

CERNLIB status

Andrii Verbytskyi¹, Ulrich Schwickerath²

¹ Max-Planck-Institut für Physik, 80805 Munich, Germany.

² CERN, 1211 Meyrin, Switzerland

E-mail: Ulrich.Schwickerath@cern.ch

Abstract.

We present a revived version of CERNLIB, the basis for software ecosystems of most of the pre-LHC HEP experiments. The efforts to consolidate CERNLIB are part of the activities of the Data Preservation for High Energy Physics collaboration to preserve data and software of the past HEP experiments.

The presented version is based on CERNLIB version 2006 with numerous patches made for compatibility with modern compilers and operating systems. The code is available in the CERN GitLab repository with all the development history starting from the early 1990s. The updates also include a re-implementation of the build system in CMake to ensure CERNLIB compliance with the current best practices and to increase the chances of preserving the code in a compilable state for the decades to come.

The revived CERNLIB project also includes updated documentation, which we believe is a cornerstone for any preserved software depending on it.

1. Introduction

CERNLIB [1] is a set of C and Fortran libraries developed between the 1970s and 2000s. It was and is widely utilized in many high energy physics (HEP), nuclear physics, and astro-particle physics experiments for data analysis for almost four decades. As a result, the experiment-specific software for many old experiments and even the access to the collected data depends on the availability of a functioning CERNLIB. For instance, the LEP experiments [2] ALEPH, DELPHI, L3, OPAL, as well as non-CERN based experiments, such as JADE [3], heavily depend on CERNLIB.

The development of the CERNLIB, historically coordinated by the CERN IT department, ended officially in 2006. Since that time, the official sources, released under GPL license (with exceptions), were no longer updated. As a consequence, the code was picked up by numerous groups and communities around the globe, and further patches and adaptations to new platforms and compilers were done in an uncoordinated manner.

Given the importance of the CERNLIB for the old experiments, the DPHEP [4] collaboration has planned and successfully executed a project to consolidate the CERNLIB developments that happened after 2006 to assure a safe future of CERNLIB.

2. CERNLIB status before 2022

The code base of CERNLIB contains 1600 kLOCs of Fortran77, and 500 kLOCs of C89. The code written in C is used primarily to provide an interface to system libraries. CERNLIB is organized in sub-packages according to their functionality. Briefly, the CERNLIB sub-packages

included basic utility subroutines and functions (KERNLIB), mathematical routines and random number generators (MATHLIB), the graphics libraries (GRAFLIB), physics analysis tools such as the PAW [5] library (PAWLIB), as well as incorporated software. Namely, the Monte-Carlo event generators (MCEGs) such as Pythia6 [6] or ARIADNE [7] were included in the MCLIB sub-package. The GEANT3 [8] simulation toolkit with a G-FLUKA [9] interface were included as GEANT321 sub-package.

The original build system of CERNLIB consisted of imake [10] configuration files and shell scripts.

The last release of CERNLIB was officially supported on a dozen of different Unix/Unix-like platforms and Windows. Older versions of CERNLIB were also supported on many other platforms, e.g. VMS. However, the level of CERNLIB support on these platforms varied and some sub-packages had only limited support for 64-bit platforms.

Therefore, at the end of the 2010s, when the 64-bit platforms started to dominate HEP computing, numerous efforts were made to allow the usage CERNLIB on the most popular platform in HEP – Linux x86_64.

After 2006 the CERNLIB builds were maintained by the Fedora (till 2016) and Debian projects. These projects also collected hundreds of patches for CERNLIB. A lot of those patches originated from the activities of Kevin McCarty to package the CERNLIB for Debian Linux. The second significant group of patches was created by Harald Vogt and his collaborators from the DESY-Zeuthen group. In addition, input from private builds by the DELPHI collaboration was considered. All those changes became the basis for the consolidation of CERNLIB in 2022.

3. CERNLIB consolidation timeline

The principal decision to consolidate and revive CERNLIB was taken in December 2021. In January of 2022, the collection of all information on the CERNLIB patch-sets and forks started. In parallel, the original CVS repositories (one per sub-package) were recovered and converted into git repositories. The resulting git repositories were put into <https://gitlab.cern.ch>. Later, the separate repositories were merged into one git repository, and the build scripts were added.

Between February and June 2022 the collected patch sets were carefully reviewed, broken down into smaller updates where needed, and added one by one to the code base.

With a sufficient number of patches it became possible to compile CERNLIB on modern and historical Linux systems on x86_64 architecture with the GCC compilers. To manage the increasing number of supported systems a continuous integration system (CI) based on GitLab CI was introduced. The development was focused on the Linux i686 and Linux x86_64 platforms with GCC at this time, as only those systems had solid support in the CERNLIB native build system.

