The High Performance Data Facility: HPDF

Graham Heyes

HPDF Technical Director

March 12th, 2024

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

HPDF - Meeting the Greatest Needs

The DOE envisions a revolutionary ecosystem – the **I**ntegrated **R**esearch **I**nfrastructure – to deliver seamless and 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

IRI will be a program within the Office of Science, executive, leadership and steering committees are being set up now.

HPDF is the fifth facility funded by the Advanced Scientific Computing Research program (ASCR) as a cornerstone of the IRI to meet the needs of data driven research

HPDF will enable analysis, preservation, and accessibility of the experimental data produced by SC facilities across ALL programs

Report published in 2023, available on OSTI https://doi.org/10.2172/1984466

Three Deeper Dives: LCLS, ESGF, EIC

LCLS-HE (BES)

- Linac Coherent Light Source at SLAC, X-ray laser
- **Science:** Visualizing a broad range of excited-state molecular dynamics
- Data needs
  - Access and management
  - Caching and orchestration
  - Storage of large data sets
  - Support for a wide range of storage and processing requirements
- Experimental data is used in conjunction with Exascale simulations enabled by ECP

ESGF (BER)

- Earth System Grid Federation
- **Science:** Facilitate advancements in Earth System Science
- Data needs
  - Access and management
  - Caching and orchestration
  - Storage of large data sets
  - Data integration and interoperability
- E3SC – Exascale Earth System Model project, links to ECP

EIC (NP)

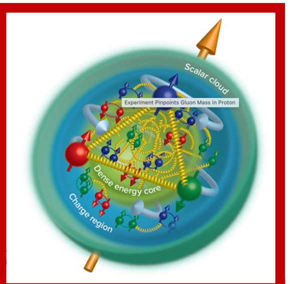

- Electron Ion Collider Facility at BNL, a joint project with JLab – distributed computing
- **Science:** Understanding the structure of protons and neutrons
- Data needs:
  - Access and management
  - Caching and orchestration
  - Realtime processing
  - Long term archive
- LQCD theory calculations enabled by ECP

ECP enabled Exascale Simulations necessary for scientific discoveries

HPDF High Performance Data Facility

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 with data flows orchestrated by the hub.

