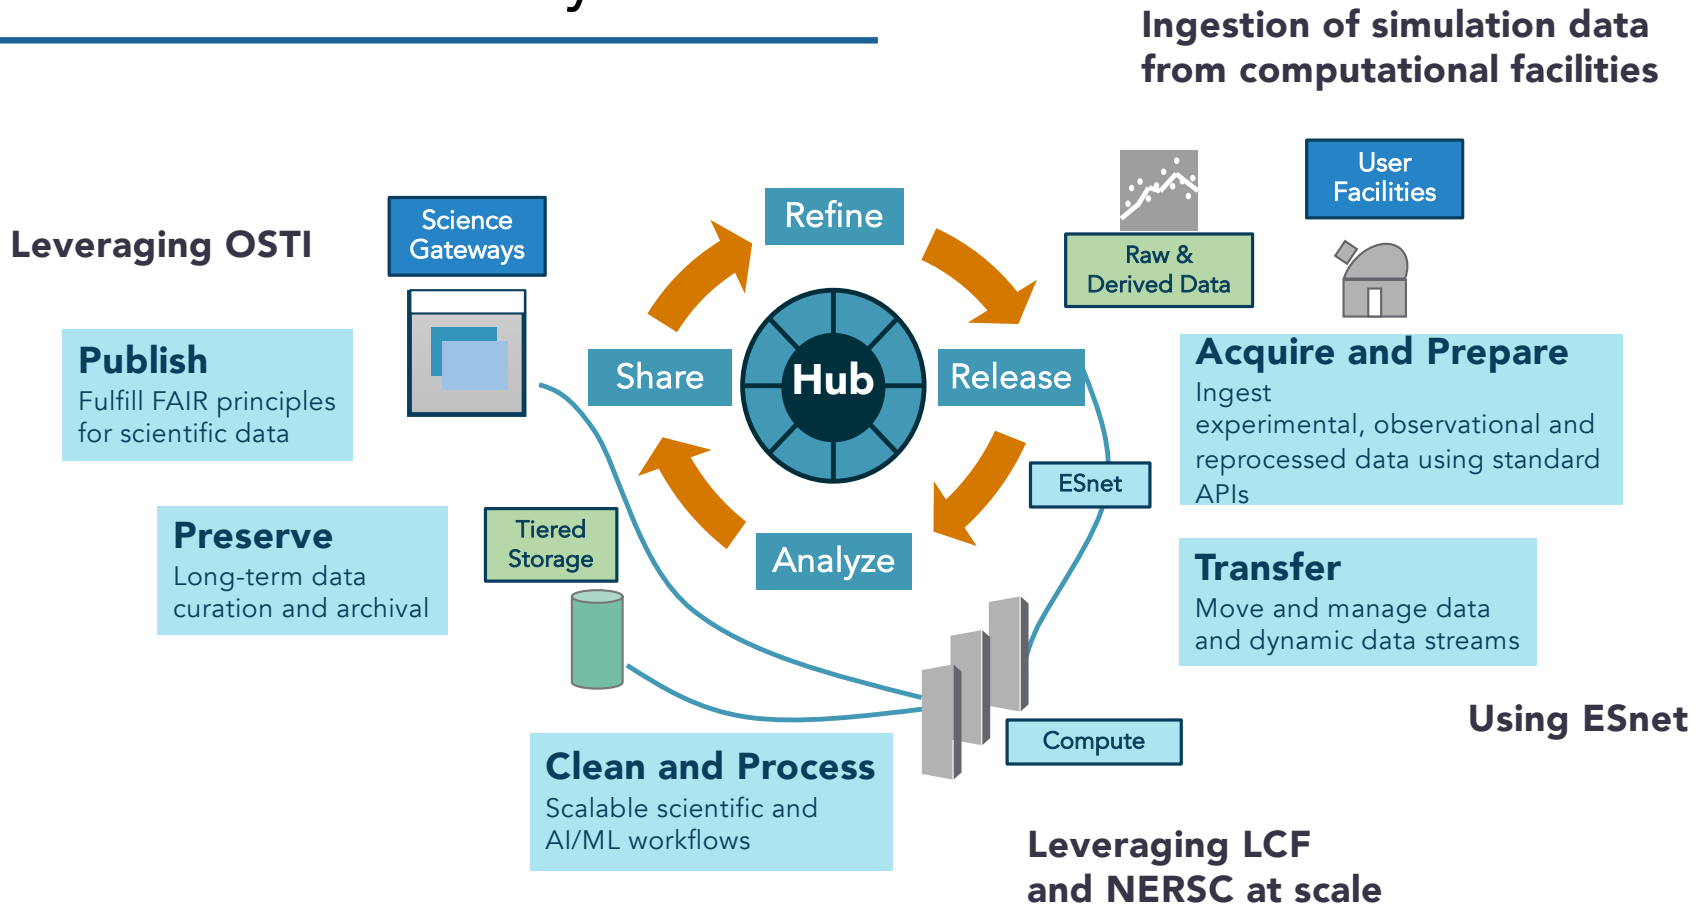

Design Overview

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

High-Level HPDF Technical Concept

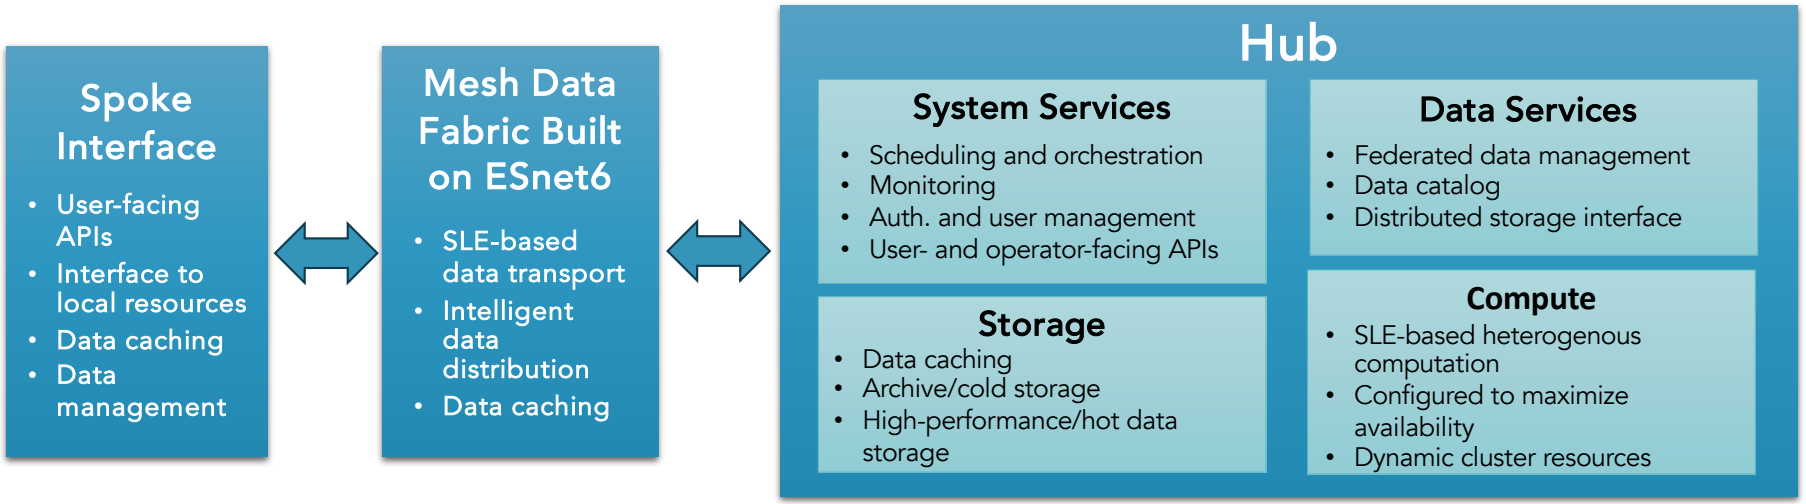

Design methodology, qualification, and approach:

- Pilot and phased delivery, enable early development, fine tune design
- Use of proven technologies to ensure a reliable, robust platform
- Hardware distributed and replicated at both sites to improve reliability and geographic diversity
- Modular heterogeneous approach to support a broad range of analysis

Approach to delivery and modularity allows composition adjustment during the design phase

