

Updates on C&SW in CERN/SFT, community from WLCG/HSF liaison role

Eduardo Rodrigues
University of Liverpool

Note:

- ❑ Largely an executive summary of the last 2 reports on CERN SFT and HSF/community activities to the WLCG Management Board
- ❑ Reports prepared by Predrag Buncic (for SFT) and myself (for HSF/community) as software liaisons

- ❑ Will end with some news on near-future activities ...

CERN SFT group – news and new structure

❑ New Group Leader (G. Ganis) and deputy (P. Buncic) since 01/07/2023

❑ New projects view:

Changes in group organization

❑ Change of ROOT PL

- A. Naumann took WP1 coordination of the NGT (Next Generation Triggers) project
- New PL: D. Piparo

❑ Simulation: flattened structure (new PL: A. Ribon)

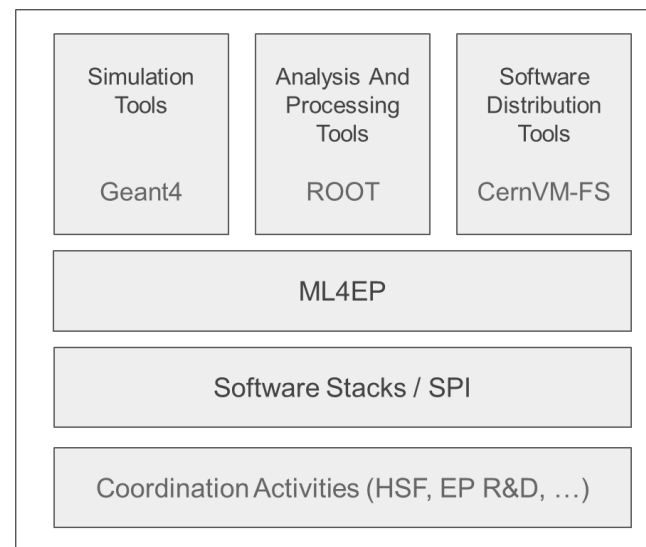
❑ Starting of SFTLabs: cradle for ideas, <20% member time

- Currently investigating possible role of Julia

❑ New Machine Learning Project

- Expectations
 - Attract interest from developing programs, e.g. NGT
 - Act as synergising point for solutions developed in the community
 - Build-up critical expertise around available technologies
- Will start from re-arranging current ML related activities from Simulation and ROOT
- Project Leader: Lorenzo Moneta

❑ Stacks (previously known as SPI), including Key4Hep (new PL: A. Sailer)



ALICE
ATLAS
CMS
LHCb
EP R&D
IT projects
HSF
FCC
...

CERN SFT Programme of Work (PoW)

- ❑ One day Programme of Work meeting held on 15/01/2024
 - ❑ <https://indico.cern.ch/event/1341656>
- ❑ SFT Project presented the planned Programme of Work for 2024
- ❑ Experiments representatives were invited and asked to provide a consolidated feedback

- ❑ The feedback was received at the Architects Forum meeting on 01/02/2024
 - ❑ <https://indico.cern.ch/event/1354953/>

 - ❑ Summaries pasted later for reference ...

LHC Experiments Feedback to SFT PoW 2024

James Catmore¹, John Chapman¹, Marco Clemencic², **Giulio Eulisse**³,
Peter Hristov³, James Letts⁴, Elizabeth Sexton-Kennedy⁴, Nicholas Styles¹

¹ Atlas, ² LHCb, ³ ALICE, ⁴ CMS

Vision for new ML/AI effort

Too early to provide concrete feedback.

Effort looks like going in the **correct direction**, do not reinvent the (Open Source) wheel, but provide **easy interfacing in ROOT for the already existing tools (and viceversa)**.

