

WLCG/HSF Workshop 2025

Monday 5 May 2025 - Friday 9 May 2025

IJCLab, Paris



Book of Abstracts

Contents

Test	1
Marionette: Data Structure Description and Management for Heterogeneous Computing	1
GWOSC and gravitational-wave data analysis training	1
Julia: Sustainability and Efficiency	2
Understanding the scale of HL-LHC physics analyses	2
Power Accounting in Heterogeneous Compute Clusters	3
ROOT: Taking stock of the Run 3 experience towards ROOT7	3
Introduction to the Virtual Research Environment: an end-user perspective	4
The European Virtual Institute for Research Software Excellence (EVERSE)	4
The EVERSE/ESCAPE use cases: evaluating and improving the quality of HEP software .	5
The EVERSE Research Software Quality Toolkit	5
The EVERSE training and recognition plan	6
Green software in HEP: benchmarks and studies on MC generators	7
Latest Developments in RooFit and Plans	7
Analysis at the HL-LHC: Data Delivery, ServiceX, and Addressing Our Analysis Challenges	8
Environmental impact, carbon and sustainability of computing in the ATLAS experiment	8
AdePT - Offloading electromagnetic showers in Geant4 simulations to GPU	9
Accelerating HEP detector simulations using G4HepEm	9
Scikit-HEP project news and future directions	10
Welcome in Orsay	10
Introduction - WLCG	10
Introduction - HSF	11
Introduction to the OPS session	11

AUDITOR. First results of the assessment by the WLCG sites	11
Accounting of jobs submitted with tokens and changes required in the WLCG accounting infrastructure.	11
GGUS: past, present, future	11
Job Allocation and Handling, status report from the WG	12
ATLAS analysis facility contribution	12
CMS analysis facility contribution	12
LHCb analysis facility contribution	12
Exploiting extremely-parallel technologies for future large-scale physics analyses demands	12
Discussion	13
Introduction and Status of DOMA	13
Rucio in DOMA	13
FTS evolution towards DC27	13
Status of DOMA BDT activity	14
Status of Monitoring of DOMA components	14
Mini data challenges	14
AI & WLCG/HSF - scope of the session	14
AI and Facilities	14
AI/ML for Physics use cases	14
Heterogeneous resources and experiment software	14
Facility view: Heterogeneous resources deployment and usage	15
Discussion	15
ALICE analysis facility contribution	15
Proposal for Interdisciplinary Analysis Facilities	15
User experience	15
Summary of the HEPiX Spring 2025 Workshop	16
Status and evolution of French sites	16
Storage relocation: INFN-CNAF experience	16
Responding Faster: Lessons from a security incident	16

WLCG Technical Roadmap - introduction	16
Technical Roadmap - setting the scene - estimates of requirements for Run4 and Run5	16
Technical Roadmap - setting the scene - estimates on capacity increase at flat budget scenario	17
Discussions on the structure and the content of the chapters	17
Discussions on the structure and content of the chapters	17
Computing Challenges for the Einstein Telescope	17
Recap and next steps - DOMA and WLCG Ops	17
Recap and next steps - WLCG/HSF Sustainability	17
Recap and next steps - HSF	18
Recap and next steps - WLCG Technical Coordination	18
Recap and next steps - Analysis at scale	18
Introduction	18
Discussion on power accounting benchmarks part 1: sites perspective	18
Discussion on power accounting and beyond part 2: inputs from sites, users and experimental collaborations	18
Training on sustainable computing at DESY	19
Overview, goals	19
HOWTO Train: the HSF Training experience	19
Open Discussion	19
User Learning for ePIC	19
Enhancing Job Monitoring with Power Consumption Metrics	20
GreenDIGIT: updates from the WLCG Environmental Sustainability Workshop	20
Token Transition Overview	21
Tape REST API, Tokens & FTS: A Unified Proposal	21
Discussion	22
Energy saving measures during peak summer at DESY	22
Environmental impact of computing in LHCb	22
Security training opportunities	22
Social Dinner	22

1

Test

HSF / 2

Marionette: Data Structure Description and Management for Heterogeneous Computing

Author: Nuno Dos Santos Fernandes¹

¹ *Laboratory of Instrumentation and Experimental Particle Physics (PT)*

Corresponding Author: nuno.dos.santos.fernandes@cern.ch

Marionette is a header-only C++ library that was designed to allow the description of arbitrary data structures that can work across heterogeneous compute devices and on the host, providing complete interoperability and convenient interfaces with no impact on runtime performance. This is achieved by decoupling the description of the data to be held from the way in which data will be stored, which enables the generation of all memory allocations, transfers and deallocations at compile-time without requiring the end user to write any sort of boilerplate code, achieving the same performance as the equivalent hand-written alternatives. Furthermore, an expressive and intuitive object-oriented interface can be offered on both host and device(s), especially since arbitrary functions can be added to the interface of both the individual objects and the entire collection of them, without any runtime performance penalty as everything is resolved at compile-time. This user-extensible interface also means that the behaviour of pre-existing data structures can be replicated, allowing for immediate porting of pre-existing code with only minor adjustments to the data types. Furthermore, since the user can override and further specialize data transfers, gradual porting to the new data structures with minimal performance loss becomes possible as efficient ways to convert to and from pre-existing structures can be provided. With a focus on flexibility, customisability and extensibility without compromising expressivity, convenience or ease of use, Marionette is designed to offer a general solution for expressing data structures, catering to a wide variety of use cases and levels of expertise, with little to no runtime impact.

