Bridging the Gap: Scalable Data Management for Scientific Workflows with IOWarp

PI: Xian-He Sun co-PIs: Koughkas, Heber, Hochhalter, Srikumar

M. Scot Breitenfeld

Scientific I/O Challenges

Material Science Workflow

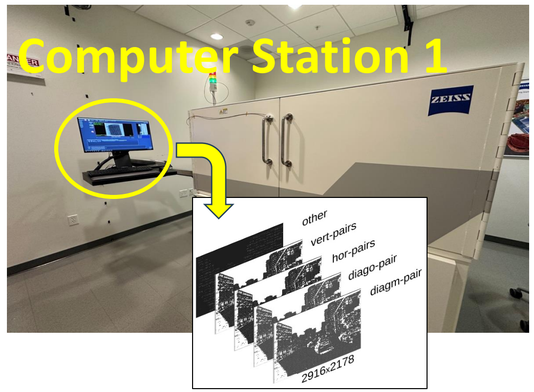

CT X-Ray Experiment

- Raw data formats depend on machine and application

Data Collection

- Computer Stations 2 & 3 are for analysis.
  - Software licensing requires different machines
- Dragonfly (station 2)
  - Image-based
  - Output file format: ORSSession (binary file)
- Simpleware Scan IP (station 3)
  - Image-based meshing for FEA/CAD analysis
  - Output file format: INP file (-> FEA)

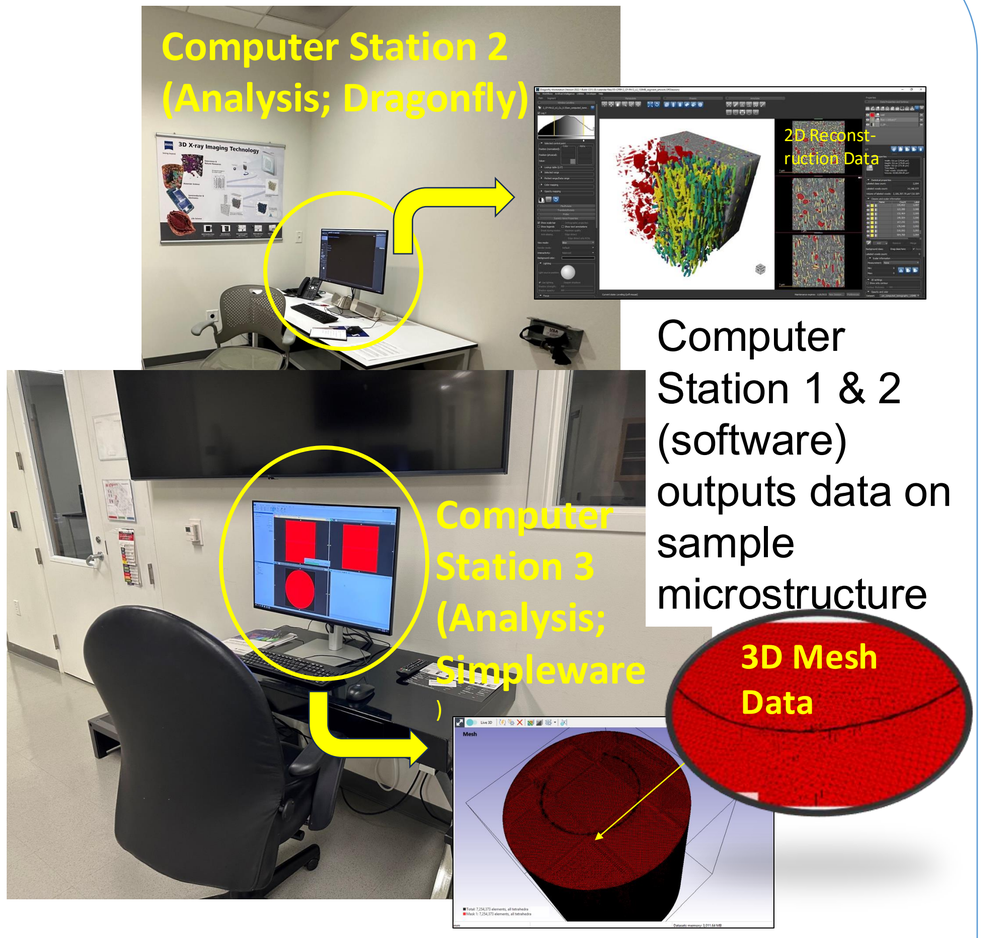

Computer Station 1 & 2 (software) outputs data on sample microstructure

Software Analysis

Data Acquisition & Management Challenges

The HDF Group

Massive Data Volume	Instruments & simulations generate terabytes/petabytes.
Storage & Infrastructure	Costly and complex management of large datasets.
Data Variety	Diverse formats (text, images, sensor data, proprietary, etc.)
Difficult Data Extraction	Lack of structure or inconsistent formatting.
Organization & Metadata	Difficulty in finding, understanding, and reusing data.

