



Contribution ID: 59

Type: **Presentation**

Study on Handling Dark Data in HPCI Shared Storage System using the WHEEL Workflow Tool

Monday, October 13, 2025 2:52 PM (11 minutes)

Overview

Understanding unnecessary data, known as Dark Data is a major operational challenge in large-scale shared storage. We propose an alternative approach that leverages HPC workflow tools to collect extended metadata at each stage of job execution, minimizing the need for fundamental system changes.

Background

The High Performance Computing Infrastructure (HPCI) project was established to provide a unified environment for accessing supercomputers deployed at universities and research institutes on Japan. This project provides network-storage service to HPCI users for the shared research data between supercomputing center among HPCI, called HPCI Shared Storage(HPSS), jointly operated by The University of Tokyo and RIKEN Center for Computational Science. HPSS began service in 2012 and will transition to 3rd generation system by 2025, to serve a total logical capacity of 95 petabytes(PB).HPSS is built upon the network-based file system known as Gfarm [1].

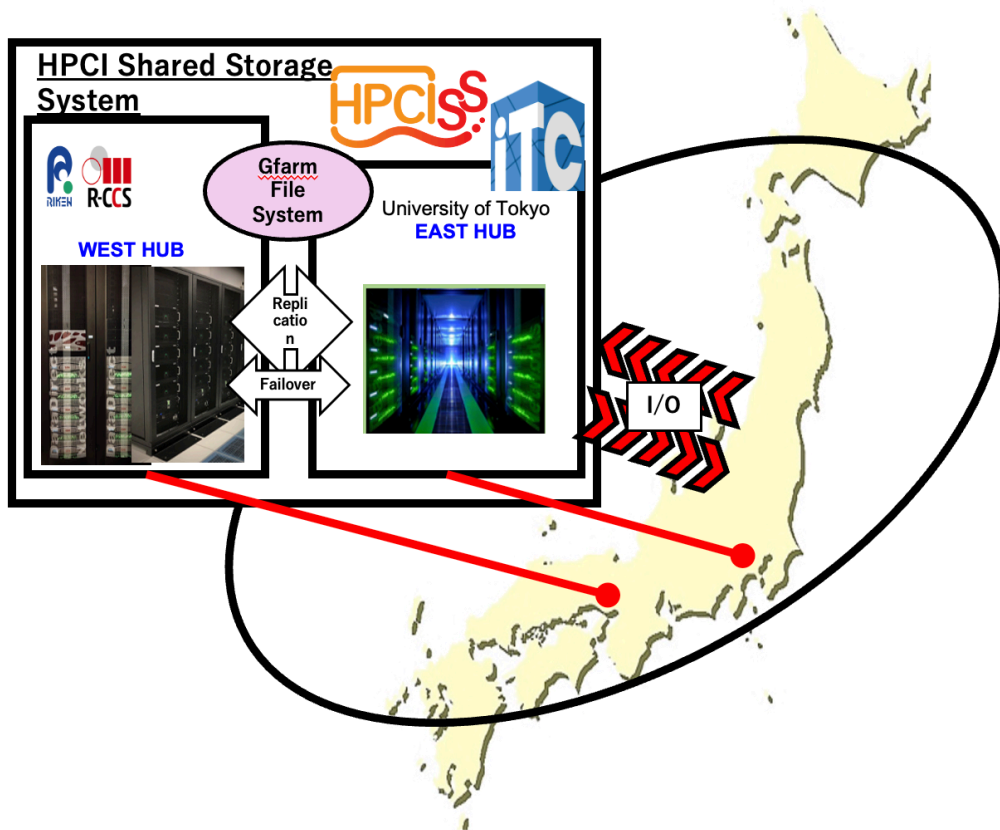


Figure 1: HPCI Shared Storage Overview

By 2024, user demand for 2nd generation HPSS surpassed the 50 PB limit, to operate under an overcommitted state. An analysis of metadata via Gfarm, revealed that much of the preserved data was Cold Data, rarely accessed despite occupying large capacity.

Discussion

Since HPSS is a free service provided to researchers with HPCI projects, maximizing its public utility is essential. Accordingly, two policies have been established

1. Prioritize users who do important research and share their data openly
2. Help users improve data management by showing infrequently used data.

To support the second policy, metadata such as access times are collected automatically and shown in dashboards like Grafana to highlight rarely used data. However, Gfarm and many HPC filesystem, such as Lustre and BeeGFS, do not provide functions to annotation custom tags or store detailed data meta information (e.g. when research data was generated). Without such extended metadata, users are unsure if certain old data and left by inactive project members, can be safely removed. This uncertainty fuels the accumulation of Dark Data, defined as data whose purpose or ownership is unclear [2]. Because users receive storage at free, there is no incentive to remove Dark Data, further contributing to Dark Data growth. A more advance data management framework is therefore critical. Ideally, expanded metadata should be automatically recorded at the time of data generation or stored storage, and should accompany the data lifecycle, from creation during computation to HPSS. However, many HPC filesystem and job schedulers (e.g. PBS) lack automated extended metadata tagging. Introducing new services (e.g. Globus, Starfish) or adding such functionality directly to Gfarm must demand extensive redevelopment and collaboration among supercomputing centers.

Proposed Method

To solve these challenges, we propose an alternative approach that gathers extended metadata at each stage of job execution by leveraging HPC workflow tools, thereby minimizing the need for fundamental system changes. By coordinating job submissions and data storage within a unified workflow framework, we can

automatically capture computational resources used, JobID, software name and version, user who executed job, and the time of execution. This metadata is annotated for all result data. We utilize WHEEL workflow tool [3], which already supports I/O operations to HPSS.