In June of 2022 the CERNLIB build system, based on imake, with a set of build scripts was complemented by an alternative CMake based build system. The comparison between the builds performed with imake and CMake setups was done using the compilation databases produced by the bear[11] tool on Linux systems. It was also checked that the resulting installations from both builds produce the same set of files.

The adoption of CMake has allowed support for a wider range of compilers and target systems to be explored. CERNLIB was compiled with the Intel One API compilers on Linux x86_64, NVidia HPC compilers on Linux x86_64, GNU compilers on Mac OSX x86_64, GNU compilers on Solaris x86_64, and GNU compilers on Linux aarch64. The variety of compilers was beneficial for a more efficient elimination of bugs and incompatibilities with modern operating systems.

Further developments were related to the improvements in the CI, which tests the CERNLIB builds starting from very old systems (Scientific Linux 3) to very modern (Fedora 37). After gradual improvements and a certain amount of bug fixes, the updated CERNLIB was used for

tests with the DELPHI software, by using it for a first proof of concept to run the entire DELPHI software stack in 64-bit mode on Linux.

After the preliminary tests, the first updated version of CERNLIB has been made available on the official CERNLIB site in November 2022 at <https://cernlib.web.cern.ch/cernlib/>. Along with this, the CERNLIB versioning schema was changed to a chronological one, with the date of the CERNLIB release in the name including release day, month and year.

After the release, the updated CERNLIB was added to the repositories of ARCHLinux, Homebrew [12] (hep-homebrew), HEPrpm [13] and spack [14], excluding any non-free licensed software parts. In addition, a release on CVMFS for LCG is planned for a future release.

4. CERNLIB validation

The validation of the new CERNLIB release is a question of highest priority. The usage of modern compilers allowed fixing some obvious bugs. The CI system helped to automate the testing. However, the most important results will come from the usage of CERNLIB with the preserved software.

It is expected that after the adoption of CERNLIB more feedback will be obtained from the users. Therefore, the suggested strategy for the validation is to use automation where it is possible and “to validate vby using it”. This is a community effort.




Figure 1. Example of a JADE event display, produced with CERNLIB 2022.

5. Impact on data preservation

As mentioned already, CERNLIB is a pre-requisite for continued full access to LEP&JADE data. The DELPHI, JADE, and OPAL software stacks were successfully recompiled with the CERNLIB 2022. While validation of software compatibility for JADE is already finished while it is still ongoing in the case of DELPHI and OPAL. No blocking issues have been reported so far.

The availability of 64-bit builds is of particular importance. One reason is that support for 32-bit is vanishing. As an example, Canonical no longer provides 32-bit Motif [15] development packages for 32-bit on recent Ubuntu versions.

The DELPHI and OPAL event displays were originally based on a commercial 3-D software package. In the past years, there was little to no hope to be able to revive these, mainly due

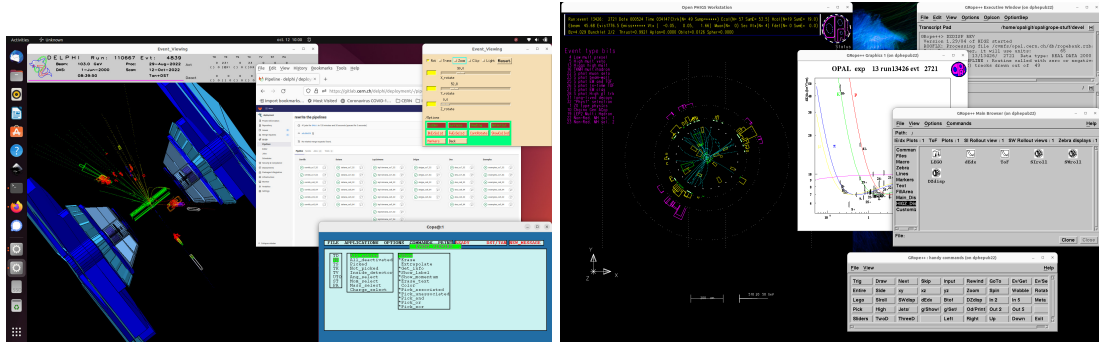


Figure 2. Event displays of DELPHI (left) and OPAL (right), produced with CERNLIB 2022 in 64-bit using modern compilers and Linux versions.

to licensing and compatibility issues of the graphics libraries with modern versions of glibc [16]. This dependency could recently be removed by extending a free, OpenGL-based back-end to the needs of data preservation. This only works properly in 64-bit mode though, and thus requires a full 64-bit CERNLIB and experiment software stack. Good progress has been made in this direction. DELPHI already released a new version of it, and for OPAL a proof of concept exists that works on Ubuntu 22 and CentOS 8. Fig. 2 shows screenshots of the two event displays.