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

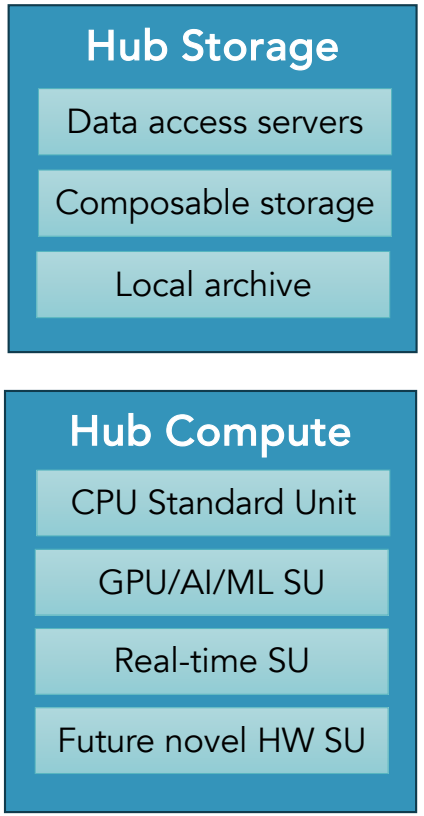

HPDF Architecture Stack

- 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

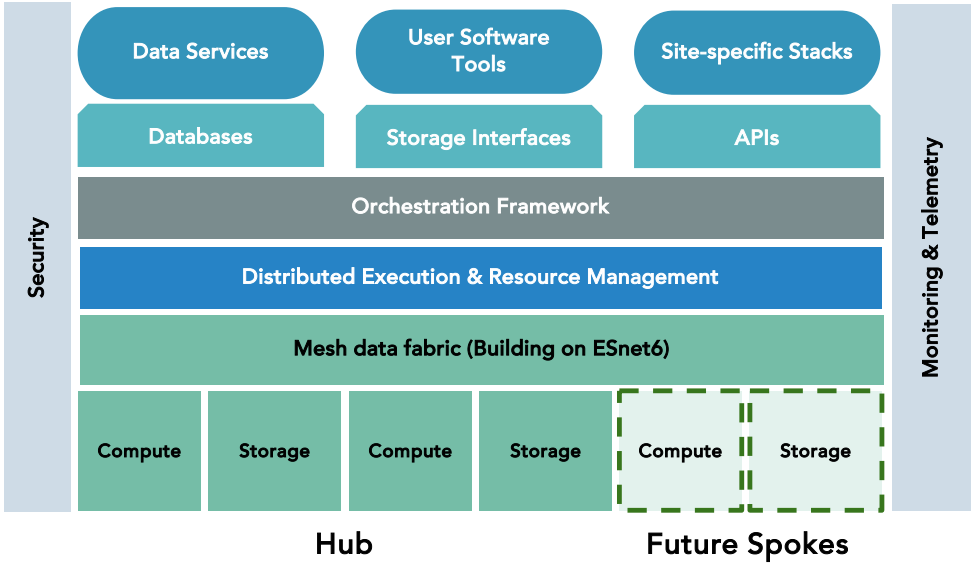

Both sites will have similar storage & compute resources, possibly at differing scales, that complement existing ASCR HPC facilities and target data-driven use cases

Preliminary Results: Scalable Data Management Infrastructure Mapped to IRI Patterns

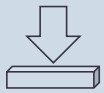

Data Capture & Storage

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

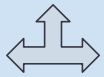

Data Management & Staging

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

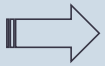

Programmable APIs

- APIs to services across entire data life cycle
- 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