Requested talk length:

HSF / 6

GWOSC and gravitational-wave data analysis training

Authors: Massimiliano Razzano^{None}; Stefano Bagnasco¹

¹ *Istituto Nazionale di Fisica Nucleare, Torino*

Corresponding Authors: massimiliano.razzano@pi.infn.it, bagnasco@to.infn.it

Since their discovery in 2015, gravitational waves have become a hot topic in physics research. Gravitational-wave data produced by the LVK Collaboration, formed by the LIGO, Virgo and KAGRA collaborations, become fully public after a grace period; combined with the relative simplicity of the data themselves (one time series of the main signal channel per each interferometer, plus some simple data quality information), this created a large community of professional scientists, students and even citizen scientists analyzing them.

The Gravitational-Wave Open Science Centre (gwosc.org), besides maintaining a site hosting the public data and relevant software, developed a large amount of training and teaching materials such

as tutorials and documentation, and organises periodical GW Open Data Workshops, “crash courses” in GW data analysis. This contribution describes such materials and activities.

Requested talk length:

15

HSF / 7

Julia: Sustainability and Efficiency

Authors: Caterina Doglioni¹; Graeme A Stewart²; Pere Mato Vila²; Samuel Cadellin Skipsey^{None}

¹ *The University of Manchester (GB)*

² *CERN*

Corresponding Authors: caterina.doglioni@cern.ch, samuel.cadellin.skipsey@cern.ch, pere.mato@cern.ch, graeme.andrew.stewart@cern.ch

There a number of studies of the general energy efficiency of different programming languages, however relatively few look at HEP specific examples. Here we present examples comparing energy efficiency of different jet reconstruction codes in different languages: specifically C++, Julia and Python. We also study the evolution of efficiency over recent releases of Julia and Python.

We also discuss general aspects of sustainability of code and show how the Julia language and ecosystem helps developers to write and maintain modular, interoperable codes that reduce the code maintenance burden.

We show that Julia is an excellent language choice, combining outstanding energy efficiency and human productivity, helping sustainability in all the most meaningful senses.

Requested talk length:

30

Plenary / 8

Understanding the scale of HL-LHC physics analyses

Authors: Alexander Held¹; Oksana Shadura²

¹ *University of Wisconsin Madison (US)*

² *University of Nebraska Lincoln (US)*

Corresponding Authors: alexander.held@cern.ch, oksana.shadura@cern.ch

The existing roadmaps and computing model plans from ATLAS and CMS for the HL-LHC area are primarily focused on the centralized aspect of computing: those steps that lead up to sets of files made available to physicists for analysis. The general approaches, resources used, and software frameworks for the area of “end-user physics analysis”, which starts from those files, are much less clearly defined, understood, or prescribed.

In order to better understand these aspects, IRIS-HEP, jointly with ATLAS AMG, CMS CAT and HSF DAAA, is designing a survey to capture the computational requirements from physics analysis use cases. This contribution will show first results from the survey and discuss the broader context of the effort. Following this survey, we aim to identify a set of physics analyses and define benchmark scenarios to extrapolate the concrete computing requirements to the HL-LHC era. The described workflows will provide more clarity about the role of analysis facilities and the kinds of services they should make available. It will allow for quantitative evaluation of analysis models and be a first step towards identifying what to do about analysis use cases that do not fit into the space set out by the benchmark examples.

Requested talk length:

20

Plenary / 9

Power Accounting in Heterogeneous Compute Clusters

Author: Emanuele Simili^{None}

Co-authors: Albert Gyorgy Borbely¹; David Britton¹; Gordon Stewart ; Samuel Cadellin Skipsey

¹ *University of Glasgow (GB)*

Corresponding Authors: gordon.stewart@glasgow.ac.uk, pepposub@gmail.com, albert.borbely@cern.ch, david.britton@cern.ch, samuel.cadellin.skipsey@cern.ch

Energy efficiency is a critical concern for WLCG operations. We present a proof-of-concept for dynamic power accounting in our heterogeneous compute clusters at ScotGrid Glasgow. Our approach leverages real-time metrics from Prometheus to attribute energy consumption to individual Virtual Organizations (VOs) based on actual core usage. By integrating hardware-specific power efficiency data, derived from static measurements across different node generations, we compute per-core power usage while accounting for architectural differences.

Our methodology distinguishes between active power (consumed by running jobs) and infrastructure overhead (idle power and other services), the latter is allocated to the hosting institute. This granular, data-driven model not only provides transparent energy allocation but also encourages system administrators to optimize resource utilization and improve overall Power Usage Effectiveness (PUE).

Our work lays the foundation for integrating energy accounting into existing monitoring infrastructures and provides insights into sustainable cluster operations.

Requested talk length:

20

HSF / 10

ROOT: Taking stock of the Run 3 experience towards ROOT7

Authors: Danilo Piparo¹; Jonas Rembser¹

¹ *CERN*

Corresponding Author: danilo.piparo@cern.ch

ROOT is a unified software package for the storage, processing, and analysis of scientific data: from its acquisition to the final visualization in the form of highly customizable, publication-ready plots. Successfully used by experiments and thousands of physicists, the ROOT Project is preparing its seventh release cycle, sustained by intense R&D activities.