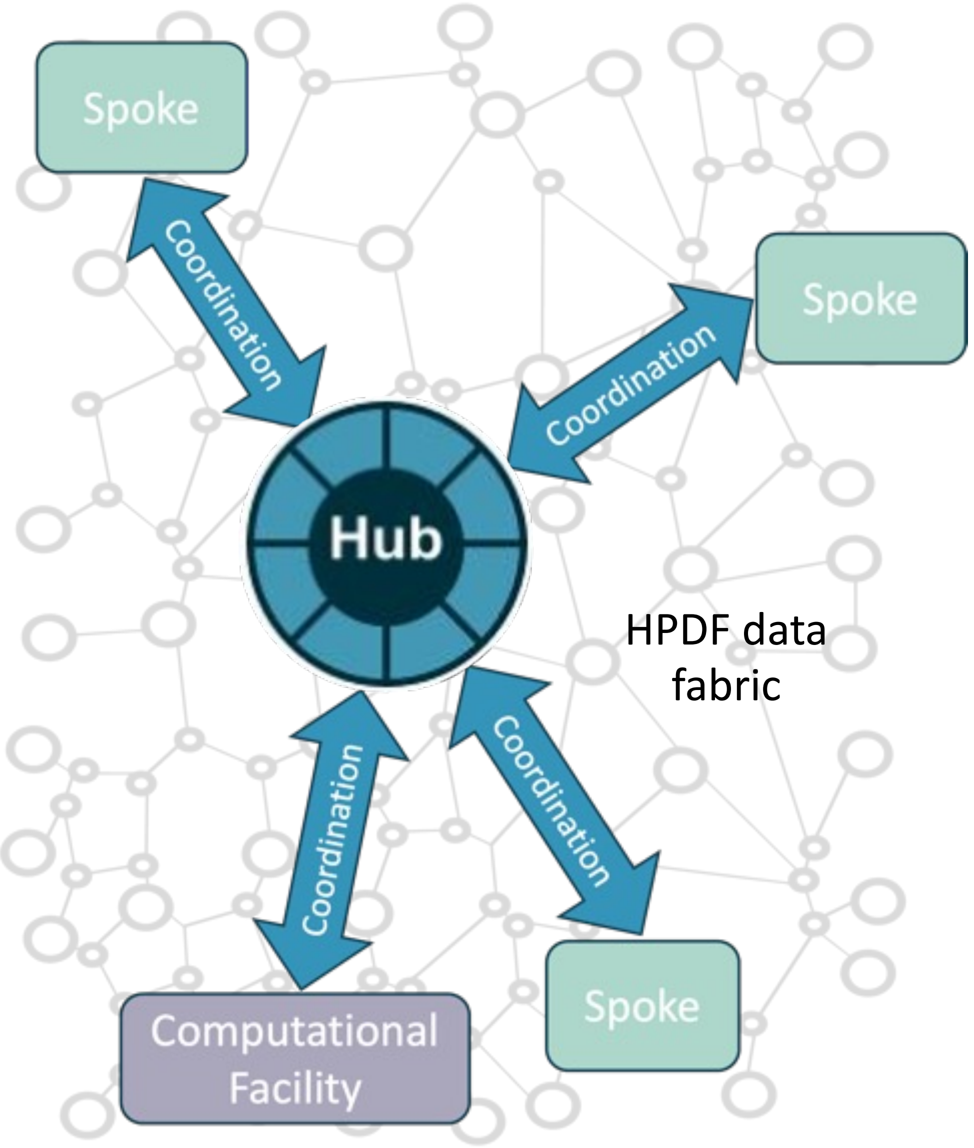

HPDF in the ASCR Ecosystem

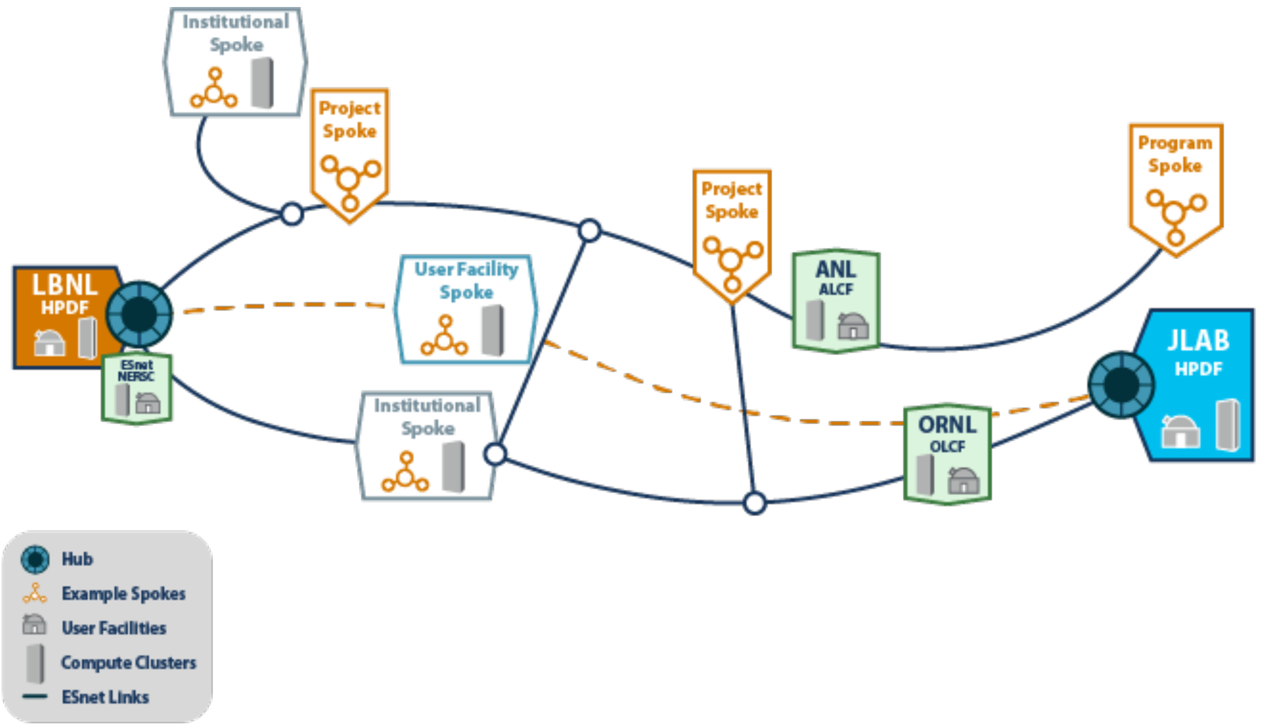

- Working with IRI, and other ASCR facilities ensures a secure, high-performance mesh data fabric that enables data and workloads to flow freely
  - Built on ESNet 6
  - Informing ESNet 7 design
- The HPDF distributed infrastructure will be designed to maximize planned availability and resilience
- Partnering 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

What is a Spoke?

- Spokes are component parts of HPDF that integrate a science "organizational unit"
  - Collaboration
  - Facility or institute
  - Instrument
  - Resource
- A spoke has two components
  - Things
    - Spoke infrastructure – how it links to the rest of HPDF
    - Contributed resources – things the organization contributes to HPDF (storage, compute)
    - Contributing resources – things HPDF contributes to the organization (services, storage, compute)
  - People
    - Support staff – HPDF funded or trained
    - Software developers – HPDF and organization
    - Science researchers and subject matter experts – people like you