Software & Data Services Strategy: Innovation & Stewardship

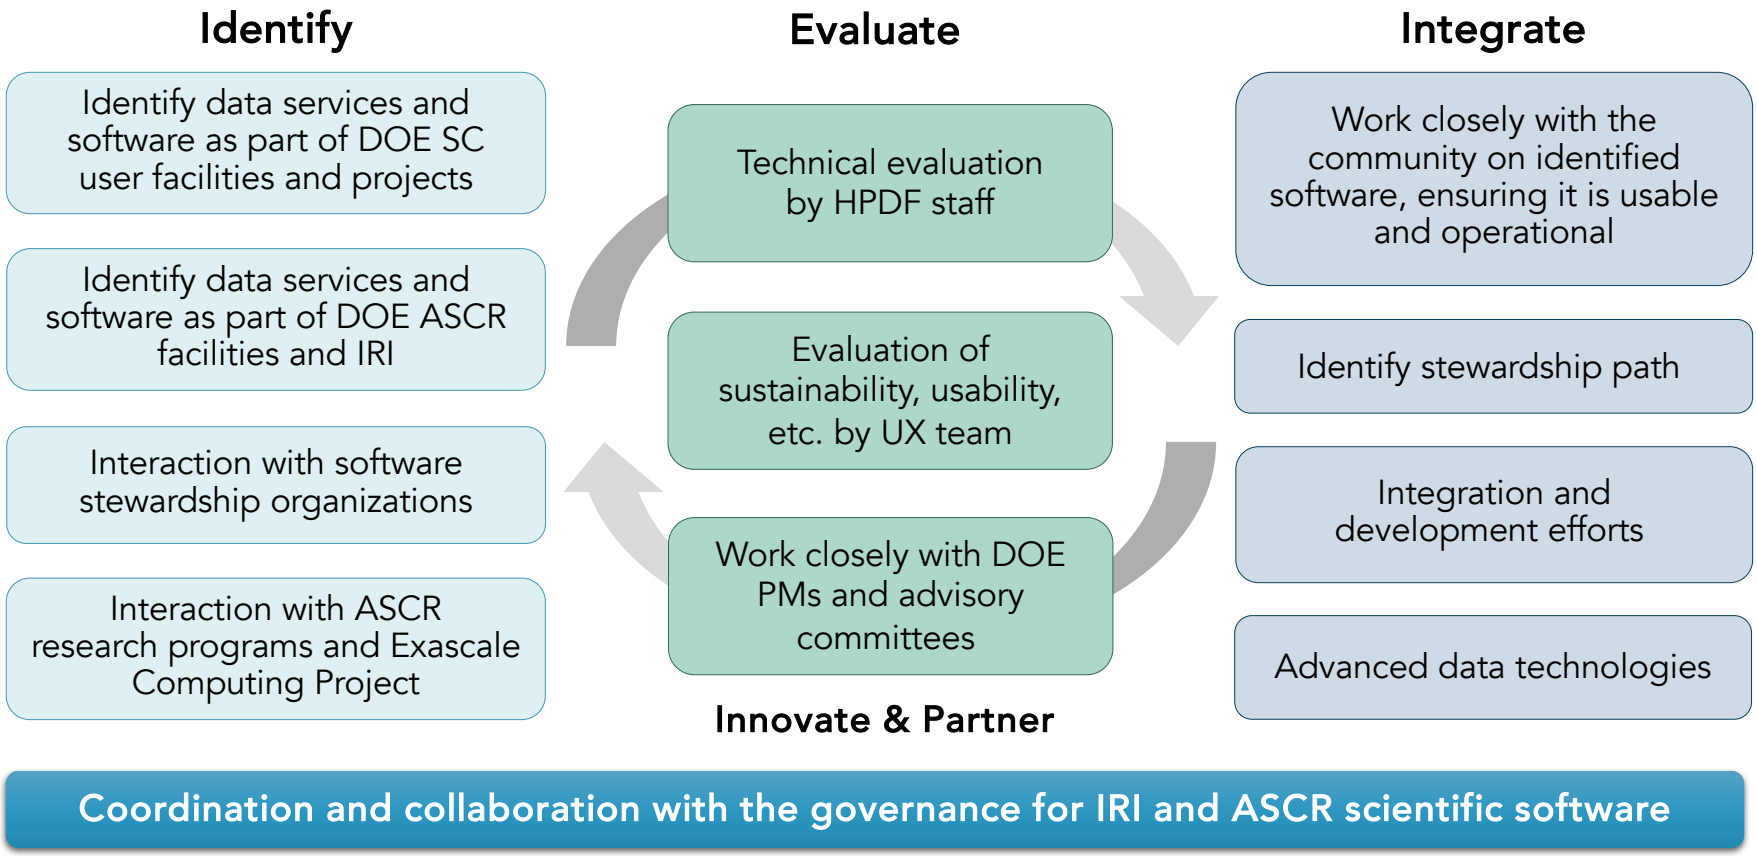

Science Engagement & Partnerships

Science Engagement & Partnerships: Critical to HPDF Success

- **ASCR Facilities & IRI:** *Meet the needs of end users*
- **DOE SC User Facilities and SC Projects:** *Understand user needs and develop strategic partnerships*
- **DOE SC Program Managers:** *Identify current and future program needs*
- **ASCR Research:** *Leverage research results and influence future requirements*
- **Open-Source Community:** *Leverage and contribute to the data ecosystem*
- **Vendors:** *Toward deploying a world-class data facility*
- **National & International Partners:** *Leverage existing efforts and establish leadership in key areas*