In this contribution, after briefly reviewing the status of the project, we'll focus on the results harvested along the R&D activities conducted so far and how those will shape the future of ROOT. The areas on which we'll concentrate will mainly be three. Firstly, we discuss how hardware accelerators can be exploited by users with ROOT and how these new features are relevant at present and future Analysis Facilities. Then, we review the current development status of RNTuple and its adoption by experiments, illustrating both advancements in terms of performance but also usability, in C++ and Python. Finally, we'll concentrate on the Python-C++ integration, at a low level, and the recent advancements in the interoperability of ROOT and Scientific Python, at the user-facing level.

Requested talk length:

20

HSF / 11

Introduction to the Virtual Research Environment: an end-user perspective

Author: Sukanya Sinha¹

¹ *The University of Manchester (GB)*

Corresponding Author: sukanya.sinha@cern.ch

One of the objectives of the EOSC (European Open Science Cloud) Future Project was to integrate diverse analysis workflows from Cosmology, Astrophysics and High Energy Physics in a common framework. This led to the inception of the Virtual Research Environment (VRE) at CERN, a prototype platform supporting the goals of Dark Matter and Extreme Universe Science Projects in compliance with FAIR (Findable, Accessible, Interoperable, Reusable) data policies. The goal of the project was to highlight the synergies between different dark matter communities and experiments, by producing new scientific results as well as by making the necessary data and software tools fully available. The VRE makes use of a common authentication and authorisation infrastructure (AAI), and shares the different experimental data (ATLAS, Fermi-LAT, CTA, Darkside, Km3Net, Virgo, LOFAR) in a reliable distributed storage infrastructure via the ESCAPE Data Lake. The entry point of such a platform (for an end-user) is a jupyterhub instance deployed on top of a scalable Kubernetes infrastructure, providing an interactive graphical interface for researchers to access, analyse and share data. The data access and browsability is enabled through API calls to the high level data management and storage orchestration software (Rucio). The VRE aims to streamline the development of end-to-end physics workflows, granting researchers access to an infrastructure that contains easy-to-use physics analysis workflow from different experiments. In this contribution, I will provide an overview of the VRE, highlight its use cases as an analyser for implementing and reproducing experimental analyses on a REANA cluster, and showcase the successful integration of an ATLAS experimental analysis workflow into the VRE platform.

Requested talk length:

HSF / 12

The European Virtual Institute for Research Software Excellence (EVERSE)

Authors: Caterina Doglioni¹; Graeme A Stewart²; Michael Sparks³; Stefan Roiser²

¹ *The University of Manchester (GB)*

² *CERN*

³ *BBC*

Corresponding Authors: msparks.work@gmail.com, caterina.doglioni@cern.ch, stefan.roiser@cern.ch, graeme.andrew.stewart@cern.ch

The EVERSE EU-funded project aims to create a framework for research software and code excellence, collaboratively designed and championed by research communities that include physics and astronomy.

EVERSE's ultimate ambition is to contribute towards a cultural change where research software is recognized as a first-class citizen of the scientific process and the people that contribute to it are credited for their efforts.

In this contribution, we will outline the aims of the network and present the achievements of the first year of the project, towards building a European network of Research Software Quality and setting the foundations of a future Virtual Institute for Research Software Excellence, including the creation of the EVERSE Network. The goal of this contribution is also to interactively discuss how these best practices can map to the HSF's sponsored project requirements.

Requested talk length:

20

HSF / 13

The EVERSE/ESCAPE use cases: evaluating and improving the quality of HEP software

Authors: Caterina Doglioni¹; James Smith¹; Michael Philip Sparks¹; Tobias Fitschen¹

¹ *The University of Manchester (GB)*

Corresponding Authors: michael.sparks@manchester.ac.uk, caterina.doglioni@cern.ch, jl.smith@cern.ch, tobias.fitschen@cern.ch

This contribution will present how the EVERSE project interfaces with the European Open Science Clusters (ENVRI-FAIR for environmental sciences, Life Sciences RI, ESCAPE for Particle physics and astrophysics, PaNOSC for Photon and neutron science and SSHOC for social sciences and humanities) through use cases of software packages or infrastructures that are in current used by researchers.

It is through these use cases that EVERSE draws best practices from the different communities, and where the various elements of the software excellence framework are tested and implemented prior to their release to the wider community. In this talk, we will describe the three ESCAPE use cases, highlight the elements that these use cases have contributed to EVERSE so far, and outline the expected improvements and pathway to obtain them by the end of the project in 2027.

Requested talk length:

10

HSF / 14

The EVERSE Research Software Quality Toolkit

Authors: Caterina Doglioni¹; Graeme A Stewart²; Michael Philip Sparks¹; Stefan Roiser²

¹ *The University of Manchester (GB)*

² *CERN*

Corresponding Authors: michael.sparks@manchester.ac.uk, caterina.doglioni@cern.ch, graeme.andrew.stewart@cern.ch, stefan.roiser@cern.ch

The Research Software Quality Toolkit (RSQKit - <https://everse.software/RSQKit/>), developed by the EVERSE project, lists curated best practices in improving the quality of your research software. It is intended for use by researchers, research software engineers, as well as those running research infrastructures involving software or involved in research software-related policy and funding.