Material Science Workflow

THE UNIVERSITY OF UTAH

The HDF Group

Tomography
(set up control system & data acquisition)

Reconstruction
(converts 2D images to 3D volume)

Analysis
(view/edit 2D slices and 3D volume datasets)

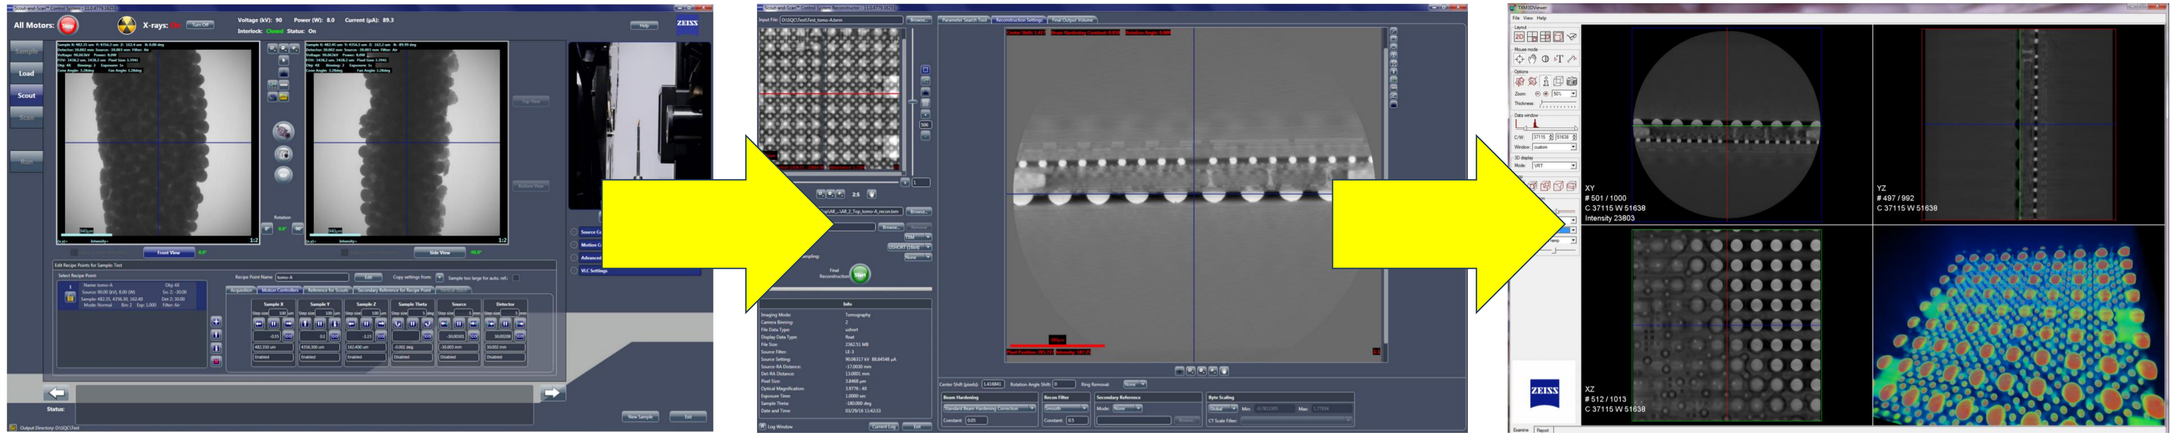

Completed with equipment at Nanofab (Raw TIFF data)

Analysis varies and is typically done on different machines.

Integration Barriers	Restrictions on connecting with external tools.
Data Quality & Cleaning	Time-consuming preprocessing for inconsistencies.
Data Veracity	Ensuring accuracy, reliability, and trustworthiness.