Community & User Engagement Strategy

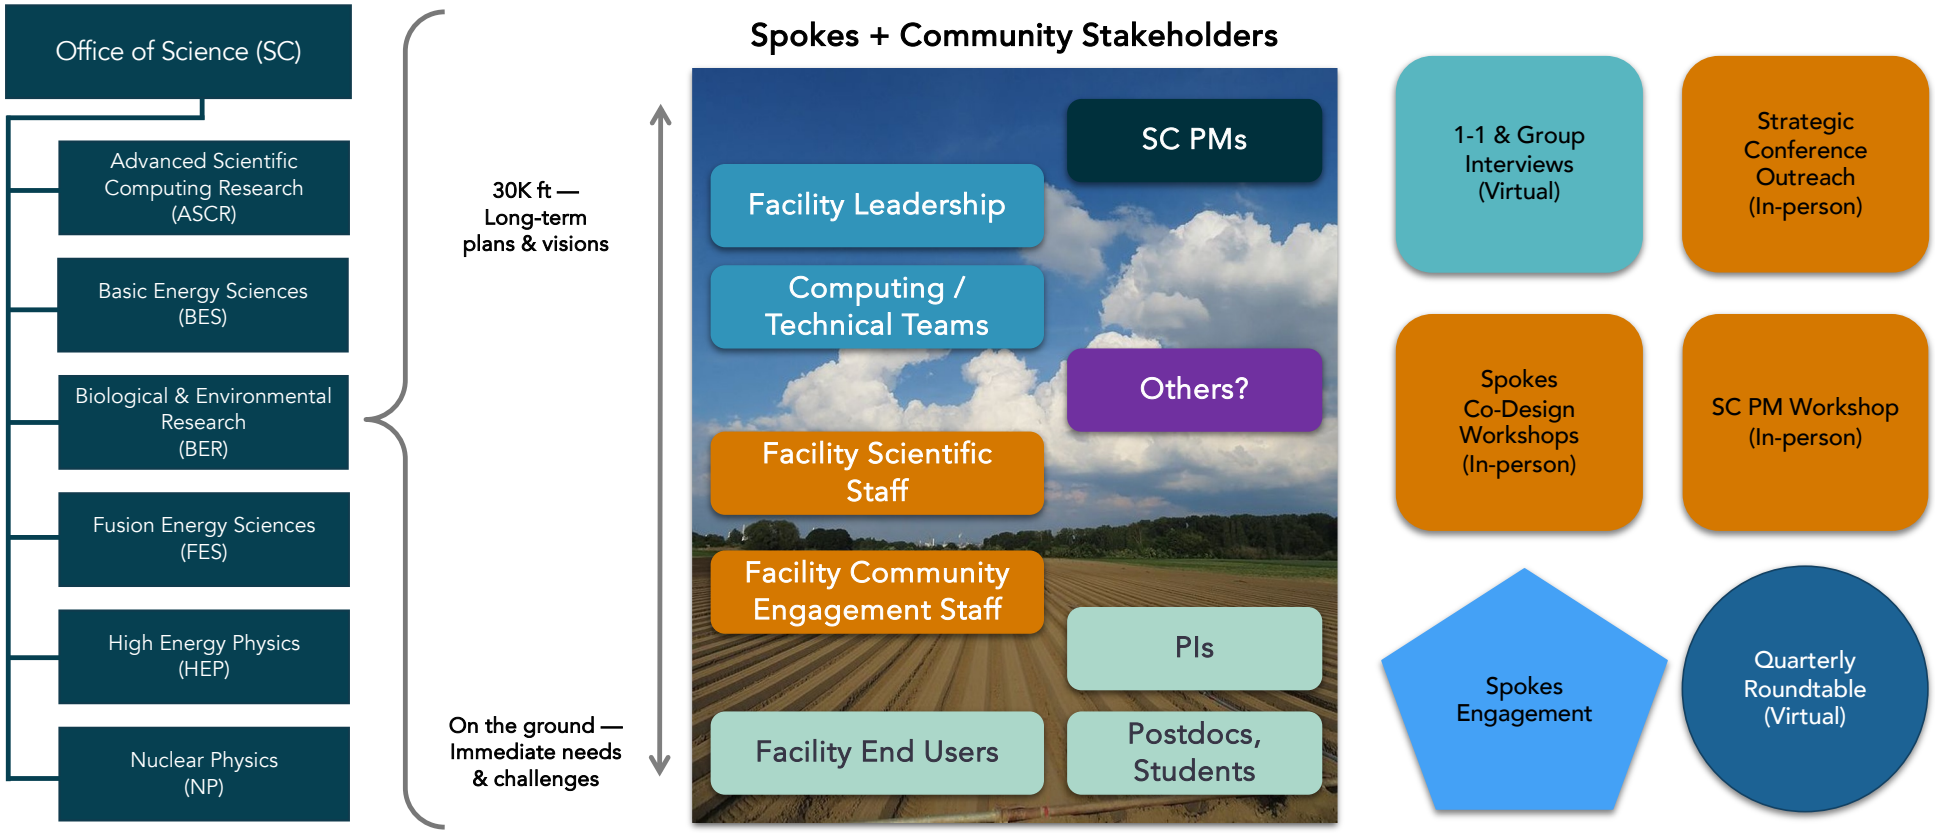

Initial Types of Engagement Activities

Amount of any type of activity will vary based on HPDF project needs & available resources

