1st Conference on Research Data Infrastructure

Poster presentations II

https://doi.org/10.52825/CoRDI.v1i.261

Published: 07 Sept. 2023

Distributed Computing and Storage Infrastructure for PUNCH4NFDI

C. Wissing[1] [https://orcid.org/0000-0002-5090-8004], B. Bheemalingappa Sagar[1], M. Blank-Burian[2], A. Drabent[3], S. Fleischer[4], O. Freyermuth[5], M. Giffels[6] [https://orcid.org/ 0000-0003-0193-3032], A. Henkel[7], M. Hoeft[3] [https://orcid.org/0000-0001-5571-1369], J. Künsemöller[8], N. Malavasi[9], C. Manazano[10], B. Roland[6] [https://orcid.org/0000-0003-3397-6475], H. Simma[1], D. Schwarz[8] [https://orcid.org/0000-0003-2413-0881], K. Schwarz[1] [https://orcid.org/0000-0002-0800-2743], N. Suvvi Neelakantaia[1], L. Vomberg[5], C. Voss[1], V. Vybornov[10], M. Wigard[2], and S. Wozniewski[11]

[1]Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany

[2]Westfälische Wilhelms-Universität Münster, Germany

[3]Thüringer Landessternwarte, Tautenburg, Germany

[4]Helmholtzzentrum für Schwerionenforschung GmbH, Darmstadt, Germany

[5]Rheinische Friedrich-Wilhelms-Universität Bonn, Germany

[6]Karlsruher Institut für Technologie KIT, Karlsruhe, Germany

[7]Universität Mainz, Germany

[8]Universität Bielefeld, Germany

[9]Ludwig-Maximilians-Universität München, Germany

[10]Forschungszentrum Jülich, Jülich, Germany

[11]Georg-August-Universität Göttingen, Germany

Abstract: The PUNCH4NFDI consortium brings together scientists from the German particle physics, hadron and nuclear physics, astronomy, and astro-particle physics communities to improve the management and (re-)use of scientific data from these interrelated communities. The PUNCH sciences have a long tradition of building large instruments that are planned, constructed and operated by international collaborations. While the large collaborations typically employ advanced tools for data management and distribution, smaller-scale experiments often suffer from very limited resources to address these aspects. One of the aims of the consortium is to evaluate and enable or adopt existing solutions. Instances of a prototype federated and distributed computing and storage infrastructure have been set up at a handful of sites in Germany. This prototype is used to gain experience in running of scientific workflows to further guide the development of the Science Data Platform, which is an overarching goal of the consortium.

Keywords: Distributed Computing, NFDI, Particle Physics, Astro-Physics, Data Management

The PUNCH4NFDI (Particle, Universe, NuClei & Hadron physics) consortium [1] brings together scientists from the German particle physics, hadron and nuclear physics, astronomy, and astro-particle physics communities to improve the management and (re-)use of scientific data from these interrelated communities. The PUNCH sciences have a long tradition of building large instruments that are planned, constructed and operated by international collaborations. Flagship examples are the particle physics experiments at the Large Hadron Collider (LHC) at CERN, which produce tens of petabytes per experiment per year, complemented by a similar amount of simulated Monte Carlo data. To store, process and analyse this data, the Worldwide LHC Computing Grid (WLCG) [2] has been established, comprising more than 160 computing centres around the world, providing more than one million CPU cores, almost 1 exabyte of disk space and around 1.5 exabytes of archival storage. By the end of the decade, the upgraded HL-LHC (High Luminosity LHC) is expected to begin taking data, requiring about an order of magnitude more CPU and storage capacity. Astronomical facilities do not yet reach the same scale but will soon catch up. Currently the radio interferometer LOFAR already relies on a federated Long Term Archive (LOFAR LTA) that holds over 50 petabytes of intermediary data products. The flagship (radio) astronomy project, the Square Kilometre Array (SKA), will have similar requirements as LHC-HL around the same time.

Although the flagship experiments are the most visible, to answer the scientific questions addressed in the PUNCH communities, many medium or small experiments are needed, each addressing a very specific scientific measurement. The amount of data produced in these experiments is typically much smaller, but issues of data management and long-term archiving are often not addressed as systematically as in the large experiments. It should be noted that the astronomy community has a long tradition of making scientific data open, including the definition of data formats and exchange protocols.

PUNCH4NFDI aims to combine the expertise of the sub-communities and to further develop the tools and to facilitate access to resources for storing, publishing and analysing data and software across all PUNCH4NFDI communities. For example, smaller experiments can benefit from the tools and methods developed for the flagship experiments. For some of the tasks, existing solutions can be adopted, while for others new solutions will be required. PUNCH4NFDI provides the forum for the knowledge transfer. Another important aspect is the ability to share existing and future computing infrastructure across the boundaries of experiments and communities for efficient use.

In this contribution, the focus is on the Compute4PUNCH and Storage4PUNCH concepts that are being developed to provide seamless and federated access to the wide variety of compute and storage systems provided by the participating communities to meet their diverse needs. Figure 1 depicts an architecture sketch. Both concepts include state-of-the-art technologies such as a token-based AAI for standardised access to compute and storage resources. For a first implementation of a PUNCH AAI, the consortium has adopted the Helmholtz AAI. Due to the diverse international participation, the AAI has to interact with existing or developing AAI structures such as the SciTokens [3] used in WLCG or EGI-Checkin and other EOSC-related implementations.

