DRAFT
Minutes: RHIC Data and Analysis Preservation Round Table

Date:May 8, 2025

Indico: https://indico.bnl.gov/event/26755/

Attendees: Eric Lancon, Jerome Lauret, Gene Van Buren, Vincent Garonne, David Gabor, Frank Geurts, Shigeki Misawa, Maxim Potekhin, Jamie Dunlop

Identified item for follow-up

Introduction (Eric)

Timeline and Infrastructure Planning

An inspirational timeline for data preservation was presented, starting from the consideration that the current computing infrastructure has a lifespan of five to seven years. It is proposed that long-term data preservation (DAP) will occur in two phases:

- **Phase I:** Development, Implementation, and Consolidation - including tool development, prototype workflows, metadata frameworks, and repository systems
- **Phase II:** Long-Term Preservation - beginning when the current infrastructure reaches end-of-life in 5-7 years.

Towards the end of Phase I, a comprehensive global reprocessing of RHIC data is proposed to create a consistent baseline dataset with unified calibration and reconstruction algorithms across all run periods. This dataset will serve as the reference dataset for the legacy of RHIC, ensuring consistency for future analyses.

Staffing and Organization

A preliminary collaborative organizational structure was outlined to support Phase I, proposing roles for three key stakeholder groups—the DAP core team, the experiments, and the computing center—whose active collaboration will be essential for the success of the DAP.

- Core Team: Including DAP Manager, Software & Workflow specialists, Repository Systems, AI Integration, Web Development, and Documentation & QA roles

- Experiments: Representatives from PHENIX, sPHENIX, and STAR to ensure alignment with experiment needs, plus training roles
- Computing Center: Providing data access support, technology monitoring, and liaison services

Participants discussed the staffing model, with Frank Geurts expressing concerns about the representation of experimental staff in the plan. To be followed up.

Data Volume and Preservation Levels

Updated data volume estimates were presented, following input from PHENIX and sPHENIX experiments.

- PHENIX: 20 PB RAW, 5 PB Analysis Objects, 10 PB for other archive
- sPHENIX: 160-300 PB RAW, 50-100 PB Analysis Objects (one processing), 50-100 PB for other archive
- STAR: No update, historical data yet to be fully quantified

Some discrepancies were identified between experiments in the reported numbers. *There is a need for clarification on what data should be preserved, particularly regarding historical and other data. Updates from STAR are needed.*

These estimates lead to the required storage need for the two preservation levels of:

- Level 3 (AO only): Estimated at 95-145 PB total
- Level 4 (RAW + AO + archive): Estimated at 435-675 PB total

It was noted that maintaining Level 3 data on disk will be challenging; further discussions are needed.

REANA Status Update

The retirement of the current PHENIX REANA instance has been announced by the computing center. The current instance is unmaintained, outdated, and lacks proper documentation.

With the OpenShift infrastructure at the data center stabilized, a new, smaller instance will be built on OpenShift pending hardware allocation from the computing center.

PHENIX expressed concerns regarding migration compatibility with CVMFS and access to external files. Further follow-up is required.

Open Data Portal Update (Vincent)

Vincent Garonne provided an update on the new deployment of the BNL Open Data Portal, which is showcased and accessible at https://rhic-dap02.sdcc.bnl.gov/ (requiring VPN access).

- The PHENIX dataset has been successfully imported from the CERN Open Data Portal to the BNL Open Data Portal
- The data is now hosted at BNL, and OSTI has assigned a new DOI
- A storage extension request has been submitted to SDCF and is currently pending approval

The need for volunteers from experiments to act as data curators was emphasized. These curators would be responsible for ensuring the quality of metadata, compliance with access policies, and the long-term preservation of data.

It was noted that examples of metadata are necessary for proper documentation and the searchability of datasets.

The discussion about the Portal included concerns regarding VPN requirements, with Frank Geurts noting that while VPNs work well for experts, they limit broader engagement from collaborations. *He suggested finding secure alternatives to provide wider access for testing and feedback. To be followed up on at a future meeting.*

Static Web (Maxim)

A presentation on static website generators highlighted their advantages over traditional Content Management Systems (CMS), particularly regarding long-term data preservation, as they require minimal ongoing maintenance. This project, developed over several years, features source code available on GitHub and has been successfully applied by HSF, NPPS, PHENIX, EICUG, and ePIC.

Key points included:

- Static website generators (such as Jekyll, Hugo, and Gatsby) are compiled once and deployed as static HTML, offering improved security, performance, and ease of maintenance.
- Content is managed in simple Markdown files, while structured data is stored in YAML or CSV format.
- The approach separates content from layout, allowing for easy theme changes without requiring modifications to the content.
- Integration with GitHub provides version control and preview capabilities

- The technology enables easy compliance with cybersecurity requirements.

The presentation highlighted the ePIC website (https://www.epic-eic.org) as a successful example of implementation.

Hardware Resources (Shigeki)

A comprehensive overview of the current state of hardware resources was presented, covering three main areas:

Tape Systems:

- The current infrastructure includes 7 legacy libraries (PHENIX/STAR) in the legacy data center with approximately 17 LTO-8 drives, and 4 new libraries (sPHENIX) in the new data center with around 100 LTO-9 drives.
- The new libraries are expected to be nearly full by the end of Run 25.
- *Open questions remain regarding the future cost and end-of-life timeline of the legacy libraries.*

Disk Systems:

- Current disk systems include approximately 95 JBOD Lustre units for sPHENIX, 13 JBOD dCache units for PHENIX, 11 JBOD Lustre units for STAR, and 20 JBOD XRootD units for STAR.
- Hardware RAID GPFS systems consist of 2 units for STAR and 2 units shared between sPHENIX and PHENIX.
- All equipment is in Building 725 MDH and follows a 5-year refresh cycle.
- Hardware RAID has a higher cost per terabyte compared to JBOD solutions.
- JBOD system costs (excluding servers) for disk components are projected to decline by approximately 10% annually.

CPU Resources:

- An estimated 5% annual increase in instructions per cycle (IPC) was noted.
- Operational costs for CPU resources rise with age, impacting their effective service life; a 7-year refresh cycle is assumed.

Some participants questioned the accuracy of certain figures presented, and it was agreed that the report should be considered a work in progress. *Future discussions will aim to refine these estimates, focusing on costs and timelines for hardware resources related to the data preservation effort and the start of Phase II*. Jamie emphasized that there is currently no budget allocated for hardware refresh.