User & Community 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

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

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

"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

Spokes

Spokes Are a Critical Component to HPDF

Spoke types based on programmatic elements

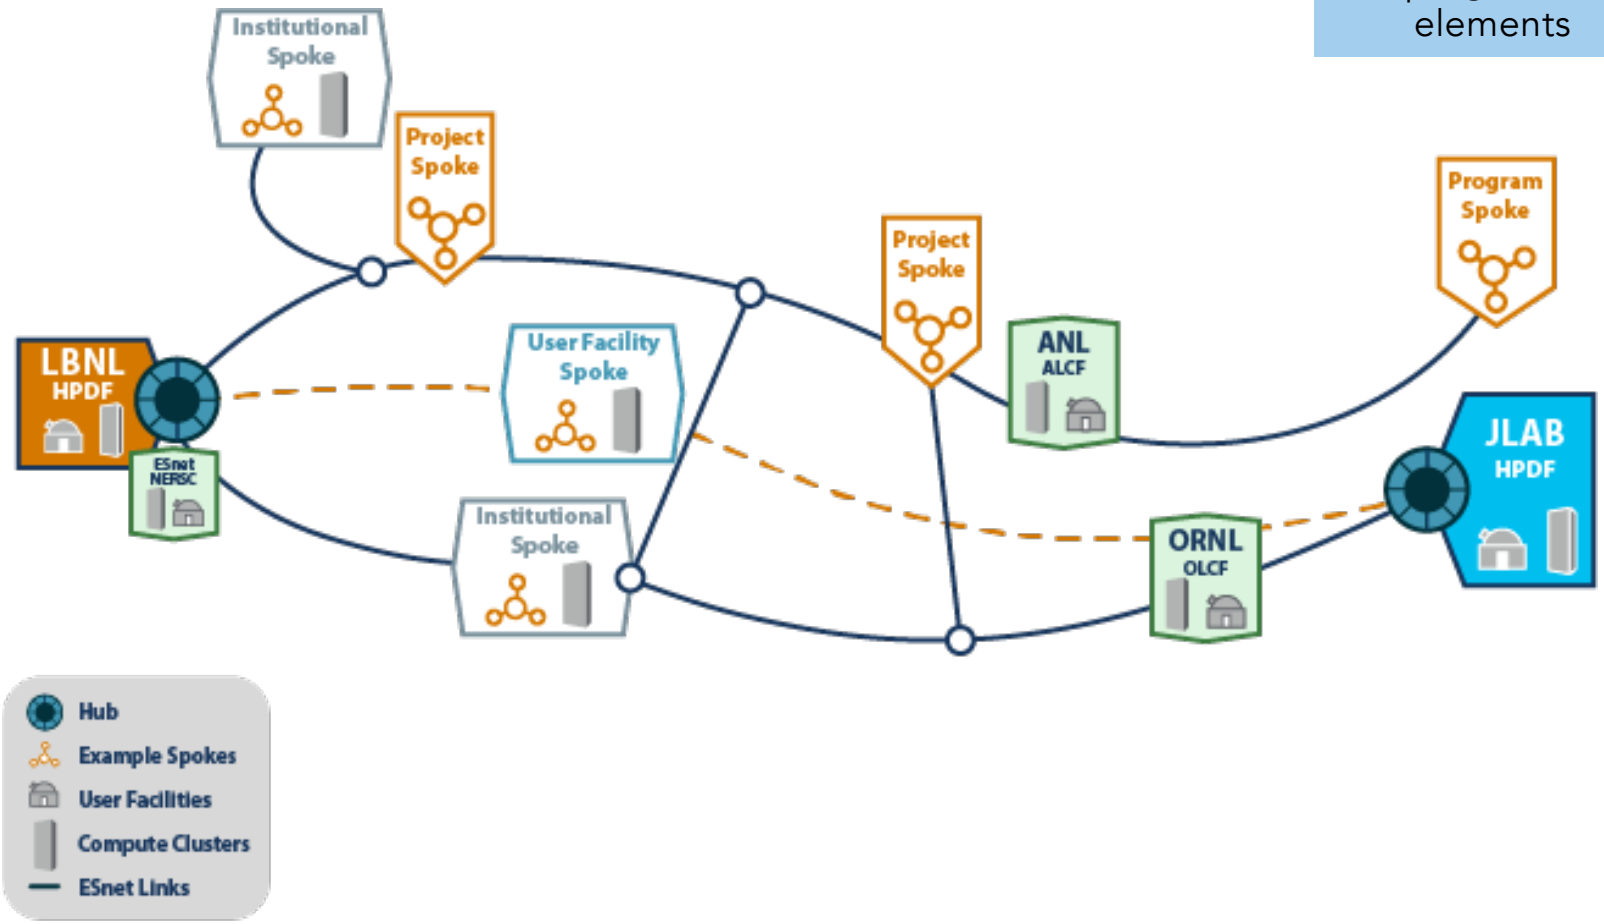

Key Facets of Spoke Design

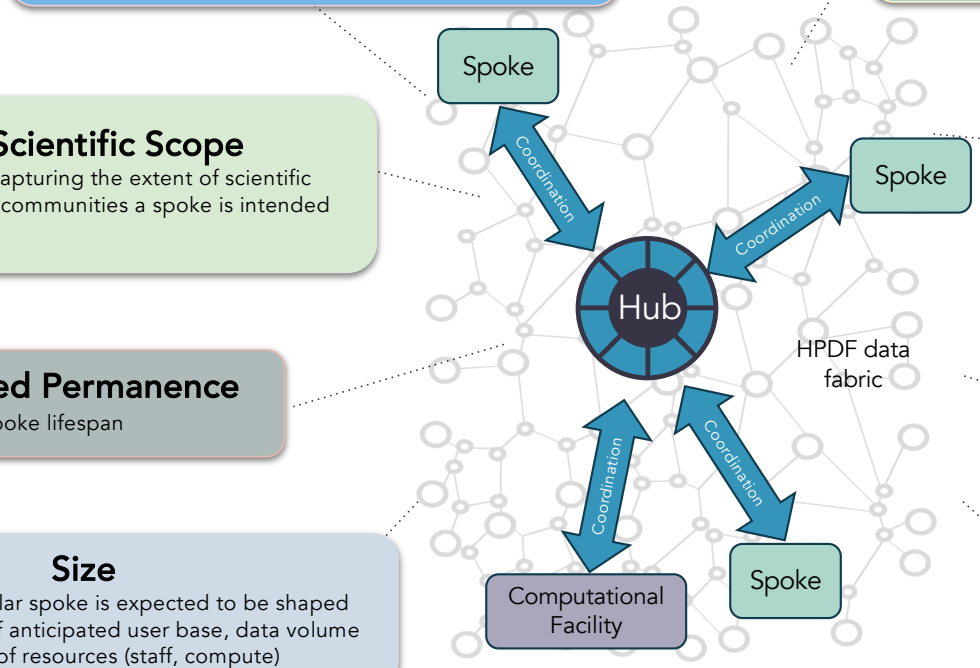

Funding Model(s)
How the spoke itself is funded, for what lifespan, and how end users are economically supported (sub-awards, allocations, etc.) to leverage its resources

Organizational Structure
How the organization or institution is designed to support its user community's full data life cycle.