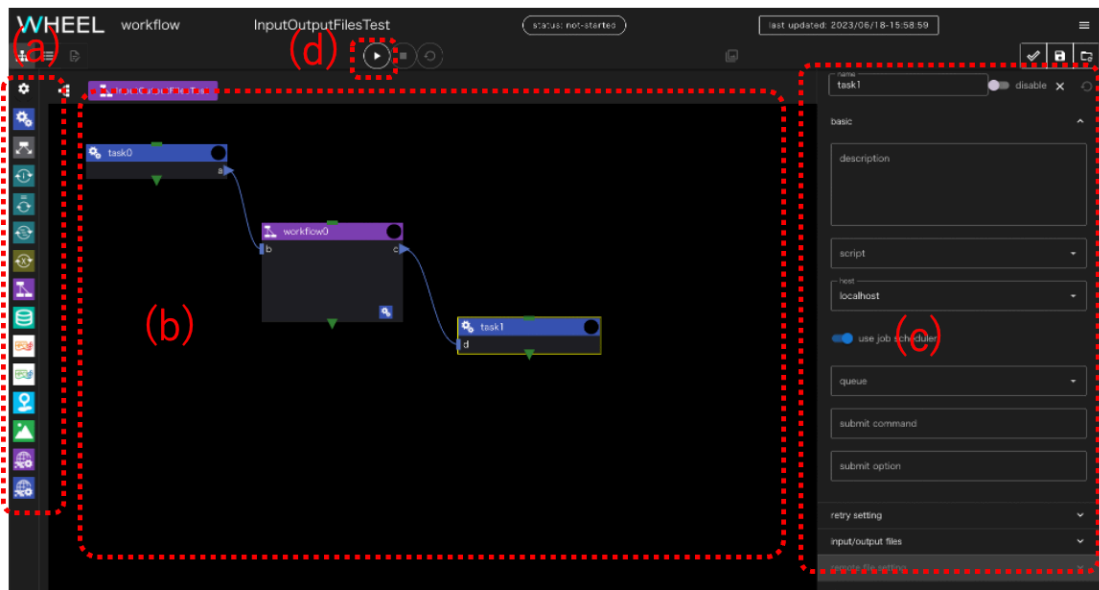


Figure 2: WHEEL's editor screen. (a): Component list, (b): Graph view, (c): Property pane, (d): Run-proj. btn.

By extending WHEEL to incorporate metadata assignment, we can automatically annotated expended meta-data, such as supercomputer names, Job IDs, software name and versions, create usernames. Existing research on workflow tools sometimes focuses on annotating extended metadata to entire workflows [4], but our goal is to directly annotate each data for clearer provenance and future integration with research data management (RDM) services. In practice, WHEEL automatically collects expanded metadata, storing these alongside file paths in NoSQL database. Because Gfarm has data archive function, so it's possible to annotate extended metadata to each packaged of data, allowing users to detailed information about the package-data later. This approach mitigates the risk of changing to Dark Data. Currently, our proof of concept writes extended metadata to a NoSQL database, however we are considering the following for future development.

1. Automated data publication, with digital object identifiers (DOI) assigned.
2. Extending Gfarm to embed extended metadata.
3. Integrating with external RDM services. (as explored in projects such as HOMER [5])
4. Extend Metadata visualization.

While this paper focuses on HPSS and WHEEL, our broader aim is to develop metadata management framework applicable across diverse HPC environments. Automating metadata annotation can simplify future data management for researchers and support practices such as RDM and open science. In particular, it can help reduce the accumulation of Dark Data in large-scale storage systems. Such efforts are expected to promote more equitable resource allocation and more efficient, transparent use of research data.

Reference

- [1] Tatebe, O., et al. (2010). Gfarm grid file system. *New Generation Computing*, 28(3), 257–275. <https://dl.acm.org/doi/10.1007/s00354-009-0089-5>
- [2] Bauer, D., et al. (2022). Revisiting data lakes: The metadata lake. In *Proceedings of the 23rd International Middleware Conference Industrial Track* (pp. 8–14). ACM. <https://doi.org/10.1145/3564695.3564773>
- [3] Kawanabe, T., et al. (2024). Introduction of WHEEL: An analysis workflow tool for industrial users and its use case on supercomputer Fugaku. In *Proceedings of the 2024 IEEE International Conference on Cluster Computing Workshops* (pp. 180–181). IEEE. <https://www.computer.org/csdl/proceedings-article/cluster-workshops/2024/834500a180/21EtRZFluLu>
- [4] Jain, A., et al. (2015). FireWorks: A dynamic workflow system designed for high-throughput applications. *Computational Materials Science*, 96, 118–124. <https://doi.org/10.1016/j.commatsci.2014.10.037>

[5] Chiapparino, G., et al. (2024). From ontology to metadata: A crawler for script-based workflows. INGGrid.
<https://www.inggrid.org/article/3983/galley/3912/download/>

Primary author: KANEYAMA, Hidetomo (RIKEN R-CCS)

Co-authors: KAWANABE, Tomohiro (RIKEN R-CCS); Mr HARADA, Hiroshi (RIKEN R-CCS)

Presenter: KANEYAMA, Hidetomo (RIKEN R-CCS)

Session Classification: Presentations Session 2: Data and Research & Data Science and Data Analysis

Track Classification: SciDataCon2025 Specific Themes: Infrastructures to Support Data-Intensive Research - Local to Global