Report from the 4th DPHEP workshop

Place and date: CERN, 2-3 October 2024
Agenda: https://indico.cern.ch/event/1432766/timetable
Workshop goals:

- Review, update, discuss, plan
- Progress report?
- CHEP2024 Contribution
- EPPSU Contribution/Statement
- Organisation and future activity

Workshop highlights

- Significant progress has been made since the last workshop, both on data preservation and opening of data
- While CMS has pioneered the publication of open data, the other LHC experiments are rapidly catching up now
- Open data policies are increasingly applied also beyond LHC
- The release of a community version of CERNLIB enabled the full resurrection of the DELPHI and OPAL software stacks; DELPHI data has been opened recently, and corresponding discussions have started in OPAL.To be noted that this was and is only possible thanks to a significant amount of voluntary works going into these projects.
- The distribution model of shared tools such as CERNLIB and OpenPHIGS is worth a review
- EDM4HEP is a proposed data format for future colliders, including FCC-EE. Other LEP experiments are encouraged to evaluate and eventually adopt the EDM4HEP format conversion efforts as ALEPH is doing, which may have a direct impact for FCC-ee developments
- Rapid changes in the computing environment remains a possible threat. An example is the unknown future of X11 and the related libraries.
- Two experiments reported to ensure DP by transferring their data to successor projects and collaborations
- In many contributions, continued funding of DP was mentioned as an issue. An example is the BaBar experiment, whose software is running on outdated hardware and there is no funding to replace it
- The report from the ICFA panel was encouraging, following a bottom up approach starting from the people involved in the practical work.
- Ongoing review of CERN OC3, which may offer some opportunities for DPHEP.

Detailed notes:

Opening:

Opening speeches by Joachim Mnich, director for research at CERN, and Cristinel, chair of the DPHEP collaboration.
Joachim kindly opened the workshop with a very supportive speech. He mentioned possible hidden treasures waiting to be discovered in archived data, and recalled how and why DPHEP was founded about 14 years ago.
Cristinel reminded that archiving data is not enough, it needs to be accessible to be useful. He also summarised the workshop goals.

LEP session

ALEPH:
The LEP session was started by a contribution from the ALEPH collaboration. ALEPH is working on a prototype to convert their DST data to EDM4HEP, which enables FCCee to, for the first time, run on real data. The last build of the software stack was for SLC4, and it kept running until SLC6. There are currently no plans to migrate the software stack to newer operating system models. Instead, the original software is run on containers, e.g. on LXPLUS. Data is written in text format, which in a second step is used to create EDM4HEP records. ALEPH collaborates with MIT, and work is ongoing to attempt to reproduce published results, using the new format.
Once fully validated, the EDM4HEP formatted data might be made available via the open data portal. There is the wish that more LEP experiments join the effort.

CERNLIB:
During the last workshop, CERNLIB was identified as a critical dependency for most of the LEP experiments while no longer being supported by CERN. In a community driven effort, the last sources were imported into CERNs gitlab repository, and external 3rd party patches were collected, reviewed and partly applied. A CMake based modernised build infrastructure has been added, as well as continuous integration to automatically create builds on various platforms and run basic tests. Known use cases for this new version of CERNLIB include the DELPHI and the OPAL software stacks, as well as several external users of e.g. PAW. Possible future risks include developments related to X11 replacements and related libraries like Motif which are used by CERNLIB.

DELPHI:
The DELPHI software stack has been rebased on the resurrected community version of CERNLIB, and runs many modern Linux distributions. The data preservation level is 4, as defined by DPHEP in https://cds.cern.ch/record/1450816/files/arXiv:1205.4667.pdf. Data is

located on EOS and on CTA at CERN, and there is at least one external copy. The software has been moved to CERN gitlab, binaries are distributed via CVMFS. Analysis preservation is a weak point as this had not been on the radar while the experiment was still operational. The event display has been resurrected based on open source software only, removing a dependency on commercial 3d libraries. DELPHI has recently changed their data access policy towards opening the data, and work is in progress to make it accessible via the open data portal.

OPAL:
Contrary to DELPHI and ALEPH, OPAL data are not open, but discussions are ongoing inside OPAL, showing openness towards this idea. The full OPAL stack has been resurrected and is at a similar level as DELPHI. As in the DELPHI case, the stack runs on many modern Linux distributions, the software stack has been moved to CERN gitlab, binaries are built via CI, and published via CVMFS. Data is on tape and on EOS. The event display, which had been given up on many years ago due to a dependency on a commercial software package, has been fully resurrected using open source software only.