Scientific Scope
Spectrum capturing the extent of scientific domains & communities a spoke is intended to support.

IRI Patterns
The type(s) of patterns the spoke is expected to have to support.

Planned Permanence
Intended spoke lifespan

Community Structure
How strongly or loosely governed or united the community is around a set of policies or goals for its data products

Size
The size of a particular spoke is expected to be shaped by the confluence of anticipated user base, data volume & velocity, quantity of resources (staff, compute)

Data/Compute Resources
The types and extent of technical functionality that the spoke supports for its user community.

Spoke Model Concepts: Tailoring to Science Needs

Spoke types based on interaction with Hub

Hub
Provides infrastructure, software services, user support

Projects and End Users
End users relying fully on Hub resources

Spoke Type E
Relies on Hub resources for preserving and publishing data

Spoke Type A
Provides edge services close to experiment or instrument and interfaces with the Hub through APIs

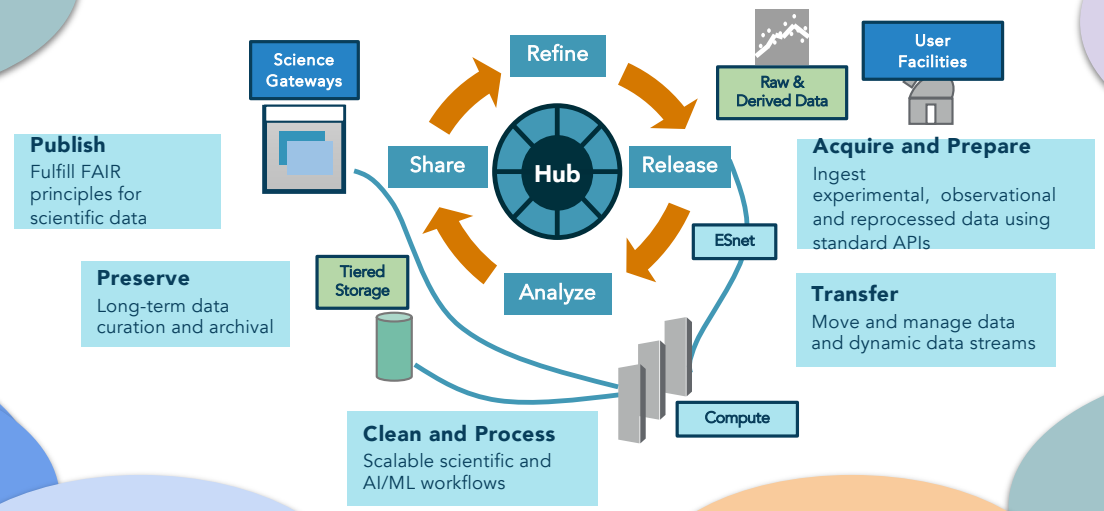

Spoke Type D
Co-designs software services w/ communi

Scientific Community
Co-designs additional software services within community

Spoke Type C
Co-designs software services with Hub

Spoke Type B
Provides computation and data services at spoke and leverages computational and/or data services at Hub

