

# Analysis Ecosystems 2

---

Graeme Stewart, for the workshop organisers



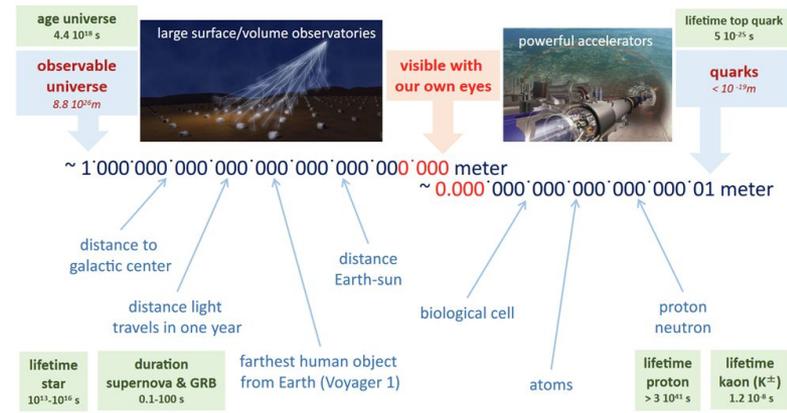
# What are the goals?

- Intense study of the Standard Model

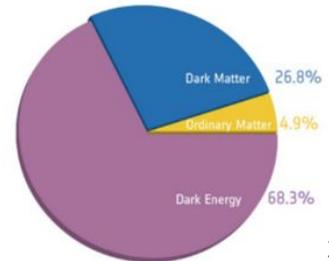
- What are the Higgs properties? LHC and HL-LHC will get us to 3-5% for some couplings
  - And a first handle on Higgs self-coupling
- Flavour sector
  - Better measurements of neutrino mixing angles (cross section is low); mass hierarchy
    - DUNE, HyperK
  - CKM matrix and CP violation
    - HL-LHC, NA62

- Beyond the Standard Model

- We know there are mysteries to uncover! (recent  $m_W$  measurements,  $\mu$  g-2, lepton universality violations)
- Low Higgs mass (aka naturalness)
- SUSY?
- Dark sector - still no idea what dark matter is!
  - Need searches from colliders, direct detection and astrophysics



A **diverse set of experiments** giving a wide range of measurements and **lots of data**  
 We need the best physics analyses we can manage (and cross-experiment too!)



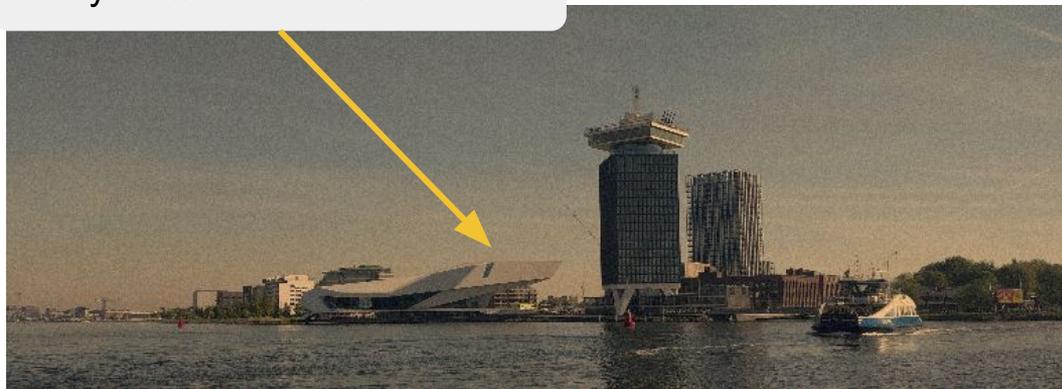
# Five Years Ago...

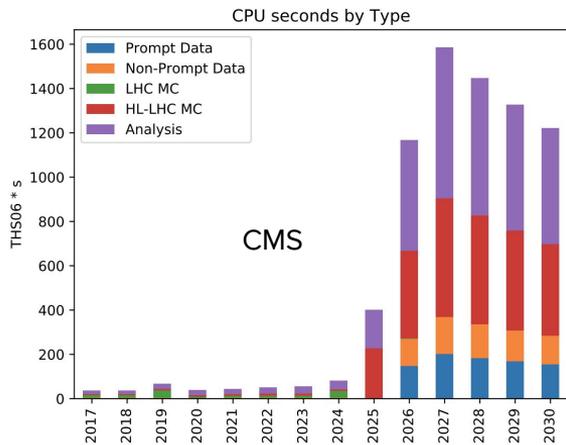
- First analysis ecosystem workshop in Amsterdam
- The big questions were
  - Evolution of ROOT
  - Relevance of data science tools
  - Links to other sciences
- Try to map out the 5-10 year view
  - We knew HL-LHC was coming
- Feed this into the *Community White Paper* view



The screenshot shows the website for the "HEP analysis ecosystem workshop" held from May 22-24, 2017, in Amsterdam. The date and location are circled in blue. The page includes a search bar, a navigation menu with options like "Overview", "Timetable", and "Meeting Participants", and several paragraphs of text. The text discusses the HEP community's history, the evolution of the ROOT ecosystem, and the workshop's purpose in examining the analysis ecosystem with a 5-10 year view.