The PHIGS story:
The LEP session was concluded by a contribution about an open source replacement for the commercial 3d libraries which OPAL and DELPHI relied on for their event displays. This was made possible by the publication of a pre-alpha level, open source prototype implementation of PHIGS, which has been extended by voluntary work to support the event displays of DELPHI and OPAL. At the time of the workshop, the code was still maintained within the DELPHI software tree on CERN gitlab. In the meantime, it has been cleaned up and finally moved to https://github.com/cern/OpenPHIGS, and the experiments using it are now moving to this distribution.

HERA:

ZEUS data is still available in the form of root ntuples. This work has been done many years ago already, and has proven to be very successful as it survived the transition from SLC5 up to current RHEL9 like operating systems. 30% of the ZEUS papers were created after the end of data taking of the experiment, and past projections about the number of papers have proven to be amazingly close to reality. The system is so easy to use that it has been successfully applied by theorists, with no ZEUS background.

H1 software is based on Fortran in the core, and C++ and ROOT as user facing interfaces. While the transition to RHEL9 is currently still ongoing, the software stack can be used from containers without problems. H1 is still active, and has already published 3 papers this year.

PETRA:

There was an update on the status of the JADE software stack, which had been resurrected in an effort taking almost one decade. JADE data was originally stored on 6500 tapes. The data and software are about to be published via the open data portal at CERN. There are 3500 pages of scanned log file entries, documenting the data taking period. The software currently runs in containers, uses the community CERNLIB in 32bit. A 64bit version has not been tried, mainly due to different priorities.

RHIC:

The PHENIX experiment at RHIC was operational until 2016, and is still publishing results. The software is currently being migrated to ALMA9. PHENIX makes extensive use of Zenodo, where one can find about 700 items related to it. One of the pain points was the extensive use of PHP, which experienced many changes over the past years. On the analysis preservation side, they focus on the two analyses with highest impact. One approach for analysis preservation is the maintenance of a web page with detailed instructions and steps needed to perform to come to the published result. During the discussion it was pointed out that it is important to start preservation efforts during the lifetime of the experiment, which is especially true for analysis preservation. Thinking about possibly opening the data early on while the collaborations are still functioning is also a plus.

SLAC:

The BABAR experiment, originally located at SLAC, had to move out after the end of data taking. The data size is about 1.5PB, and is located at GridKA in Karlsruhe, with copies at CERN and at IN2P3. The software is hosted at UVIC, and there is still an Oracle database at SLAC which needs to be dumped. The hardware hosting the software stack is old and out of warranty, however, there are no funds to replace it. The analysis framework does not compile on 64 bit, and there is a dependency on older software which makes the migration to recent operating systems difficult.

KEK: Belle and Belle II

Belle collected about 1.6PB of data between 1999 and 2010. Belle and Belle II are two different collaborations, albeit with some significant overlap. A merge of the Belle data into Belle II was approved, and the format of Belle data converted into Belle II format. The Belle software does not run (yet) on RHEL9, but can be run in CC7 containers. Replacement of the KEK's main computer system, which involves data migration, happens approximately every 4 years.

Belle II started in 2019, and continues to collect data. The run 2 phase just started. Belle II has already collected about 3PB of data. Distributed data management is based on Rucio. There was a DP workshop in 2022, a follow-up still to be done.

BESIII:

BESIII is actively taking data, and started to think about data preservation. They are aiming for a level4 preservation mode, preserving raw data. The plan includes an AI based data ecosystem, which shall allow users to get started quickly with the data analysis.

CERN Open Data Portal

The next talk was more on open data development. The number of records in the CERN open data portal has increased by a factor of 4 recently, and new experiments have joined, namely JADE and DELPHI. Both are expected to hit production in the coming weeks. The portal, which can now as well be reached under the cern domain (https://opendata.cern), experiences about 7 updates per year. Work is in progress to support cold data. The idea is that for large data sets only a part is kept on disk, while the rest is stored on tape and is recalled only on request. Several challenges need to be addressed to make this work, for example the need to notify users when the requested data is finally available. Work is going as well into the creation of KPIs and detailed monitoring. Standard CERN IT best practices are applied, and the portal makes extensive use of other CERN IT provided services, such as the monitoring system and the OpenSearch service.

REANA

REANA is an open source platform for reproducible containerised data analyses. It is provided at CERN as a service, but several other installations outside of CERN exist, both in particle physics (e.g. University of Chicago) and in other scientific disciplines (e.g. AIP Potsdam in astronomy). REANA comes both with a command-line client and a web frontend interface allowing researchers to submit, inspect and manage their analysis workflow runs. A recently added feature allows users to share workflows with colleagues. REANA integrates with GitLab with the idea of allowing users to use "preproducible" workflow practices from the start of their analyses, facilitating their future reproducibility. One use case of relevance to the data preservation activities is "continuous reuse", allowing data curators to periodically execute data usage examples (e.g. daily or weekly) in order to ensure that the data usage examples keep being valid and that the preserved data can be reused in the future.