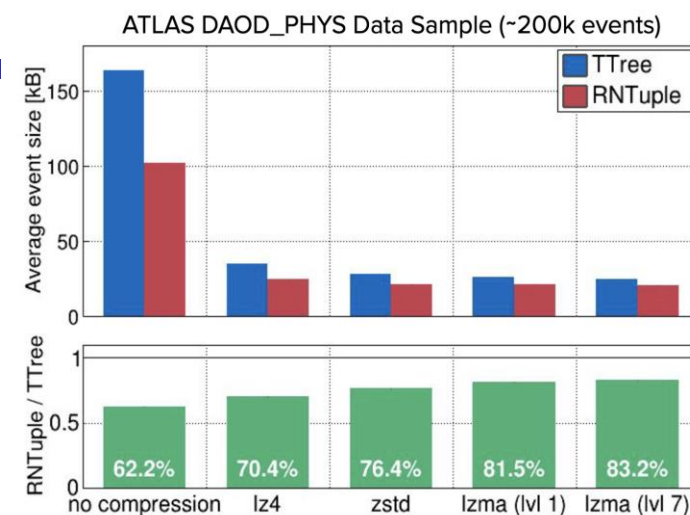
Ability to **avoid format conversion from ROOT files to do the training** would be a massive advantage (RNTuple reader library usable with major ML packages?)

Inference should be as transparent as possible. In particular:

- analysis steered from Python **and** (particularly) for production-level reconstruction/simulation
- **as few dependencies** as possible
- emphasis on performance and memory footprint


CERN SFT updates – Analysis / ROOT

- ❑ Support for C++20 since ROOT 6.28 (6.28/00 is LCG 103) , which has new features to support ML and improve RooFit
- ❑ Several other releases such as 6.30.04
 - Bug fixes and major performance improvements, targeting the C++ interpreter and Python bindings
- ❑ RNTuple (validation):
 - Technical work on large-scale testing effort started with teams from EP-SFT and IT-ST (12/2023)
 - First successful tests of the IRIS-HEP's Analysis Grand Challenge performed
 - After the AGC tests, continue with workloads from experiments
- ❑ RooFit acceleration with GPUs and automatic differentiation
- ❑ Development of the Python interoperability layer and histograms depends on EP R&D full
- ❑ Productive and successful [RNTuple workshop](#) with 4 LHC experiments + DUNE
 - The key features requested by experiments are already implemented
 - Discussed path to production involving the key software people
 - Large scale testing waiting for approval procedure to converge and result in resource allocation in IT



1st ROOT Hackathon: the Fixathon

- ❑ Took place at CERN on February 13-14
- ❑ Main objectives - all achieved:
 - Close as many ROOT issues as possible
 - Connect further with ROOT community, inclusively
 - ~30 people, from early career physicists to senior scientists
 - Onboard early career colleagues teaching how to contribute to a large open-source scientific software



ROOT

All the experiments would like to **thank Axel** for his work over the years and **welcome Danilo** as new project leader.

As a general remark **ROOT remains a central** component of our software stacks.

All experiments welcome **RNTuple** prioritization in the plan. A few critical features:

- Copy-less bulk reading (ALICE). **Performance is the feature** sometimes.
- Evolution path for **polymorphic types** (CMS, see also November plan at ROOT workshop)
- **Unaligned friends** (LHCb)

ROOT 6.32.0 release should be early (i.e. CMS: August, LHCb: September, ALICE: November) to be adopted for 2025 datataking.

Security of the Web GUI: it would be good to see explicitly mentioned a path to a more secure solution (e.g. the one Jupyter uses).

Everyone feels that **Windows support should be descoped** to “Community Effort”.

Everyone welcomes the rebasing on clang-repl. In general, **reducing customisation on top of CLANG / LLVM** is very welcome.

ALICE: **memory overhead** of cling should still be on the radar and it would appreciate progress on the **TBufferFile 1GB limit**.