Material Science Workflow

THE UNIVERSITY OF UTAH

The HDF Group

Raw TIFF Data

PY File & Custom Code Outputs

```
            if not hasattr(self, '_headers_buffer'):
                self._headers_buffer = []
            self._headers_buffer.append(("%s %d %s\r\n" %
                    (self.protocol_version, code, message)).encode(
                        'latin-1', 'strict'))

    def send_header(self, keyword, value):
        """Send a MIME header to the headers buffer."""
        if self.request_version != 'HTTP/0.9':
            if not hasattr(self, '_headers_buffer'):
                self._headers_buffer = []
            self._headers_buffer.append(
                ("%s: %s\r\n" % (keyword, value)).encode('latin-1', 'strict'))

        if keyword.lower() == 'connection':
            if value.lower() == 'close':
                self.close_connection = True
            elif value.lower() == 'keep-alive':
                self.close_connection = False
```

python™

Custom Code

Organized, Relevant Result

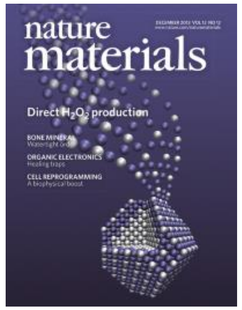

GitHub

globus

Published Results

INP data (volume mesh)

INP data (volume mesh)

Research Hub

STEP, IGES,STL,3MF (CAD Applications)

SYNOPSYS® SIMPLEWARE

CENTER FOR HIGH PERFORMANCE COMPUTING

SIMULIA ABAQUS

STEP, IGES,STL,3MF (CAD Applications)

SOLIDWORKS

Research Hub & Continued Analysis

Analysis, Interpretation & Reproducibility Challenges

The HDF Group

Data Complexity	High dimensionality, heterogeneity, and complex relationships.
Scalability of Algorithms	Need for methods to handle massive datasets.
Interpretation of Results	Translating complex analysis into scientific insights, AI/ML tools.
Spurious Correlations	Risk of finding false relationships in large datasets..
Bias & Uncertainty	Identifying and quantifying potential issues..
Developing Appropriate Methods	Tailoring analysis to unique data characteristics, AI/ML ingestion.
Versioning & Provenance	Tracking data and analysis workflows for reliability..
Data Sharing & Accessibility	Ensuring findings can be validated.

IOWarp

IOWarp

The HDF Group

Aimed At Resolving Scientific Workflow Challenges

1. **Diverse Data Repositories**
   Efficiently connect with various external repositories/platforms for seamless data integration and transfer

2. **Mapping Legacy Data to Unified Format**
   Automatic standardizes data for efficient transfers and interoperability

3. **Interoperable Content Delivery by Any I/O Interface**
   Provides adaptive interfacing and complex data manipulation

4. **Workflow-Aware, Hardware-Optimized Content Organization**
   Optimizes content placement and access using machine learning to enhance workflow performance and storage efficiency

IOWarp

The HDF Group

Aimed At Resolving Scientific Workflow Challenges

5 **Optimizing Task Scheduling in Heterogeneous Environments**

An advanced task scheduler for allocating tasks based on data access patterns and resource usage.

6 **Accelerate I/O with Advanced Storage Hardware Interfaces & System Resource Information Extraction**

- Leverage modern storage technologies to boost I/O performance
- Use data tracking and examination to improve resource use and boost performance

7 **Natural Language Analytics Interface & Streamlining AI/ML Framework Integration**

- Integrates an intuitive, natural language-driven assistant (WarpGPT) for advanced data analytics and AI-driven scientific workflows
- Enables AI/ML frameworks for seamless data access and integration into workflows

IOWarp

The HDF Group

The framework, example, tutorials, and documentation can be found at https://github.com/iowarp and https://grc.iit.edu/research/projects/iowarp/

Four Key Components

1. Platform Plugins Interface (PPI)
2. Content Assimilation Engine (CAE)
3. Content Transfer Engine (CTE)
4. Content Exploration Interface (CEI)

Platform Plugins Interface (PPI)

The HDF Group

What is it?