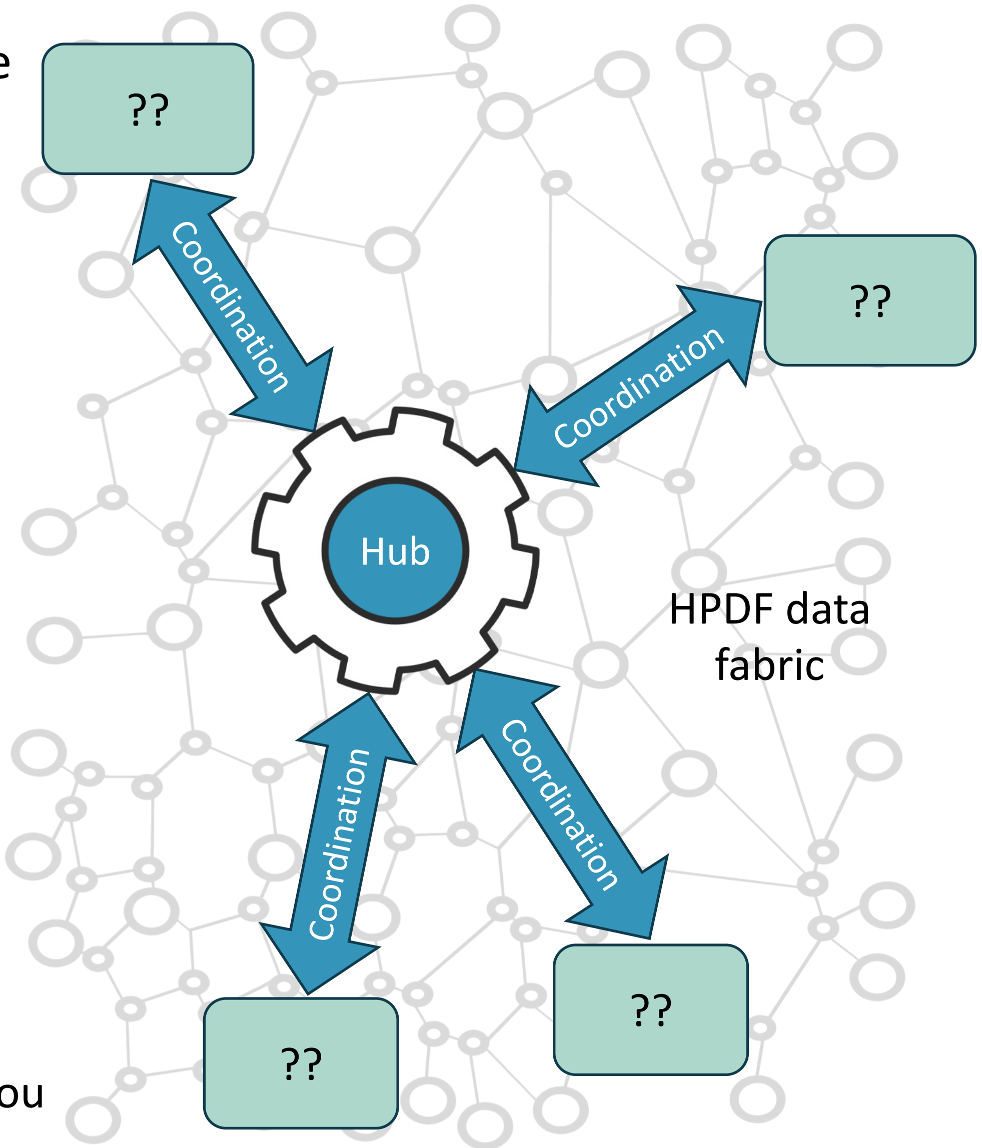

HPDF Will Address SC Priority IRI Science Patterns

Drivers	IRI Patterns	Exemplars
Supporting data curation, repositories, and archives	Data Integration-intensive, Long-Term Campaign	BER: ESS-DIVE, NMDC HEP: DESI
Supporting data processing and analysis pipelines	Time-Sensitive, Data Integration-intensive, Long-Term Campaign	BES: Light Sources, NCEM NP: EIC HEP: DUNE
Data federation, sharing, and collaboration	Data Integration-intensive, Long-Term Campaign	BES: Materials Project BER: ESGF, KBase, NSF: OSG NP: GLUEX
Real-time streaming and processing	Time-Sensitive	BES: Light Sources FES: DIII-D, PPPL/KSTAR HEP: Direct Dark Matter Searches NP: EIC, FRIB

Time-Sensitive — Data Integration-intensive — Long-Term Campaign

HPDF Will Set the 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.