CERN SFT updates – Stacks (was SPI)

- ❑ Keeping up with changes (CentOS -> AlmaLinux, MacOS, ARM) and demands from experiments
- ❑ “Turnkey Software”
 - Key4hep integrated software stack adopted for ECFA Higgs/Top/EW factory studies
 - Prototypes possible solution for FCC
- ❑ Experiments feedback of SFT’s 2024 Programme of Work:

Stacks

LCG Releases:

- **ATLAS** very happy with the prompt and efficient collaboration with SFT in the preparation of LCG releases.
- **LHCb** hopes for a lcgmake replacement.
- CMS and ALICE do not use them, nevertheless they are happy to participate in LIM.

Spack:

- **LHCb: looking forward to** testing Spack to replace lcgmake. If proven to be desirable in SFT, they will start using it ASAP.
- **ATLAS: happy to experiment** with it, but not baseline for now.
- **ALICE:** not interested until the missing “**developer story**” and “production HEP stack story” is there.
- **CMS: still valid feedback** given in detail in [March 2023](#) (“not the right tool for packaging large software stack”).

Key4HEP:

- **Not used as “turnkey solution”.** Some **components are** / will be used though (e.g. **DD4Hep, Gaudi**).

CERN SFT updates – CernVM(FS)

- ❑ Focus of work on scalability issues surfacing with increased usage
- ❑ Major release 2.11 addresses performance and operational issues
- ❑ Collaboration with [EESSI](#) and EuroHPC aimed at lowering the threshold for access to HPCs
 - Jointly gave an online tutorial on CVMFS best practices on HPCs: > 100 participants, very positive feedback



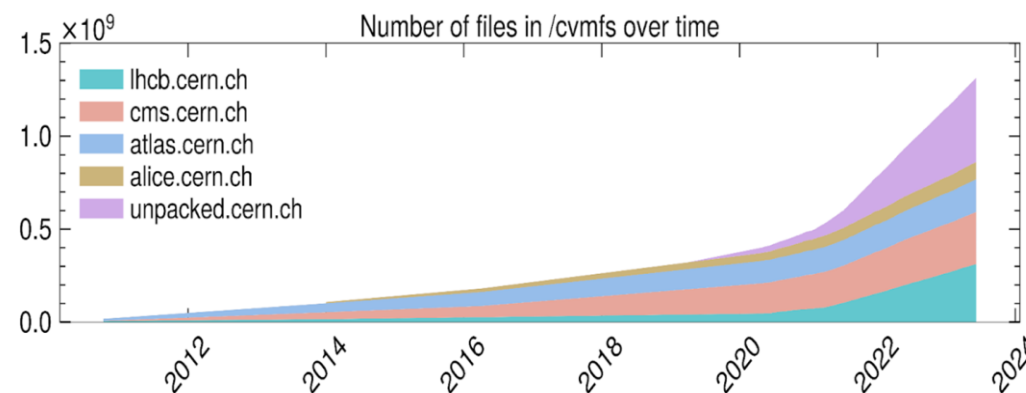
7. Would you recommend this tutorial to friends or colleagues?

● Absolutely not	0
● Probably not	0
● I'm not sure	1
● Maybe	4
● Definitely	48



- ❑ Experiments feedback of SFT's 2024 Programme of Work:
“In general happy with the performance and developments ...”

- ❑ New TECH student joined 02/2024,
bringing the team finally back to nominal size (3 FTEs)
 - Will focus on container tools





General news

- ❑ Many improvements on physics (models)
- ❑ Improved portability on VecGeom geometry and development of a new surface model geometry better suited for GPUs
- ❑ Exploring new approaches for fast simulation using ML in collaboration with ATLAS, LHCb, Openlab & IBM
- ❑ Extended the testing and validation suite – geant-val.cern.ch

Meetings and tutorials

- ❑ Held Geant4 Technical Fora at CERN, Collaboration Meeting in Japan
- ❑ Training courses at CERN: for Beginners (April), and Advanced (October)
- ❑ Other courses held in Italy, France, Japan, and South Korea

CERN SFT updates – simulation – GPU developments

- ❑ Two prototypes for offloading electromagnetic showers on GPUs
- ❑ See Ben Morgan's presentation from earlier today ... ;-)