The interface that boosts interactivity and adaptability by linking external tools and services

Features

External Service Integration: Provides plugins to connect with:

- **Integration with Resource Schedulers** (e.g., Slurm): For awareness of job allocations and resource mapping.
- **Workflow Managers** (e.g., Pegasus): For task mapping and provenance tracking.
- **Monitoring Systems** (e.g., Ganglia, Prometheus): Gather system telemetry for intelligent decision-making.
- **Custom Library Support**: Allows users to integrate their own libraries for tasks such as custom data transformations, analysis, tracing, encryption, AI.
- **Enabling Complex Orchestration**: Facilitates coordination between data management (IOWarp) and overall job/workflow execution.

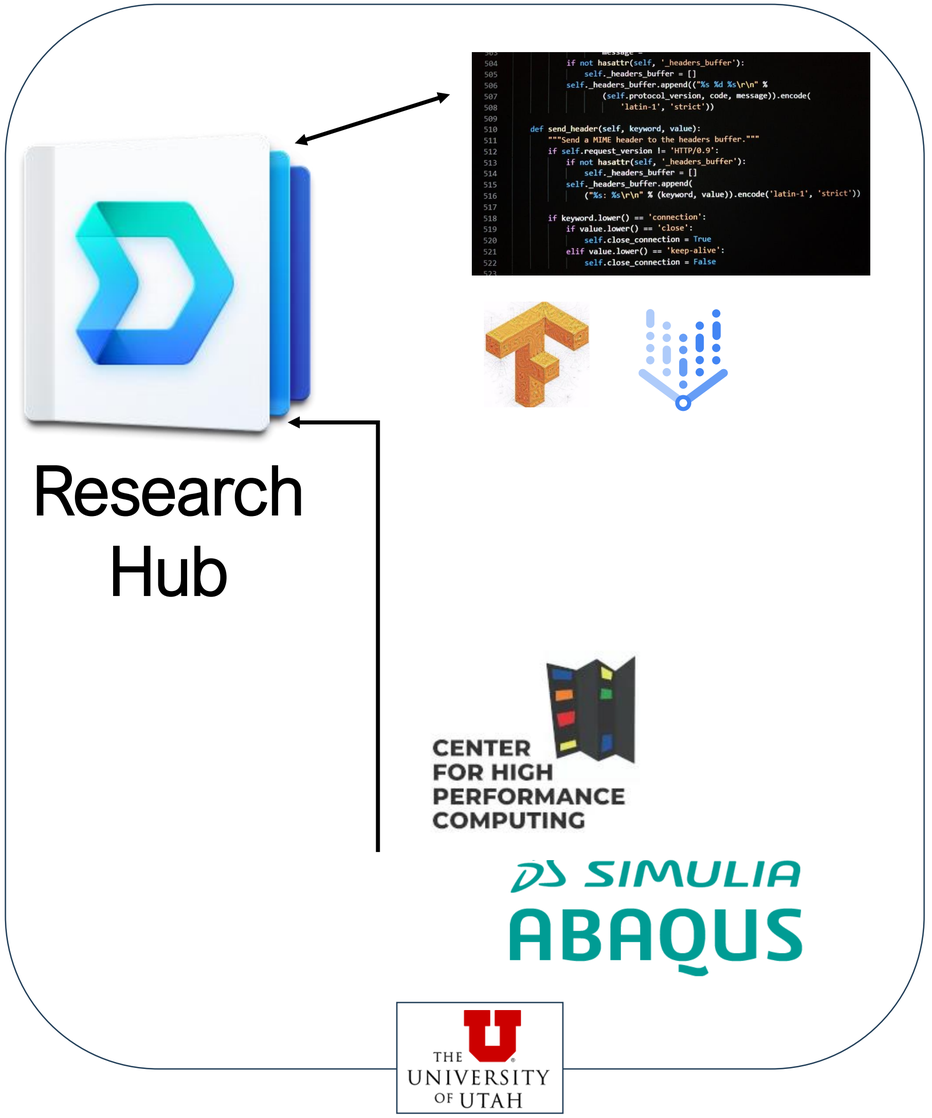

Content Assimilation Engine (CAE)

The HDF Group

What is it?

Distributed systems ingest data from various sources with different formats. Normalizing this data simplifies processing.

Features

- **Ingestion:** Connects to diverse data repositories (e.g., file systems like PFS, object stores like S3, HSDS, streaming sources, databases).
- **Transformation:** Converts data from original formats (e.g., HDF5, NetCDF, Parquet, CSV) into a unified internal representation ("Content").
- **Semantic Tagging:** Attaches metadata (e.g., data type, structure, origin, units, transformations applied) to the unified representation.
- **Benefit:** Simplifies diverse data formats for a consistent view in downstream components.