CERN analysis preservation framework

"The CERN Analysis Preservation (CAP) platform manages and safeguards research outputs from the collaborations by capturing knowledge, data, code, and documentation. It features advanced search functionality, enabling easy access to preserved analyses, and simplifies the preservation process for researchers. Recent updates include new tools for preservation administrators and the ability to easily share analysis components, enhancing collaboration within the CERN community (access to it requires a CERN account for now). By streamlining research preservation and collaboration, CAP ensures that valuable scientific insights are readily accessible, driving innovation and discovery forward.

CERN Open Data: Policy to implementation

After the release and endorsement of the LHC policy file by the large LHC experiments, we have entered the implementation phase of the policy. Meanwhile, the idea is extended to the smaller LHC experiments, as well as non-LHC experiments. In some cases, due to differences in the communities, the original policy is not one to one applicable to these experiments, requiring the creation of different policies which are inspired by the one endorsed by the LHC experiments (not directly mentioned during the workshop, but this is the case for DELPHI). The working group meets once per year now, and the focus is generally on L3 preservation as storing the raw data is too expensive both in resources and required knowledge. About 5.5PB of data is currently stored in the OD portal, and resources have been granted for the first 5 years until 2025. Discussions for the next period are ongoing.

LHC:

Alice:
The Alice experiment reported that they are deploying a new software framework featuring their A02d format. It allows them to shrink the size of the data significantly. A data set taking 1.2PB initially now requires 52TB of data only. The total size of Alice data is determined by the trigger rate, which increases for each run period. As of 2030, storing Alice open data will require several PB per year which is not sustainable, therefore some rethinking will have to be done in the next years.

ATLAS:
ATLAS released Open Data for Research for the first time in July, putting out the 2015 and 2016 proton-proton collision dataset and significant accompanying simulation, which totals 65 TB. The documentation of these data is an ongoing process, including ensuring that documentation is sufficient to understand the complexities of systematic uncertainties that are required for scientific results and the limitations of the PHYSLITE format in which the open data were released. ATLAS is considering a workshop on Open Data in 2025, similar to what CMS did in 2024. The collaboration is also actively discussing the long-term preservation of Run 1-3 data, for which these open data constitute the minimum level of support.

LHCb:
LHCb makes extensive use of gitlab continuous integration for analysis preservation. While Run 1 open data has been fully published already, this model is not sustainable for the following years because it would require way too much storage space. Therefore, LHCb is working on an Ntuple Wizard, which uses standard experiment resources and data to create ntuples. Work on this is in progress, in collaboration with the CERN IT Opendata team.

CMS:
CMS has been pioneering open data. Usage of open data can be tracked via DOIs. Their nano-aod format is a big step forward as it does not require CMS specific software to be read. It's based on flat ROOT trees. Open data is also meant to be used for educational purposes. Examples are given at different levels on the portal. Annual workshops are organised on the use of open data in 2020. CMS is very interested in using cold storage to cover a fraction of their open data storage requirements.

Other

Antares:
This is a neutrino experiment, located in the mediterranean sea, which has been taking data between 2000 and 2022. It has been superseded by KM3NET, consisting of 2 sites, ORCA and ARCA, which are currently under construction. Despite different collaborations, the DP strategy is that Antares will hand over their data to the follow-up collaboration. There is about 600 TB of data, located in CCLyon. They plan to keep the legacy software. The main issue is the longer term funding.

Multi-experiment talk: punch
PUNCH allows combining data from different experiments, and doing a common analysis. This was demonstrated on open data from ATLAS and CMS, running a Higgs search and reproducing the discovery of it.

Data lifecycle - what ?
This was a presentation related to the work of the ICFA panel. Preservation is not the end but needs resources long term. The panel's mandate is to do meta work, and come up with recommendations, bridging the past to the future. It was pointed out that the DP community is rather small which makes it more difficult to go to big conferences. In this field, the inreach is as important as the outreach, raising awareness inside the collaborations to think about DP early on. The goal is to recommend future directions, which should be based on experience, following a bottom up approach, exploiting experiences of the people doing the actual work, concrete and specific. A survey has been set up on the conference page, and it is still possible to contribute to it.

Discussions:

- There may be options for archiving via the CERN library
- Advertising the DPHEP status is important, e.g. it may not be widely known that LEP data is still available and accessible
- Advertising for educational purposes was briefly discussed. One option is to reach e.g. teachers via the CERN teacher's program
- An article in the CERN courier should be considered, to give an update on the current status
- Gerardo Ganis kindly agreed to jump in for a presentation at CHEP
- The Data LifeCyle panel will initiate a white paper on good practices, DPHEP will contribute
- DPHEP will contribute to ESPPU (a short document) first draft in january; deadline March 31st, 2025