- ❑ AdePT (CERN/Warwick)
 - ❑ Uses existing VecGeom for geometry and G4HepEm for electromagnetic physics
 - ❑ Develops a new surface-based model in VecGeom for GPUs (based on bounded surfaces)
 - ❑ Starting integration with experiments
- ❑ Celeritas (ORNL/FNAL/ANL/LBNL)
 - ❑ Uses VecGeom for geometry and own electromagnetic physics implementation on GPUs
 - ❑ Can use their own ORANGE as surface geometry model (based on unbounded surfaces)
 - ❑ Originally developed for nuclear reactor simulations
 - ❑ not (yet) applicable for large and complex HEP detectors
 - ❑ Implements particle transport in magnetic field-map on GPUs
 - ❑ Ongoing integration with experiments – CMS the most advanced
 - ❑ Can use the same handling of hits (i.e. sensitive detector code) on the CPU by collecting scoring information on GPUs and sending this information back to CPU

- ❑ Geant4 assessment – see next slide ...



Draft Geant4 Assessment Report of the GPU R&D projects AdePT & Celeritas

Alberto Ribon
CERN EP-SFT

AF meeting, 1 February 2024

Geant4 assessment of R&D GPU projects AdePT and Celeritas

❑ Assessment held on December 13-14 2023

❑ Waiting for the official report

❑ Quotes from the current draft report:

*“The panel **does recommend continuing discussions** and exchanges between the two projects about findings, difficulties and solutions. The panel **does not recommend** however **merging** – at least for the time being – the activities of the two projects into a single one.”*

*“The panel proposes a **delta-assessment by December 2024**.”*

*“The two projects have demonstrated the need to evolve in depth the **geometry representation**, shown as the main **limiting factor** today. The two projects anticipates they will be ready by September. December leaves then some time to perform batteries of tests. For this delta-assessment, the panel considers proposing a series of **application benchmarks** together with a series of distributions of various **key observables** and **performance** measurements.*

General Considerations

- *“The panel congratulates the two projects for the impressive amount of work and investigations done. It has appreciated the open discussions and the constructive exchanges, including those on the issues and difficulties met by the projects.”*
 - *“The panel considers that **no key items have been overlooked**. Both projects showed they are well aware of the issues, of the developments and investigations to be carried on, and of the volume of effort to address these.”*
 - *“The projects are capable of LHC-scale detector simulations and **no show-stoppers have been identified today**, for what physics, performances or scalability are concerned.”*
 - Note: so far, this is based on test simulations, not yet on production simulations
- “The projects still have important tasks to address, in particular on the surface-based geometry approach, its memory consumption and scalability to large structures.”*

Simulation

All experiments **happy with their feedback loop with Geant4**. Detailed feedback will be given through the usual channels. Everyone is using / looking forward to use Geant4 11.2.

Threading model changes in Geant4 should be presented early and discussed together with the experiments.

Everyone is interested in **GPU related developments**. Particular interest in G4HepEM, Adept & Celeritas integration.

Everyone is looking forward improvements in VecGeom.

CMS: **parallel initialization** before the start of a run would be very welcome.

ATLAS: We are glad to see that the improvement of the **ParticleHP** physics is on the plan of work for this year.

LHCB: happy to see will to integrate experiments validations in **geant-val**, but it has to go beyond “test beam” data.

ALICE: we welcome Geant4 / **FLUKA** integration. It will allow us to streamline / remove our abstraction layer.

The HEP Software Foundation facilitates cooperation and **common efforts** in High Energy Physics software and computing internationally.



HSF organisation

- ❑ **New conveners appointed in the various WGs**
 - 2024 plans have been presented at several coordination meetings
- ❑ **Many software training activities (presented at the last 2023 LHCC meeting)**
- ❑ **Several Workshops (co-)organised, see next slides ...**

- ❑ **The HSF is turning 10!**
 - Internal review under way ... executive summary to be presented in May at the joint WLCG/HSF workshop at DESY ...