In a prototype distributed infrastructure setup, heterogeneous HPC, HTC and cloud compute resources provided by the community are dynamically and transparently integrated into a federated HTCondor-based overlay batch system using the COBaLD/TARDIS resource meta-scheduler [4]. Traditional login nodes and a JupyterHub provide entry points to the full landscape of available compute resources. Scientific software is dis-

tributed using the latest container technology and the CERN Virtual Machine File System (CVMFS) [5], which has been proven to allow replication of once centrally installed software to tens of thousands of globally distributed compute nodes via layers of HTTP Squid caches.

In Storage4PUNCH, community-supplied storage systems, mainly based on dCache [6] or XRootD technology, is federated into a common infrastructure using methods well established in the wider HEP community. Existing solutions that allow coherent management of files, which can be grouped into datasets, across multiple geographical locations will be evaluated. These systems provide methods for 'technical metadata' to manage the files, such as file sizes, checksums or file locations. Systems capable of describing the data more in terms of scientific content are also within the scope of this evaluation and prototyping. Existing caching technologies will also be evaluated for deeper integration. The combined Compute4PUNCH and Storage4PUNCH environment will enable a wide range of researchers to perform resource-intensive analysis tasks.

The current prototype spans a handful of sites in Germany that provide storage and computing resources. In addition to the purely technological evaluation of the integrated components, the first real scientific workflows are being carried out. The experience gained will help to guide further integration steps and identify the developments needed to meet the requirements of scientific applications. Over time, the prototype should incorporate more and existing data sources and thus provide a more uniform access to them. Ultimately, this distributed infrastructure should become the backbone of the Science Data Platform for the participating community, which is one of the overarching goals of the PUNCH4NFDI consortium.

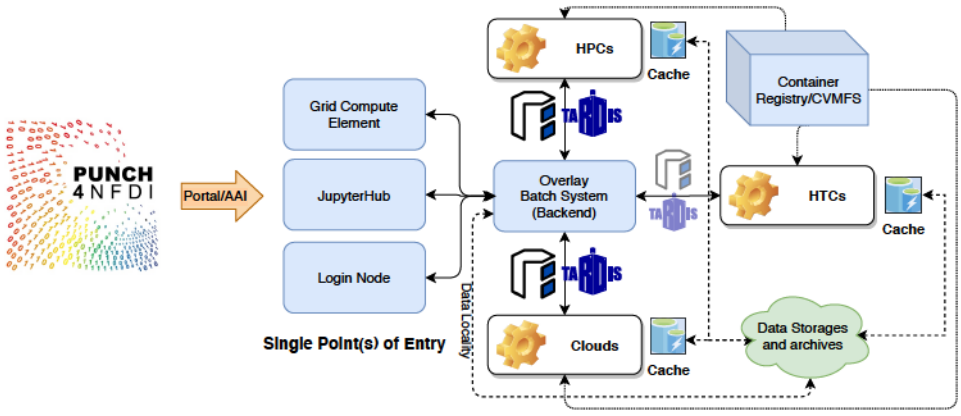

Figure 1. Architecture diagram of the distributed computing and storage PUNCH4NFDI infrastructure.

Competing interests

The authors declare that they have no competing interests.

Funding

PUNCH4NFDI is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Projektnummer 460248186.

Acknowledgements

The authors want to thank the entire PUNCH4NFDI consortium members for the fruitful discussions and the exchange about requirements and experiences with the prototype infrastructure.

References

[1] The PUNCH4NFDI Consortium, *Punch4nfdi consortium proposal*, version v1 without funding tables, This is the version documenting the work plan at the proposal stage. The reduction in funding led to a re-shaping of the work programme that is documented elsewhere., Sep. 2020. DOI: 10.5281/zenodo.5722895. [Online]. Available: https://doi.org/10.5281/zenodo.5722895.

[2] I. Bird, "Computing for the large hadron collider," *Annual Review of Nuclear and Particle Science*, vol. 61, no. 1, pp. 99–118, 2011. DOI: 10.1146/annurev-nucl-102010-130059. eprint: https://doi.org/10.1146/annurev-nucl-102010-130059. [Online]. Available: https://doi.org/10.1146/annurev-nucl-102010-130059.

[3] Weitzel, Derek, Bockelman, Brian, Basney, Jim, Tannenbaum, Todd, Miller, Zach, and Gaynor, Jeff, "Capability-based authorization for hep," *EPJ Web Conf.*, vol. 214, p. 04 014, 2019. DOI: 10.1051/epjconf/201921404014. [Online]. Available: https://doi.org/10.1051/epjconf/201921404014.

[4] M. Giffels, M. Fischer, A. Haas, *et al.*, *Matterminers/tardis: The escape*, version 0.7.0, Feb. 2023. DOI: 10.5281/zenodo.7680164. [Online]. Available: https://doi.org/10.5281/zenodo.7680164.

[5] J. Blomer, P. Buncic, G. Ganis, N. Hardi, R. Meusel, and R. Popescu, "New directions in the cernvm file system," *Journal of Physics: Conference Series*, vol. 898, no. 6, p. 062 031, Oct. 2017. DOI: 10.1088/1742-6596/898/6/062031. [Online]. Available: https://dx.doi.org/10.1088/1742-6596/898/6/062031.

[6] T. Mkrtchyan, O. Adeyemi, P. Fuhrmann, *et al.*, "dCache - storage for advanced scientific use cases and beyond," *EPJ Web Conf.*, vol. 214, A. Forti, L. Betev, M. Litmaath, O. Smirnova, and P. Hristov, Eds., p. 04 042, 2019. DOI: 10.1051/epjconf/201921404042.