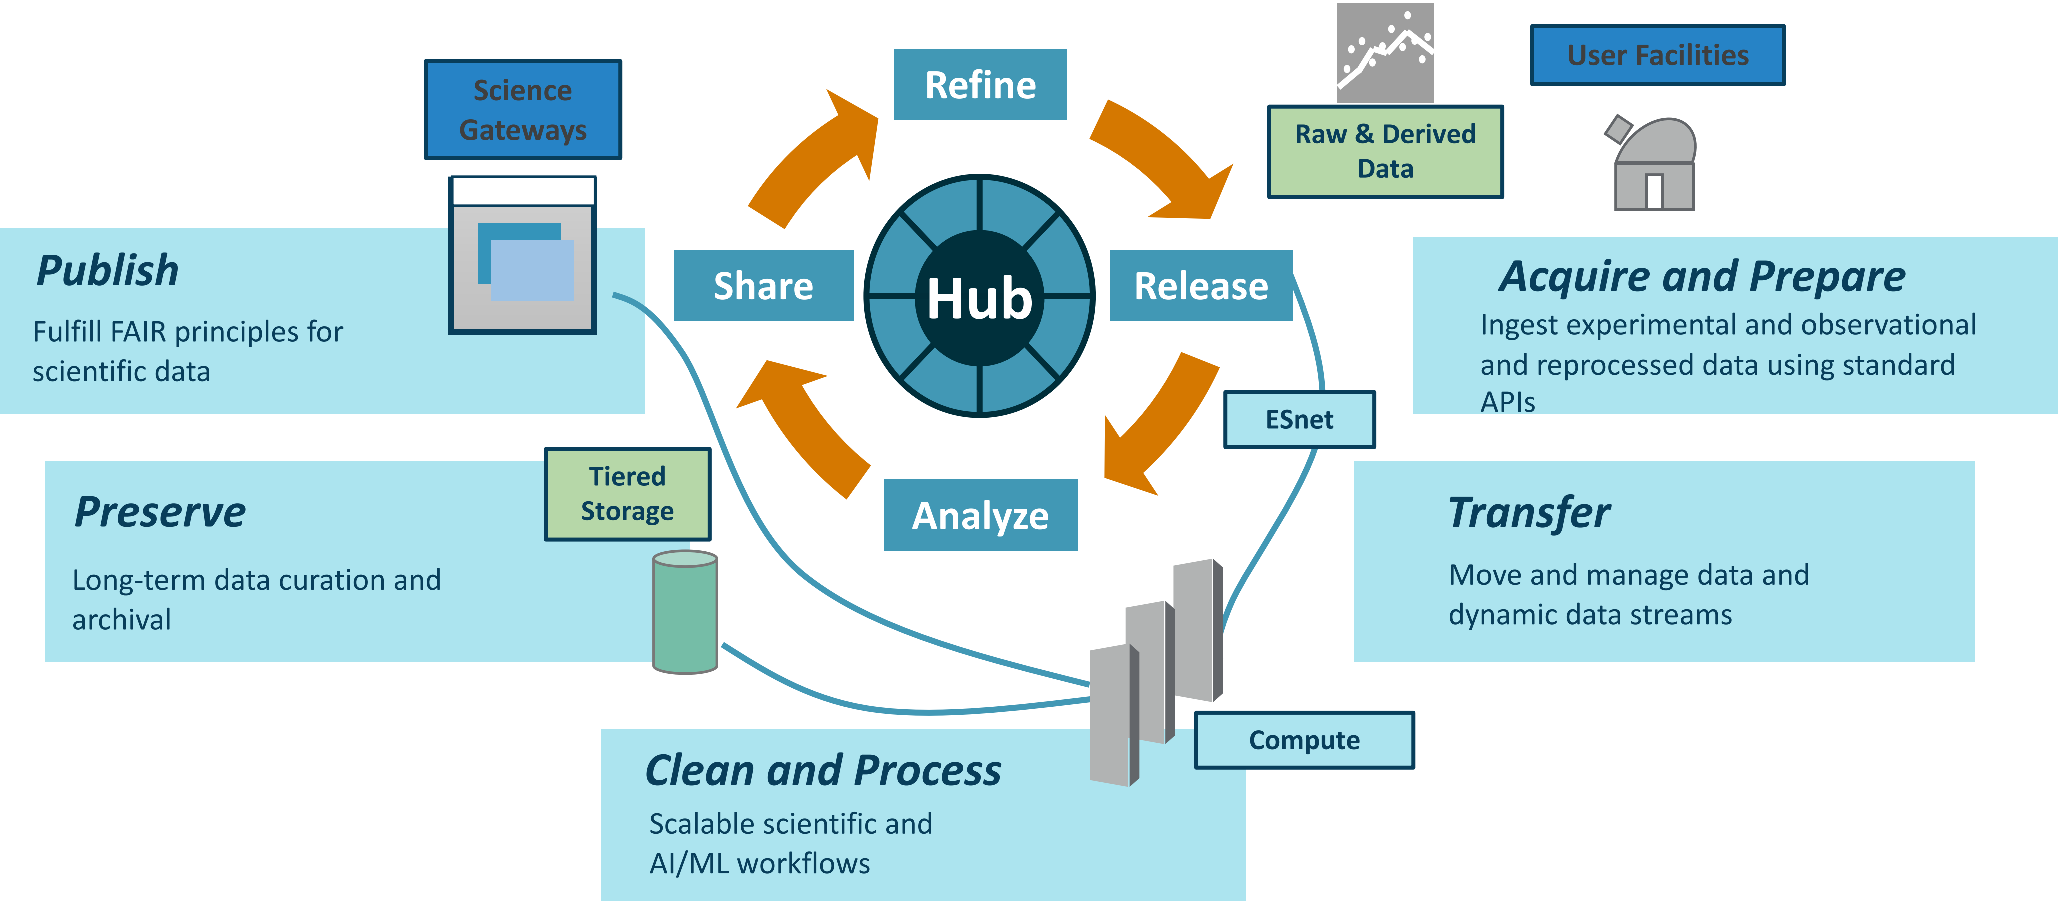

Key Facets of Data Requirements

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

Design Overview

HPDF High Performance Data Facility

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

Spoke Interface

- User-facing APIs
- Interface to local resources
- Data caching
- Data management