Engagement with other bodies such as the LHCC / WLCG MB


- ❑ **Preparation of LHCC Focus Session on MC Generators**
 - Related-ish to the Event generators' and N(n)LO codes' acceleration workshop in Nov. '23
 - Several preparation meetings held with representatives from the Generators community and WLCG MB software liaisons
 - Report finalised and circulated among community members, presentation at Feb. LHCC meeting

HSF activities: new WG JuliaHEP

- ❑ Initiative after CHEP 2023 to organise a new working group for Julia, which is a very promising language for HEP
 - ❑ See paper [[2306.03675](#)] on the potential
- ❑ First JuliaHEP workshop organised in Erlangen, Germany
 - ❑ 20 people in person + 30 people remote
 - ❑ Great atmosphere and many useful presentations and discussions
 - ❑ Contributions from key Julia developers (including one of the language founders, Stefan Karpinski)
- ❑ Developed training material to help people get started
 - ❑ Plus other tutorials in HPC with Julia, statistics and ML
- ❑ Presentations of large scale projects being tackled in Julia (e.g., new event generators or Bayesian toolkit)
- ❑ Many pieces of “glue” that help connect Julia for the HEP world: ROOT file reader, Geant4, Jet Finding, PDG, etc.
- ❑ Discussion of priorities moving forward that will help focus community work coherently

JuliaHEP Launches

After a lot of rising interest in Julia for HEP in the last few years, the HSF has started a new **JuliaHEP** working group.



We just published a new paper **Potential of the Julia programming language for high energy physics computing** and we're planning the first **JuliaHEP Workshop** in November. Keep an eye out for upcoming Julia events in the **calendar!**



HSF activities highlights



- **Google Summer of Code 2023**
 - 66 proposals for ~39 projects received -> 17 projects with good students accepted
- **Software Training:**
 - Many activities, see dedicated report from last LHCC meeting
- **Data Analysis:**
 - WG held several meetings on Open Data activities, LHC and beyond
 - Discussions started on HS3 (HEP Statistics Serialization Standard)
 - Meeting late November to be advertised shortly
- **Event Generators and Detector Simulation:**
 - Held/plan several meetings on detector simulation, see Indico category
 - Participation in the Geant4 technical forum
 - Tuning workshop, ~90 participants
- **Papers:**
 - “Training and Onboarding initiatives in High Energy Physics experiments”, HSF-TN-2023-01
- **Articles:**
 - The Phoenix Event Display got featured in a CERN Courier, see <https://cerncourier.com/a/event-displays-in-motion/>

HSF activities highlights

Event generators' and N(n)LO codes' acceleration Workshop 13/11/2023

- ❑ **Topics:** How to address difficulties of higher-order calculations
 - Use of heterogeneous hardware, particularly GPUs, How to exploit HPCs effectively
 - Discussion on CERN's involvements in generator software
- ❑ **Outcomes were presented to the LHCC at the Feb. 2024 meeting**

PyHEP “Python in HEP” series of workshops

- ❑ **2023 was the first year with 2 workshops !**
- ❑ **PyHEP.dev 2023**, a new, “developers workshop”, to plan a coherent roadmap and make priorities for the upcoming year
- ❑ **PyHEP 2023**, the usual online “users workshop” (~670 registrations)
- ❑ **PyHEP 2024** – registration is open, as well as the call for abstracts ! Same for **PyHEP 2024.dev** !



Joint WLCG-HSF workshop

- ❑ **Co-organising this joint workshop @ DESY in May 13-17 - Registration is open !**
- ❑ **Programme planning is advancing and will include**
 - Joint session with WLCG on Analysis Facilities
 - Community Software Projects
 - Training Community Lessons and Plans
 - HSF's 10th birthday and plans for evolution

HSF activities highlights – PyHEP.dev workshops



Indico for the 2023 workshop

PyHEP.dev is an in-person, informal workshop for developers of Python software in HEP to plan a coherent roadmap and make priorities for the upcoming year. It complements the [PyHEP Online](#) workshop, which is intended for both developers and physicists.