Modernisation of CERNLIB enables the adoption of modern tools, e.g. automatic building and testing for the whole software stacks depending on CERNLIB. In the case of DELPHI, a CI has been implemented which allows rebuilding of the whole stack for a variety of operating system versions on changes, including basic checks for completeness of the stack. The resulting artifacts can directly be used for deployment. A similar approach is planned for OPAL in the near future.

6. Further Plans

While initial testing looks very promising, and no major issues have been identified so far, more evaluations are needed to confirm that the new code base works and can be used to reproduce old results from the different experiments. This implies further advertising of it's availability and timely bug fixing as reports come in. Maintenance will have to continue, on a voluntary basis, to adapt the code and build infrastructure to upcoming changes. This includes upstream software evolution, for example in terms of compilers, and also changes in the computing architecture, in particular new generations of CPUs.

7. Conclusions

The presented approach to revive CERNLIB combines experiences from various experiments that depend on it with modern tooling, allowing to automate the creation of the libraries and tools. This makes it easy to be ported to additional platforms while maintaining backward compatibility, making it fit for data preservation for the coming 10 years and beyond.

Acknowledgements

We are grateful to the authors of patches to CERNLIB that were created since the official support of CERNLIB was terminated and to the Fedora and Debian projects that maintained the builds of CERNLIB in their repositories.

We are grateful to the authors of the Pythia6, ARIANDE, LEPTO [17], JETSET [18] and fritiof [19] packages who agreed to re-license their codes under GPL license for CERNLIB.

References

- [1] R. Brun et al., CERNLIB: short writeups. <https://cds.cern.ch/record/450356>.
- [2] O. Callot and P. Charpentier, The LEP detectors, *C. R. Acad. Sci.*, 4 **3** (2002), 1131.
- [3] S. Bethke and A. Wagner, The JADE Experiment at the PETRA e^+e^- collider - history, achievements and revival (2022), [arXiv:2208.11076](https://arxiv.org/abs/2208.11076).
- [4] Z. Akopov et al., Status Report of the DPHEP Study Group: Towards a Global Effort for Sustainable Data Preservation in High Energy Physics (2012), [arXiv:1205.4667](https://arxiv.org/abs/1205.4667).
- [5] R. Bock et al., PAW: Towards a physics analysis workstation, *Comput. Phys. Commun.* **45** (1987), 181.
- [6] T.Sjostrand, S. Mrenna and P.Z. Skands, PYTHIA 6.4 Physics and Manual, *JHEP* **05** (2006), 026, [arXiv:hep-ph/0603175](https://arxiv.org/abs/hep-ph/0603175).
- [7] L. Lonnblad, ARIADNE version 4: A Program for simulation of QCD cascades implementing the color dipole model, *Comput. Phys. Commun.* **71** (1992), 15.
- [8] R. Brun et al., GEANT3 (1987).
- [9] A. Fasso et al., FLUKA: Present status and future developments, *Conf. Proc. C* **9309194** (1993), 493.
- [10] X.org, C preprocessor interface to the make utility , 2022. <https://www.x.org/releases/individual/util/>.
- [11] L. Nagy, bear, 2023. <https://github.com/rizotto/Bear>.
- [12] M. Howell et al., The Missing Package Manager for macOS (or Linux), 2022. <https://brew.sh>.
- [13] T. Hahn and A. Verbytskyi, Deployment of High Energy Physics software with a standard method, *J. Phys. Conf. Ser.* **2438** (2023), 012024, [arXiv:2210.17261](https://arxiv.org/abs/2210.17261).
- [14] T. Gamblin et al., The spack package manager: bringing order to hpc software chaos, *SC '15: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–12. (2015), doi:10.1145/2807591.2807623.
- [15] Integrated Computer Solutions, Inc., "Motif Window Manager", 2023. <https://motif.ics.com/>.
- [16] glibc developers, 2023. <https://www.gnu.org/software/libc/>.
- [17] G. Ingelman, A. Edin and J. Rathsman, LEPTO 6.5: A Monte Carlo generator for deep inelastic lepton - nucleon scattering, *Comput. Phys. Commun.* **101** (1997), 108, [arXiv:hep-ph/9605286](https://arxiv.org/abs/hep-ph/9605286).
- [18] T. Sjöstrand, High-energy physics event generation with PYTHIA 5.7 and JETSET 7.4, *Comput. Phys. Commun.* **82** (1994), 74.
- [19] B. Nilsson-Almqvist and E. Stenlund, Interactions Between Hadrons and Nuclei: The Lund Monte Carlo, Fritiof Version 1.6, *Comput. Phys. Commun.* **43** (1987), 387.