Mesh Data Fabric Built on ESnet6

- SLA-based data transport
- Intelligent data distribution
- Data caching

Hub

System Services

- Scheduling and orchestration
- Monitoring
- Auth. and user management
- User- and operator-facing APIs

Data Services

- Federated data management
- Data catalog
- Distributed storage interface

Storage

- Data caching
- Archive/cold storage
- High-performance/hot data storage

Compute

- SLA-based heterogenous computation
- Configured to maximize availability
- Dynamic cluster resources

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

- The hub computing infrastructure must combine high availability with flexibility while supporting time-critical and real-time workflows
- To achieve this, our design is based on the concept of “standard units,” hardware elements following well defined architectures targeting specific use cases
  - Batch jobs, AI/ML intensive workflows, streaming real-time data, dynamic reconfiguration
  - CPU/GPU flavors to run existing optimized code
- The hub will incorporate several standard units in a mix that meets the science needs but can evolve over time
- This is not a one-size-fits-all approach, but allows tailoring to need and lowers the barrier to HPDF use
- We will implement composable storage configured to minimize the need for users to modify existing code

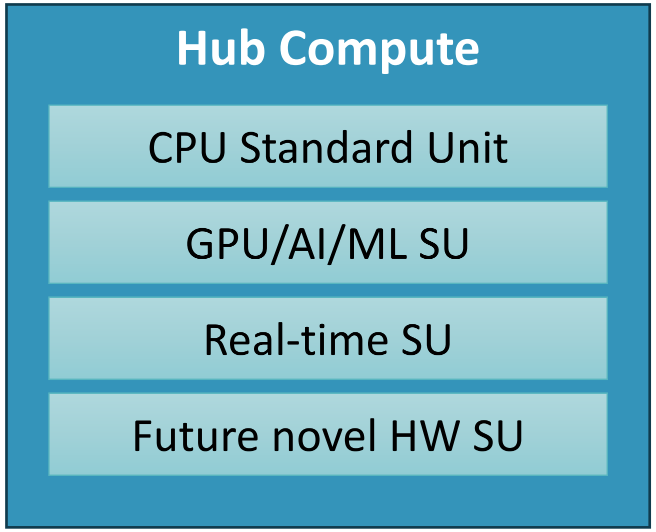

HPDF
High Performance Data Facility

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

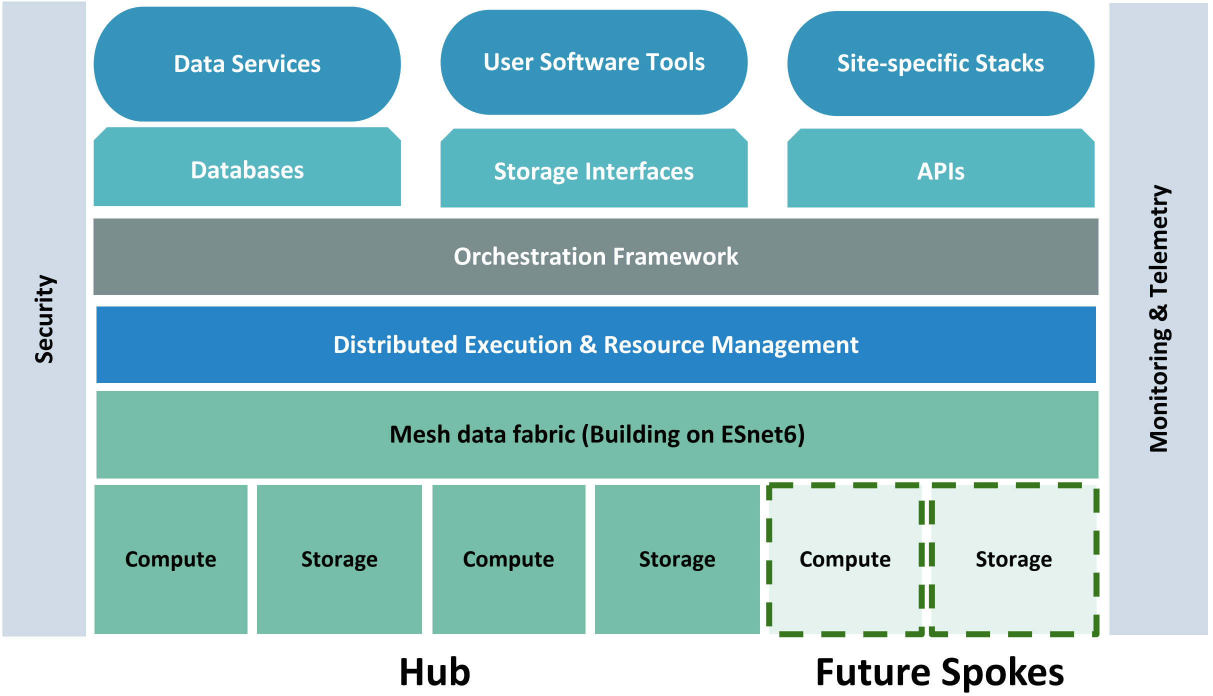

Both sites will have similar storage & compute resources, at differing scales, that complement existing ASCR HPC facilities and target data-driven use cases

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 Life Cycle Services

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

Data Management & Staging

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

Data Repository & Archiving

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

Programmable APIs

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

Data Analysis & AI

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

Time-Sensitive   Data Integration-intensive   Long-Term Campaign

HPDF Governance and Execution

Summary

HPDF Over the Next 10 Years

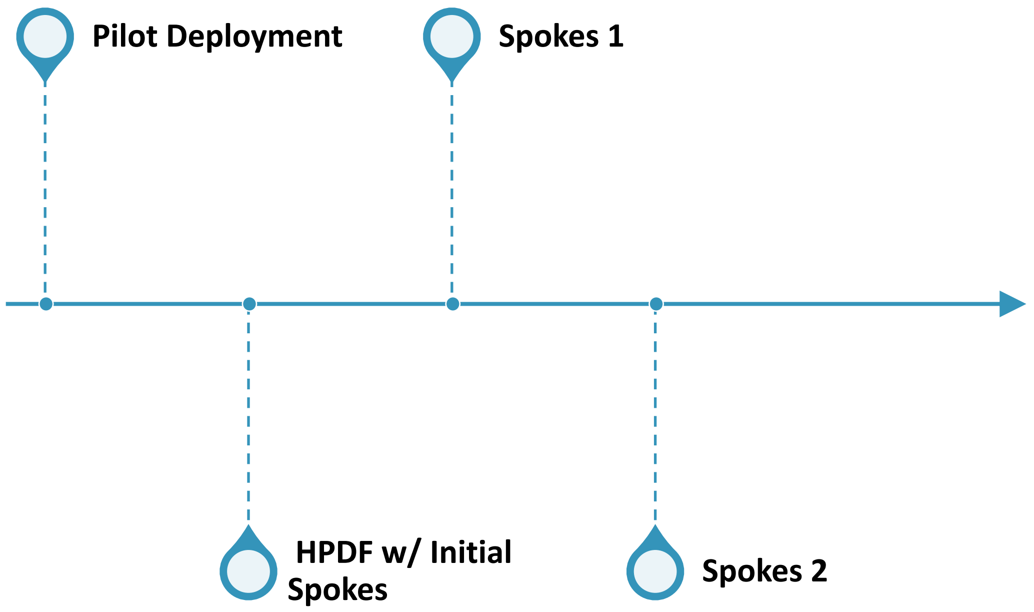

HPDF Era:

- Address: Urgency drives early spokes
- Train: Expand data workforce
- Establish: Data infrastructure component of IRI and ASCR Ecosystem
- Engage: Deep conversations with users
- Innovate: Usher in new era and new ideas

Spokes 1 and 2 Era

- Address: Extensive data needs
- Train: Data workforce at spokes
- Establish: Seamless user experience through IRI
- Engage: National and international partners
- Innovate: Storage and data technologies

Continuous Innovation Cycle is Critical to HPDF

EJFAT
ESnet-Jefferson Lab FPGA Accelerated Transport

IRI & ASCR Facilities

Research & development (ASCR, etc.)

HPDF
High Performance Data Facility

Vendors

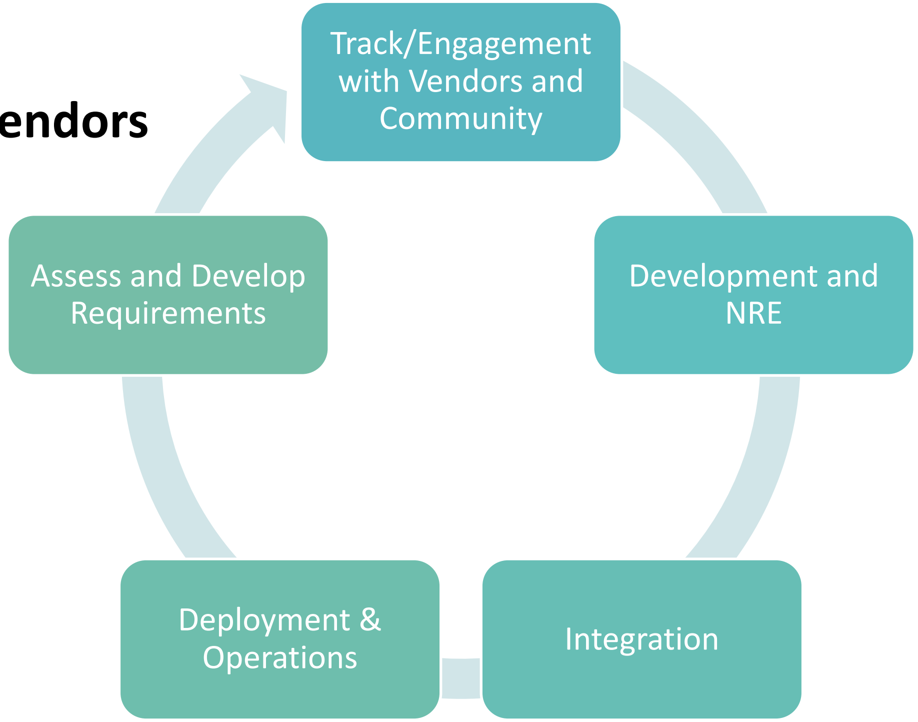

Open-source community

Data science is driving innovations in software and hardware for storage, data cataloging, and accessibility.

HPDF will drive storage technologies with science use cases and vendor partnerships

What does this mean for you?

- The addition of HPDF makes JLab a multi-purpose lab and a member of the ASCR family
- In my view nothing but positive outcomes for nuclear physics
  - Some clearly visible some harder to see
- Being the hub of HPDF requires:
  - High network bandwidth – we just upgraded to dual 100 Gbit/s more to come – priority bump
  - Power resilience – adding substations and utility feeds – possible more power options for CEBAF
  - High availability – lessons learned will be applied to all lab computing (already seeing improvements)
- The HPDF project brings:
  - Influx of people with new ideas, experiences, and skills
  - R&D projects that benefit both HPDF and NP – like EJFAT, applied AI/ML, advanced computing technologies
  - Things that are core to HPDF's mission but were secondary to the NP mission, like FAIR data
  - Heightened visibility across a broader science community
  - Positive changes in policies and procedures as the lab transitions to a multi-funding source model
  - The opportunity to be more than one-deep in many support areas

Overall, it's an exciting journey to be starting out on!

Q&A