- ❑ Series of HSF workshops part of the [PyHEP WG](#) activities
- ❑ .dev: new workshops since 2023 complementing the “good old” PyHEP “users” workshops held since 2018 and online since 2020
- ❑ Planning/coordination meeting for *developers* of HEP software in Python
- ❑ [Jul 25–28, 2023 at Princeton](#) (hosted by IRIS-HEP, 38 participants)
- ❑ [Aug 26–30, 2024 in Aachen](#) (IRIS-HEP & ErUM-Data-Hub, targeting 40–50). Registration & call for abstracts open !

- ❑ Preparation: participants set goals for themselves in [GitHub Issues](#), starting a month before the workshop, and followed up on these plans during and after the workshop
- ❑ Format: introductory (lightning) presentations followed by discussions in topical groups
- ❑ Result: [minutes taken here](#). The 2024 edition will also have a formal conclusions document



Cross-experiment projects – IRIS-HEP (USA)

- ❑ Got funded for a 2nd 5-year period – congratulations!

- ❑ Analysis Grand Challenge (AGC) demonstration event took place mid September 2023
 - Multiple implementations for AGC tasks available
 - Reference based on coffea (with optional ServiceX)
 - ROOT RDataFrame implementation (ML functionality added this year)
 - Julia-based implementation
 - Columnflow-based Python implementation
 - AGC implementations being used for performance studies on analysis facilities

- ❑ More AGC to be held ...

- ❑ Data Analysis Pipeline “IDAP”
 - New acronym to more clearly distinguish different efforts that were formerly all happening under the AGC label
 - IDAP to focus on activities that are happening within IRIS-HEP and relate to the pipeline that is used to scale analysis to the HL-LHC. The more outward-facing activities will keep running under the name AGC
 - First meeting took place in February

LHCC request of topics for 2024

LHCC request of topics to be covered in 2024

- ❑ PyHEP.dev developer workshops – hear on outcome regarding the planned roadmap and priorities for the upcoming year
- ❑ CVMFS – collaboration with EESSI and EuroHPC to "train" HPCs to use CVMFS and lowering the threshold for adoption
- ❑ ROOT – issue of more person-power expected
- ❑ Geant4 – assessment in December on Adept vs Celeritas

Charge on Analysis Facilities

- ❑ Follow-up from focus session in 2023
- ❑ The experiments, sites and the HSF Analysis Facilities Forum must engage in a dialogue towards defining use cases and establishing realistic benchmarks for Analysis Facilities
- ❑ LHCC charged the WLCG at their last meeting to oversee the continuous process and to report regularly to the LHCC on the progress and steps taken as described at ->
- ❑ List of questions must be defined first that seek to define the expectations from experiments for Analysis Facilities

Follow-up on the Focus Session on Analysis Facilities held at the 154th LHCC in June 2023

Authors: Catherine Biscarat & Gonzalo Merino, former and current lead reviewers of WLCG at the LHCC, February 2024

Agreed with: LHCC and HL-LHC sw & computing panel

Context

A focus session on Analysis Facilities has been organized during the 154th LHCC week, where the current R&D work happening within the community and discussed in the HSF Analysis Facilities Forum has been presented. The LHCC reported the following:

- The HSF Analysis Facilities Forum has focused on the technical building blocks that can make the user experience more productive. Novel concepts of Analysis Facilities, integrating interactivity, scalability (off-load) and Machine Learning tools are being built.
- The LHCC **recommends** that experiments engage in the process of developing and defining the structure of the future Analysis Facilities and requests they produce a document which defines the use cases in order to establish realistic benchmarks. This process should be coordinated with the HL-LHC Computing and SW review panel. The document is expected to be regularly updated in the process towards HL-LHC.