These practices are informed by software excellence and quality in the context of research; with a focus on FAIR software, Open Research, community development and software engineering practices at different tiers of research software (analysis scripts, prototype tools and research software infrastructure). RSQKit links to tools and resources which support best practices. It includes software quality dimensions and links to indicators and tasks to guide the usage of the best practice, as well as links to training resources and existing guides and materials.

This contribution will introduce RSQKit and its aims and architecture. It is presented in a way to attract and consider feedback from researchers who code from the particle physics community and from WLCG infrastructure experts.

Requested talk length:

15

HSF / 15

The EVERSE training and recognition plan

Authors: Caterina Doglioni¹; Graeme A Stewart²; Kenneth Brian Rioja³; Stefan Roiser²

¹ *The University of Manchester (GB)*

² *CERN*

³ *IT-FTI*

Corresponding Authors: graeme.andrew.stewart@cern.ch, kenneth.brian.rioja@cern.ch, caterina.doglioni@cern.ch, stefan.roiser@cern.ch

The EVERSE project aims to collect, enhance and curate training resources aligned with domain-specific practices, create a long-term training activity supported by community services and platforms and establish a framework for recognizing Trainers and RSEs.

This contribution will describe how EVERSE plans to collect and provide training, guidance and education to researchers, software developers, and other stakeholders in the research community to help them understand the importance of software and code quality, as well as how to apply established best practices and standards for assessing, verifying, and improving the quality of their software and code.

We will describe the registry of training initiatives compiled so far, and how this registry will be maintained and presented through the TESS infrastructure. We will also outline recent activities in terms of recognition of software activities and roles.

Our aim is to get feedback from the community at the workshop and to highlight missing aspects in training and recognition that we can bring back to the project.

Requested talk length:

15

Plenary / 16**Green software in HEP: benchmarks and studies on MC generators****Authors:** Luis Villar¹; Tobias Fitschen²**Co-authors:** Caterina Doglioni ²; James Smith ²; Michael Sparks ²¹ *University of Manchester*² *The University of Manchester (GB)***Corresponding Authors:** msparks.work@gmail.com, caterina.doglioni@cern.ch, tobias.fitschen@cern.ch, jl.smith@cern.ch, luis.villar-2@student.manchester.ac.uk

In this talk, we will describe the studies undertaken at the University of Manchester to estimate and improve the energy efficiency of computing hardware and software used by students and researchers.

The goal of these studies is to build an understanding of the environmental impact of particle physics research focusing on two fronts:

- 1) the carbon cost of the hardware uses for high power computing hardware and the local computing cluster
- 2) the energy efficiency of data analysis software and machine learning models in “big data”-related scientific fields including as high-energy particle physics.

The focus of this contribution will be the energy efficiency of scientific software algorithms and MC generation packages, taking Herwig, ML data compression and top tagging algorithms as examples. We will discuss different tools and benchmarks and review their methodologies.

We will then describe our plans towards a lifecycle analysis for computing hardware, and work undergoing to estimate the power consumption of our local cluster more precisely.

Requested talk length:

15

HSF / 17**Latest Developments in RooFit and Plans****Author:** Jonas Rembser¹¹ *CERN***Corresponding Author:** jonas.rembser@cern.ch

RooFit is a software package written in C++ for statistical data analysis that is part of ROOT. It is widely used in the High Energy Physics (HEP) community, with the most prominent users being the LHC collaborations. Recent RooFit development has focused on performance improvements and supporting new statistical analysis approaches to enable cutting-edge analyses, such as combined Higgs measurements with ATLAS or CMS. In this contribution, the development pillars that helped

to achieve this goal are elaborated on. The first pillar is code optimization and refactoring to optimally use both CPU and GPU resources. Then, there is supporting Automatic Differentiation (AD) with Clad, a compiler plugin for Clang. Furthermore, RooFit now provides new Python interfaces to include ML models as likelihood surrogates, enabling Simulation-Based Inference (SBI). Finally, this contribution will also report on the development status of the Minuit2 library for numerical minimization since it is a key dependency of RooFit, and the two packages are developed hand-in-hand to implement performance-optimal statistical analysis workflows for HEP.

Requested talk length:

30

Plenary / 19

Analysis at the HL-LHC: Data Delivery, ServiceX, and Addressing Our Analysis Challenges

Authors: Artur Cordeiro Oudot Choi¹; Benjamin Galewsky²; Gordon Watts¹; Ilija Vukotic³; Kyungeon Choi⁴; Peter Onyisi⁴; Roger Janusiak⁵

¹ *University of Washington (US)*

² *Univ. Illinois at Urbana Champaign (US)*

³ *University of Chicago (US)*

⁴ *University of Texas at Austin (US)*

⁵ *University of Washington*

Corresponding Authors: gwatts@uw.edu, ponyisi@utexas.edu, ilija.vukotic@cern.ch, rogerja@uw.edu, bengal1@illinois.edu, artur.cordeiro.oudot.choi@cern.ch, kyungeon.choi@cern.ch

As the HL-LHC era approaches, the scale and complexity of data present challenges for analysis workflows within ATLAS and other HL-LHC experiments. This contribution reports on recent developments in ServiceX, a cross experiment utility, and its role as a data delivery and transformation service within the analysis ecosystem. Designed to bridge the gap between centrally produced datasets and user-level analysis code, ServiceX now supports a broader range of input formats—including custom non-ROOT ATLAS data—and has seen significant improvements in reliability, scalability, and ease of use.