TXM (ZEISS proprietary microscopy data format; machine default)
TIFF (Universal for material scientists)
DICOM (Standard medical format in radiology)
BIN (Binary file format)
OBD, INP (Abaqus)
3MF,IGES, STEP, STL (Solid Works)

THE UNIVERSITY OF UTAH

Content Transfer Engine (CTE)

The HDF Group

What is it?

Manage efficient data movement and placement, critical for performance in systems with distributed memory and hierarchical storage..

Features

- **Manages Data Flow**: Orchestrates movement between workflow stages and across storage tiers (RAM, NVMe, SSD, PFS, etc).
- **Multi-Tiered I/O**: Explicitly uses the storage hierarchy, placing frequently accessed or latency-sensitive data on faster tiers.
- **Optimized Placement**: Employs strategies (building on Hermes' DPE) to decide where data should reside, potentially using access patterns or ML predictions.
- **Leveraging Fast Interconnects**: Designed to use high-performance networking (like RDMA) and direct hardware access (like GPUDirect, CXL interfaces) for low-latency, high-bandwidth transfers.

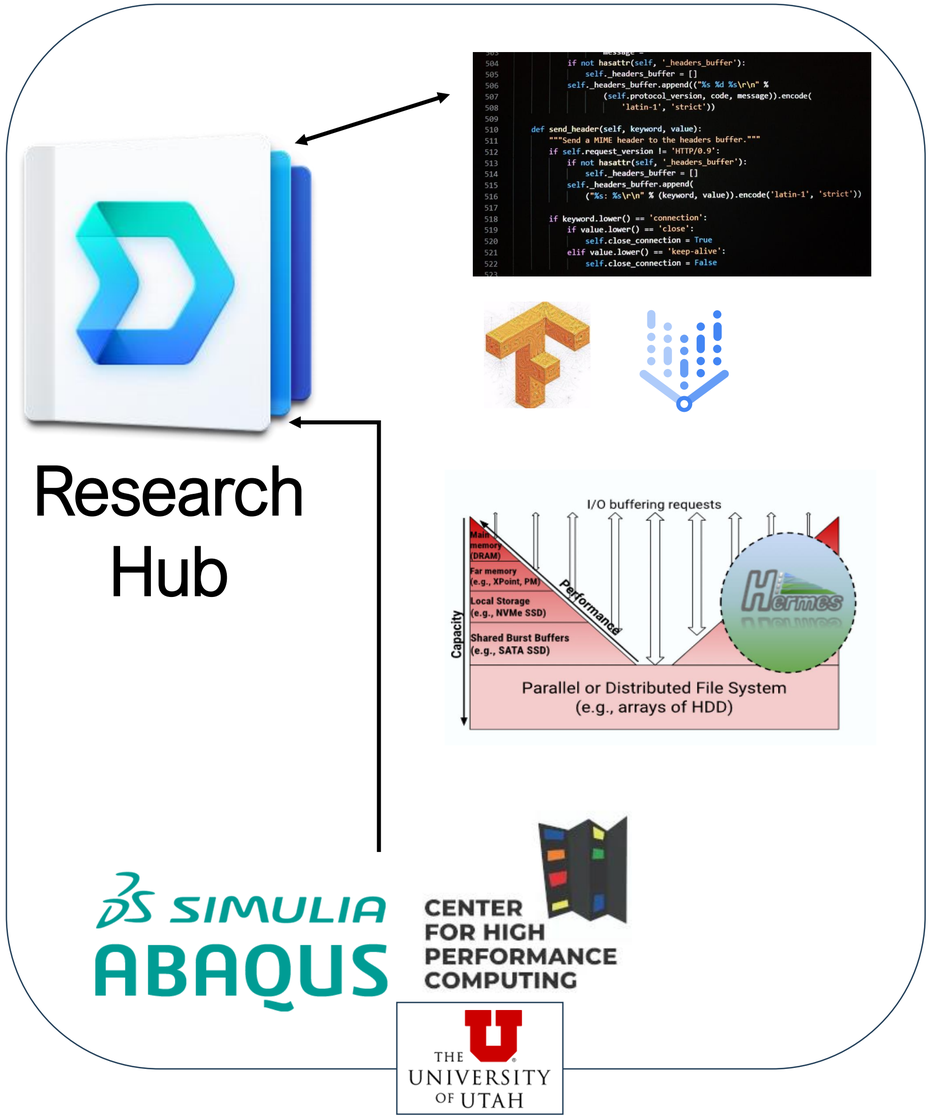

Content Exploration Interface (CEI)

The HDF Group

What is it?