Ensuring Engagement Across Spectrum of Spokes, Projects & Communities

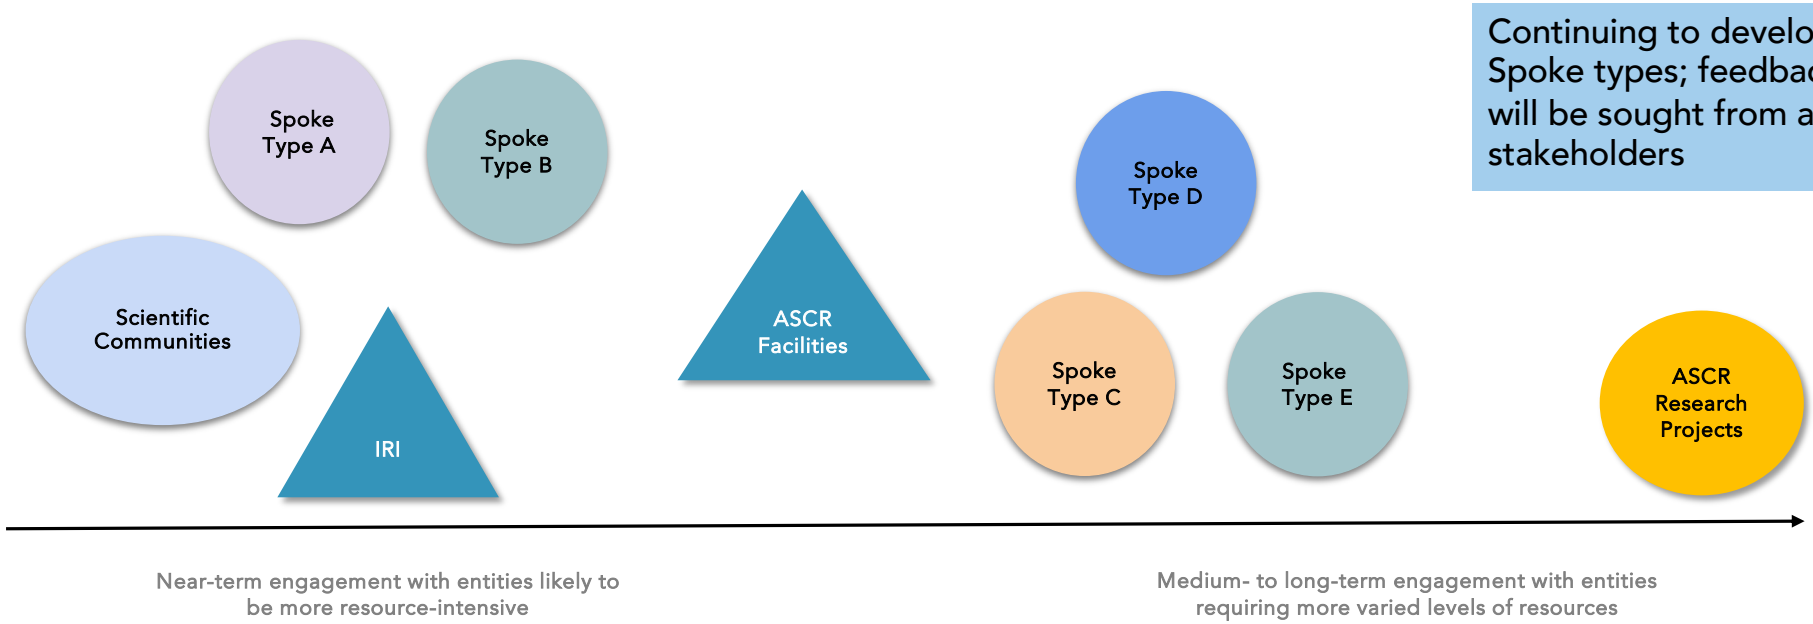

Continuing to develop Spoke types; feedback will be sought from all stakeholders

Plan to consistently engage with entities who have varied resource needs (e.g., compute, data, people) over time as HPDF offerings develop and grow

Long-Term Vision for HPDF: Extending Hub & Spokes with Partnerships

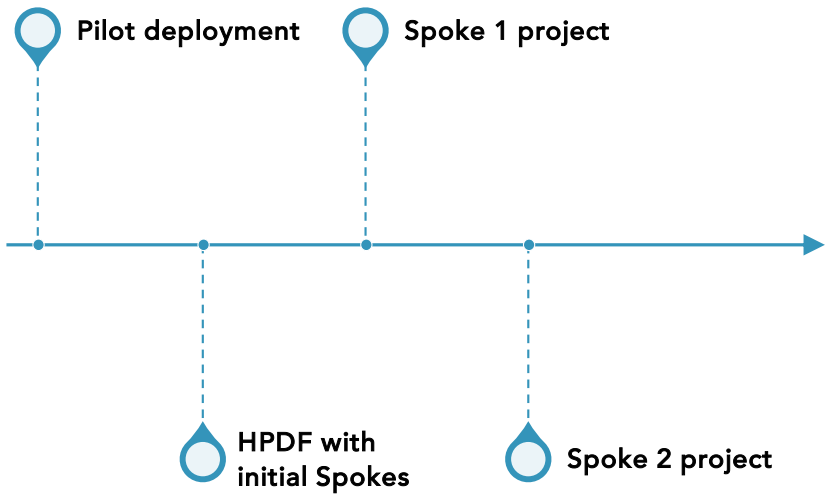

Initial Spokes

- Part of the HPDF Project – serving urgent community needs and demonstration of functionality
- Engagement with various stakeholders will establish Spoke candidates

Spokes 1 and 2 Projects

- Anticipate a large programmatic component and (possibly) joint project funding between ASCR and relevant program offices
- Early Spokes will be in tight partnership with the HPDF Project, contributing to technical innovations and eventually operational HPDF

HPDF Hub and Spokes will be intellectual partners, advancing the data life cycle of efforts across the SC complex

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)
  - FES PI meeting (June, in-person)
  - HPDF/IRI workshop (July, in-person)
  - Small-group interviews with groups identified through initial HPDF workshop (summer/fall, virtual)
  - Supercomputing '24 (November, in-person)

Please reach out if you have a DOE program or community meeting that we can participate in for outreach!

Stay Informed

https://hpdf.science

https://linkedin.com/company/doe-hpdf

https://www.youtube.com/@doe-HPDF

Share your thoughts!
Answers will be provided via a website FAQ within a few weeks.