We highlight progress in several areas: improved bullet-proofing of the data transformation infrastructure to avoid errors at scale; new convenience utilities for introspecting datasets (e.g., listing available branches); and tighter integration with common ATLAS frameworks such as TopCPTools, enabling more standardized analysis environments. We also share results from recent scaling tests demonstrating performance at analysis-scale workloads, and discuss how ServiceX fits into the broader analysis software stack, helping to address long-standing challenges of reproducibility, data access, and user efficiency.

This talk will provide a view of the design of ServiceX, the technical and usability gains made in recent months, and what's coming next as we prepare for the HL-LHC's analysis demands.

Requested talk length:

30

Plenary / 20

Environmental impact, carbon and sustainability of computing in the ATLAS experiment

Author: Rodney Walker¹

Co-author: Zach Marshall²

¹ *Ludwig Maximilians Universitat (DE)*

² *Lawrence Berkeley National Lab. (US)*

Corresponding Authors: rodney.walker@physik.uni-muenchen.de, zach.marshall@cern.ch

In this contribution, we will describe the efforts within the ATLAS experiment to evaluate and mitigate various aspects of the environmental impact of ATLAS computing sites, such as building awareness in the experiment community, adjusting aspects of the computing policy, and modifications of data center configurations, either in ways that take advantage of particular features of ATLAS work or in generic ways that reduce the environmental impact of the computing.

Requested talk length:

10

HSF / 21

AdePT - Offloading electromagnetic showers in Geant4 simulations to GPU

Authors: Andrei Gheata¹; Benjamin Morgan²; John Apostolakis¹; Jonas Hahnfeld³; Juan Gonzalez Caminero¹; Mihaly Novak¹; Severin Diederichs¹; Stephan Hageboeck¹; Witold Pokorski¹

¹ *CERN*

² *University of Warwick*

³ *CERN & Goethe University Frankfurt*

Corresponding Authors: andrei.gheata@cern.ch, mihaly.novak@cern.ch, j.gonzalez@cern.ch, jonas.hahnfeld@cern.ch, benjamin.morgan@cern.ch, john.apostolakis@cern.ch, witold.pokorski@cern.ch, severin.diederichs@cern.ch, stephan.hageboeck@cern.ch

The Geant4 simulation throughput of LHC experiments is limited by increasing detector complexity in the high-luminosity phase. As high-performance computing shifts toward heterogeneous architectures such as GPUs, GPU-accelerated particle transport simulations offer a potential way to improve performance. Currently, only electromagnetic showers can be offloaded to GPUs, making an efficient CPU-GPU workflow essential. In this contribution, we present state-of-the-art detector simulations for LHC experiments using GPUs, outline the outstanding challenges, and discuss future directions.

Requested talk length:

20

HSF / 22

Accelerating HEP detector simulations using G4HepEm

Author: Mihaly Novak¹

Co-authors: Benjamin Morgan ²; Jonas Hahnfeld ³

¹ *CERN*

² *University of Warwick (GB)*

³ *CERN & Goethe University Frankfurt*

Corresponding Authors: mihaly.novak@cern.ch, ben.morgan@warwick.ac.uk, jonas.hahnfeld@cern.ch

Geant4 based detector simulations make a significant contribution to the overall computing budget of the LHC experiments. The individual experiments have been investing considerable effort in making their simulations more and more efficient. These performance optimisations are now even more important in order to cope with the special computing challenges of the HL-LHC era.

G4HepEm is one of the R&D projects that have been launched with the goal of contributing to this effort. It provides an efficient simulation of the electromagnetic shower, tailored for HEP detector simulations, in the form of a Geant4 extension. A significant performance improvement (~20 %) of the ATLAS and CMS full detector simulations has been achieved recently after integrating G4HepEm into the ATLAS Athena and CMS-SW frameworks while preserving the accuracy of the results. The motivations, ideas and results obtained for ATLAS and CMS will be presented.

Requested talk length:

HSF / 23

Scikit-HEP project news and future directions

Author: Eduardo Rodrigues¹

¹ *University of Liverpool (GB)*

Corresponding Author: eduardo.rodrigues@cern.ch

Scikit-HEP is a community-driven and community-oriented project with the goal of providing an ecosystem for particle physics data analysis in Python fully integrated with the wider scientific Python ecosystem. The project provides many packages and a few “affiliated” packages for data analysis. It expands the typical Python data analysis tools for particle physicists, with packages spanning the spectrum from general scientific libraries for data manipulation to domain-specific libraries. An overview of where the project is will be presented. Future developments and matters of sustainability will be discussed.