Comprehensive **querying and indexing** component to facilitate advanced data analysis.

Features

- Enables human- and AI-driven queries to extract insights from scientific data.
- Advanced querying capabilities
- Builds and maintains metadata for fast, efficient data retrieval
- WarpGPT, a language model tailored for scientific data

```
➤User: Show me the average grain boundary length from all samples for the last week.

☞Response: Here is the average grain boundary length from all samples for the last week:

Date       | Average Length (nm)
-----------|----------------------------------------
2025-04-22 | 5.2
2025-04-23 | 6.8
2025-04-24 | 7.5
2025-04-25 | 6.1
2025-04-26 | 4.9
2025-04-27 | 5.5
2025-04-28 | 7.0

➤User: Are there any correlations between grain length and thermal conductivity greater than 100 W/mK?

☞Response: After analyzing last month's data, no correlation appears evident.
```

New Frontier with IOWarp

- **Beyond Basic Storage**: It's not enough to store and retrieve data effectively.
- **Focus on Insight**: Provide a workflow tool for analyzing data to discover meaningful patterns and knowledge actively.
  - **Smart Analysis**: Employ intelligent methods, such as AI/ML and complex simulations, to gain a deeper understanding of the data.
  - **Full HPC Utilization**: Leverage all the aspects of High-Performance Computing (parallel processing, CPU/GPUs, fast interconnects, memory tiers) for complex analysis.
- **Outcome-Driven**: Focus on making discoveries faster, improving predictions, and making informed decisions.
  - Use data-driven insights while keeping data flow simple and effortless

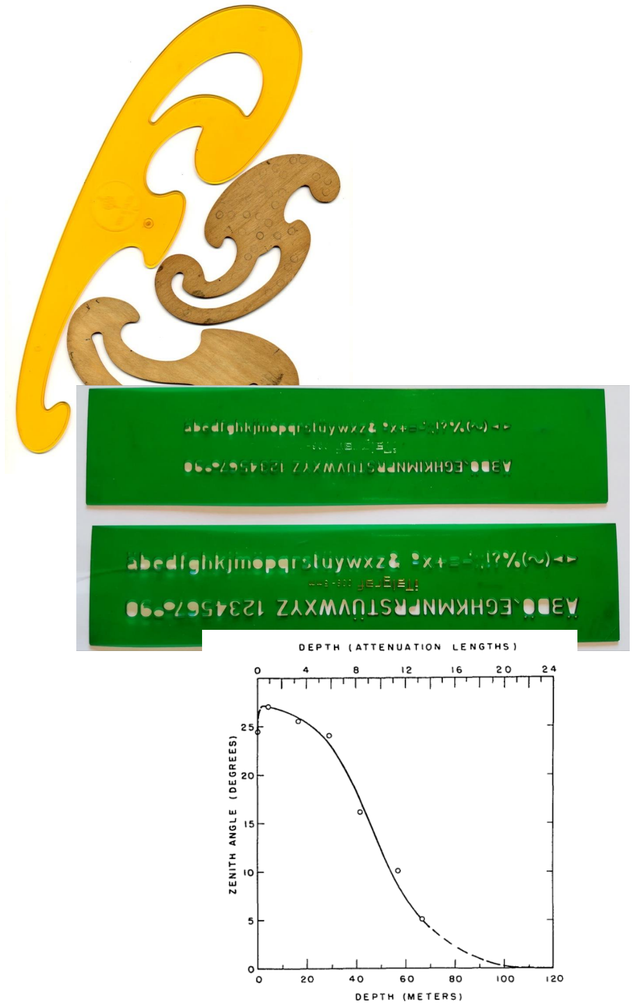

THANK YOU!

Questions & Comments?

We gratefully acknowledge the support of the National Science Foundation (NSF), Office of Advanced Cyberinfrastructure, under Grant CSSI-2411318. This funding enables us to develop IOWarp, a next-generation data management platform that Accelerates scientific discovery. We sincerely thank the research community, our partners, and stakeholders whose collaboration is essential to IOWarp's success in democratizing data access in the age of intelligent computing.