Requested talk length:

Plenary / 24

Welcome in Orsay

Corresponding Authors: chambert@ipno.in2p3.fr, michel.jouvin@ijclab.in2p3.fr

Requested talk length:

Plenary / 25

Introduction - WLCG

Corresponding Authors: tboccali@cern.ch, david.britton@cern.ch, simone.campana@cern.ch

Plenary / 26

Introduction - HSF

Corresponding Authors: paul.james.laycock@cern.ch, eduardo.rodrigues@cern.ch

WLCG / 27

Introduction to the OPS session

Corresponding Author: julia.andreeva@cern.ch

WLCG / 28

AUDITOR. First results of the assessment by the WLCG sites

Corresponding Authors: thomas.hartmann@cern.ch, alessandro.pascolini@cern.ch, kleinemuehl@uni-wuppertal.de, max.fischer@kit.edu, michael.boehler@desy.de, michael.boehler@cern.ch, maria.alandes.pradillo@cern.ch

AUDITOR (Accounting Data Handling Toolbox for Opportunistic Resources) is a flexible and extensible accounting system designed to support a wide range of use cases and infrastructures. Its integration with APEL enables it to function as a generic component within the WLCG accounting infrastructure, tracking the usage of various types of site computing resources. Several WLCG sites have evaluated AUDITOR, and the outcomes of these assessments will be reviewed and discussed.

Requested talk length:

WLCG / 29

Accounting of jobs submitted with tokens and changes required in the WLCG accounting infrastructure.

Corresponding Author: tom.dack@cern.ch

Requested talk length:

WLCG / 30

GGUS: past, present, future

Corresponding Authors: aliaksei.hrynevich@cern.ch, pavel.weber@cern.ch, guenter.grein@kit.edu, aliaksei.hrynevich@kit.edu

Requested talk length:

WLCG / 31

Job Allocation and Handling, status report from the WG

Corresponding Author: antonio.perez.calero.yzquierdo@cern.ch

Plenary / 32

ATLAS analysis facility contribution

Corresponding Author: emma.torro.pastor@cern.ch

Requested talk length:

Plenary / 33

CMS analysis facility contribution

Corresponding Author: oksana.shadura@cern.ch

Plenary / 34

LHCb analysis facility contribution

Corresponding Author: nicola.skidmore@cern.ch

Plenary / 35

Exploiting extremely-parallel technologies for future large-scale physics analyses demands

Corresponding Author: jonas.rembser@cern.ch

The HEP experiments and community face complex computing and storage requirements which are expected to increase by several factors with the advent of HL-LHC. By the end of Run3, the machine will have accumulated roughly 10% of the total dataset. These data, stored in the ROOT data format, are then analysed by the various experiment communities, with varying levels of coordination and centralisation.

Very often, the computational complexity of these physics analyses is such that parallelisation must be employed to ensure a reasonable time-to-insight for the physicists. Batch computing on the grid is a well-oiled workflow that has enabled obtaining a plethora of important physics results. In recent times, interactive scheduling approaches have emerged that lower the programming entry barrier and allow for a more ergonomic end-user experience. In this context, a specific combination of storage, computing resources, software distribution and distributed engines can be identified and is sometimes referred to as an “analysis facility”.

In this contribution, we highlight existing technologies such as ROOT, SWAN, EOS, XRootD and others that represent concrete building blocks available in production today to help the community in addressing the expected complexity of HL-LHC analyses. We also display how various current and future physics use cases can benefit from such technologies. These are used both in analyses facilities and other distributed computing infrastructures to minimise the time-to-insight and enable ergonomic and user-friendly distributed computing.

Requested talk length:

Plenary / 36

Discussion

Corresponding Authors: nicola.skidmore@cern.ch, alessandra.forti@cern.ch

Requested talk length:

WLCG / 37

Introduction and Status of DOMA

Corresponding Authors: johannes.elmsheuser@cern.ch, katy.ellis@stfc.ac.uk

Requested talk length:

WLCG / 38

Rucio in DOMA

Corresponding Author: martin.barisits@cern.ch

WLCG / 39

FTS evolution towards DC27

Corresponding Authors: steven.murray@cern.ch, mihai.patrascoiu@cern.ch

Requested talk length:

WLCG / 40

Status of DOMA BDT activity

Corresponding Authors: petr.vokac@cern.ch, mihai.patrascoiu@cern.ch

WLCG / 41

Status of Monitoring of DOMA components

Corresponding Author: borja.garrido.bear@cern.ch

Requested talk length:

WLCG / 42

Mini data challenges

Corresponding Authors: katy.ellis@stfc.ac.uk, johannes.elmsheuser@cern.ch

Plenary / 43

AI & WLCG/HSF - scope of the session

Corresponding Authors: paul.james.laycock@cern.ch, alessandro.di.girolamo@cern.ch

Plenary / 44

AI and Facilities

Corresponding Author: ilija.vukotic@cern.ch

Plenary / 45

AI/ML for Physics use cases

Corresponding Author: tboccali@cern.ch

Plenary / 46

Heterogeneous resources and experiment software

Corresponding Author: pcalafiura@lbl.gov

Plenary / 47

Facility view: Heterogeneous resources deployment and usage

Corresponding Authors: ben.dylan.jones@cern.ch, rodney.walker@physik.uni-muenchen.de

Requested talk length:

Plenary / 48

Discussion

Plenary / 49

ALICE analysis facility contribution

Corresponding Author: maarten.litmaath@cern.ch

Plenary / 50

Proposal for Interdisciplinary Analysis Facilities

Corresponding Author: thomas.hartmann@cern.ch

Next to HEP other scientific communities have emerged producing large amounts of data and requiring correspondingly processing power. While Analysis Facilities can be optimized to specific workflows, the underlying risk is a highly specialised and inflexible lock-in. By aiming for an interdisciplinary approach concentrating on established industry standards and establishing generic interfaces, Analysis Facilities can profit from a broader usage scope.

Requested talk length:

Plenary / 51

User experience

Corresponding Author: albert.borbely@cern.ch

A users perspective on typical problems faced throughout an analysis cycle, the ad-hoc solutions implemented, and the limited view a user has of the Grid and how it is displayed to them.

Requested talk length:

WLCG / 52

Summary of the HEPiX Spring 2025 Workshop

Corresponding Author: jose.flix.molina@cern.ch

WLCG / 53

Status and evolution of French sites

Corresponding Author: laurent.duflot@in2p3.fr

WLCG / 54

Storage relocation: INFN-CNAF experience

Corresponding Author: andrea.rendina@cnafl.infn.it

Requested talk length:

Plenary / 55

Responding Faster: Lessons from a security incident

Corresponding Author: jose.carlos.luna@cern.ch

NOTE: this contribution was actually part of the *Facilities* session, which overflowed into the afternoon.

Requested talk length:

WLCG / 56

WLCG Technical Roadmap - introduction

Corresponding Authors: letts@physics.ucsd.edu, alessandro.di.girolamo@cern.ch

WLCG / 57

Technical Roadmap - setting the scene - estimates of requirements for Run4 and Run5

Corresponding Author: ben.couturier@cern.ch

WLCG / 58

Technical Roadmap - setting the scene - estimates on capacity increase at flat budget scenario

Corresponding Author: andrea.sciaba@cern.ch

WLCG / 59

Discussions on the structure and the content of the chapters

Corresponding Authors: letts@physics.ucsd.edu, alessandro.di.girolamo@cern.ch

WLCG / 60

Discussions on the structure and content of the chapters

Corresponding Authors: letts@physics.ucsd.edu, alessandro.di.girolamo@cern.ch

Plenary / 61

Computing Challenges for the Einstein Telescope

Corresponding Author: paul.james.laycock@cern.ch

Plenary / 62

Recap and next steps - DOMA and WLCG Ops

Corresponding Authors: katy.ellis@stfc.ac.uk, panos.paparrigopoulos@cern.ch

Requested talk length:

Plenary / 63

Recap and next steps - WLCG/HSF Sustainability

Corresponding Authors: david.britton@cern.ch, caterina.dogliani@cern.ch

Requested talk length:

Plenary / 64

Recap and next steps - HSF

Corresponding Author: graeme.andrew.stewart@cern.ch

Requested talk length:

Plenary / 65

Recap and next steps - WLCG Technical Coordination

Plenary / 66

Recap and next steps - Analysis at scale

Corresponding Authors: nicola.skidmore@cern.ch, alessandra.forti@cern.ch

Plenary / 67

Introduction

Corresponding Authors: david.britton@cern.ch, caterina.dogliani@cern.ch

Talk to set the scene and summarise the actions from the WLCG environmental sustainability workshop in December 2024: <https://indico.cern.ch/event/1450885/timetable/>

Plenary / 68

Discussion on power accounting benchmarks part 1: sites perspective

Plenary / 69

Discussion on power accounting and beyond part 2: inputs from sites, users and experimental collaborations

HSF / 70

Training on sustainable computing at DESY

Corresponding Author: christoph.wissing@desy.de

A few years ago, primarily young scientists of the DESY particle physics division founded a forum for sustainability to enhance awareness of the environmental impact of research and to propose measures to reduce the energy footprint. Since computing is a substantial consumer of resources in particle physics, the forum initiated a series of workshops on computing with the aim of training scientists in best practices for efficient computing and software. Roughly twice per year, there is a workshop for beginners and newcomers to DESY with an introduction to the local computing facilities and hands-on exercises employing tools like ROOT, Git, and HTCondor. About once per year, there is a workshop on a slightly more advanced topic, such as software testing and performance, or using and building containers. The workshops are organized by a team of scientists from the IT department and from the experimental groups. All exercises are constructed along examples that are typical for particle physics and employ, e.g., CMS Open Data.

Requested talk length:

HSF / 71

Overview, goals

Corresponding Author: mhernande1@bnl.gov

HSF / 72

HOWTO Train: the HSF Training experience

Corresponding Author: mhernande1@bnl.gov

The HSF Training group has built a fruitful learning environment within the high-energy and nuclear physics community through the organization of numerous training events. This talk will share practical insights gained from years of experience in planning and executing training events. We have learnt that organizing effective training requires careful planning, continuous adaptation, and the crucial role of community feedback. We will highlight strategies for engaging participants, adapting to diverse skill levels, and building a sustainable training ecosystem that upskill the next generation of researchers and software developers.

Requested talk length:

HSF / 73

Open Discussion

HSF / 74

User Learning for ePIC

Corresponding Author: hszumila@jlab.org

The ePIC Software User Learning group supports the ePIC Collaboration by providing resources and opportunities for new user onboarding, software training events, and maintaining documentation. The ePIC Software employs a User-Centered Design model and supports the detector simulation and geometry design, analysis of physics and event reconstruction, and provides benchmarks to ensure a cohesive and consistent framework is available to users. This talk will discuss the strategies and resources available so far for developers with plans for future growth and support.

Requested talk length:

Plenary / 75

Enhancing Job Monitoring with Power Consumption Metrics

Corresponding Authors: natalia.diana.szczepanek@cern.ch, domenico.giordano@cern.ch

To support the sustainability of WLCG compute infrastructure, we propose a strategy to extend the current job monitoring system to include energy consumption data. Currently, WLCG monitoring systems primarily focus on traditional job metrics such as CPU time, memory usage, runtime, and failure rates.

However, they do not capture job-level power consumption, as this data is typically managed by fabric-level monitoring systems and is not directly accessible within jobs.

Our proposal is to bridge this gap by including node-level power measurements within the job data reports.

We have successfully demonstrated the feasibility of this idea with a prototype approach that uses standard WLCG grid jobs submitted to a set of pilot sites (DESY, AGLT, Glasgow).

Although this approach is valid for any kind of job payload, we have leveraged the HEPBenchmark Suite as payload, which organically captures both the performance of the node in terms of HS23 and a set of utilization metrics, including CPU load, memory usage, frequency, and power consumption. This setup enables correlation between performance and energy efficiency metrics. It provides a more comprehensive view of job efficiency, revealing opportunities for optimization.

This is just the start of an R&D process that would need the involvement of the whole WLCG community to implement that strategy.

Requested talk length:

Plenary / 76

GreenDIGIT: updates from the WLCG Environmental Sustainability Workshop

Corresponding Author: gergely.sipos@egi.eu

GreenDIGIT (<https://greendigit-project.eu/>) is a TECH-01-01 Horizon Europe project that started in March 2024 to pursue environmental sustainability within digital services and service ecosystems that research infrastructures (RIs) rely on. GreenDIGIT brings together institutes from 4 digital

RIs: EGI, SLICES, EBRAINS, and SoBigData to address the pressing need for sustainable practices in digital service provisioning.

Central to GreenDIGIT's mission is the identification of good practices in emphasized textow-impact computing, which involves evaluating the broader computing landscape and identifying opportunities for improvements across various stakeholder groups both within, and beyond the RIs represented in the project consortium. The project will establish reference architectures and design principles that facilitate environmental impact considerations throughout the entire RI lifecycle, ensuring that sustainability is embedded within the entire design-implementation-operation-termination phases of RIs at the level of nodes, services and components.

Moreover, GreenDIGIT will develop and validate innovative technologies and methodologies that empower digital service providers to reduce energy consumption and/or overall operational environmental impact. By supplying technical tools for both providers and researchers, the project promotes the design and execution of environmentally conscious digital applications, aligning with principles of Open Science and FAIR data management.

Through education and support initiatives GreenDIGIT will foster a culture of sustainability among researchers and service providers, equipping them with best practices for lifecycle management and operation. By mapping the landscape of environmentally friendly computing, and by offering policy and technical recommendations, the project not only enhances the sustainability of research infrastructures but also contributes to the broader goal of mitigating climate change.

The first tangible outcome of the project is a landscape analysis that was conducted within the 4 participating RIs, and within other RIs that operate significant digital infrastructures. The presentation will share the first findings from this landscape report, which included responses from 15 EGI sites, many being present in WLCG too.

In the second phase of the project EGI members will extend the EGI-WLCG service configuration database (GOCDB) with environmental metrics to fuel "green workload optimisation strategies" in HTC and Cloud infrastructures. Based on the metrics the DIRAC workload manager (CNRS), the AI4OS AI/ML framework (CSIC), the Terraform DevOps tool (SZTAKI) and a Data Center Energy scheduler (CESNET) will be extended to optimise job/VM/task execution according to green strategies.

In summary, the GreenDIGIT project is a transformative initiative that champions environmental sustainability in digital research services, offering various outcomes and collaboration opportunities to the WLCG community.

This contribution will present updates since the last presentation at the WLCG Environmental Sustainability Workshop in December, focusing on (a) Framework for managing the full digital RI lifecycle from the environmental perspective (b) Software tools for power usage efficiency

Requested talk length:

WLCG / 77

Token Transition Overview

Corresponding Author: maarten.litmaath@cern.ch

WLCG / 78

Tape REST API, Tokens & FTS: A Unified Proposal

Corresponding Author: mihai.patrascoiu@cern.ch

WLCG / 79

Discussion

Plenary / 80

Energy saving measures during peak summer at DESY

Corresponding Author: thomas.hartmann@cern.ch

This contribution will detail initial plans for reducing peak day loads at the DESY computing center during the summer.

Requested talk length:

Plenary / 81

Environmental impact of computing in LHCb

Corresponding Authors: francesco.sborzacchi@cern.ch, hgcms@cern.ch, henryk.giemza@cern.ch

This contribution will present the activities in LHCb online and offline towards environmentally sustainable computing.

WLCG / 82

Security training opportunities

Corresponding Author: david.crooks@stfc.ac.uk

Requested talk length:

83

Social Dinner

in La Mère Catherine

Plenary / 84

Some final words ...