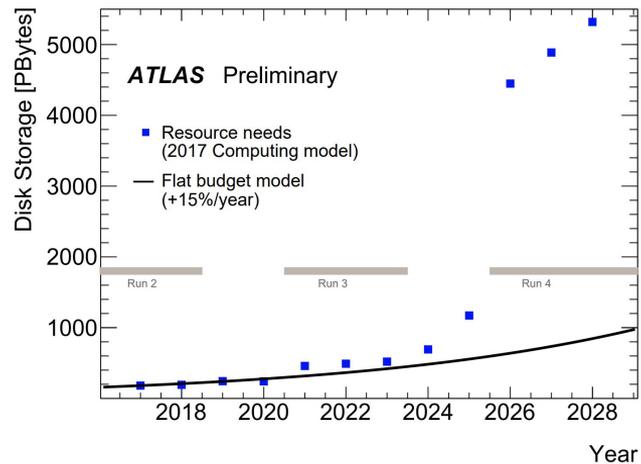
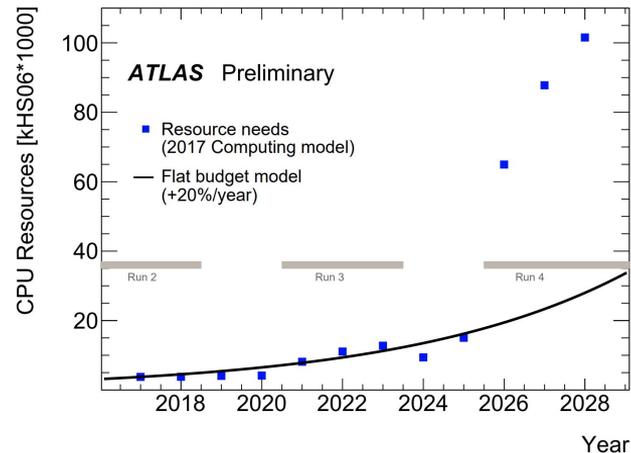
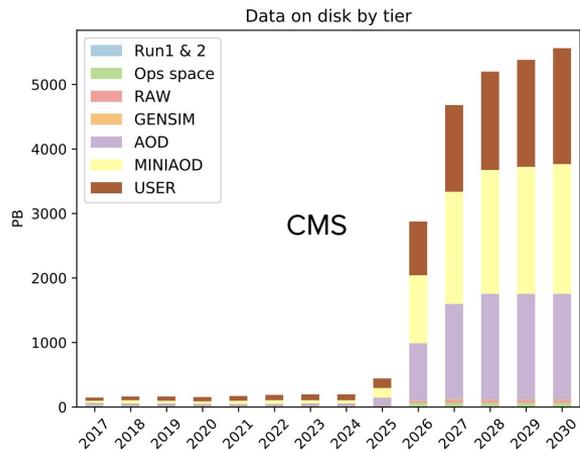
Many of us were here...





# Challenges we faced...

- Plots from the [CWP](#) in 2017
- We knew these were something of a 'worst case scenario'



# AE1 Topics

1. Today's analysis ecosystem and toolset
2. The analysis ecosystem vision
3. Concrete technologies
4. Missing pieces

- There was a lot we didn't know
- “Modern” machine learning was only just starting
- Data formats, size and performance was clearly an issue
- We knew technology and software would evolve... but not quite in which ways!

09:30	--- Registration / Welcome Coffee and Crossaints ---		08:00	The analysis ecosystem vision - Axel Naumann (CERN) Brian Paul Bookeman (University of Nebraska-Lincoln (US)) Gordon Watts (University of Washington (US)) (until 10:30) 0	
10:00	Workshop Introduction and Goals - Elizabeth Sexton-Kennedy (Fermi National Accelerator Lab. (US)) 0	Analysis Ecosystem Introduction-2.pdf Google Slides for Intro	09:00	What hardware and computing facilities do we expect in 5-7 years for analysis? - Fons Rademakers (CERN) 0	
10:20	Today's analysis ecosystem landscape and toolset - Andrea Rizzi (INFN Sezione di Pisa, Universita' e Scuola Normale Superiore, P) Atilia Krasznahorkay (CERN) Paul James Laycock (CERN) (until 12:30) 0		09:20	Perspectives on HEP data analysis software, by an ex-high energy physicist turned (big) data scientist - Max Baak (CERN) Max Baak (CERN) 0	
10:20	ROOT's Place and its Capabilities in the HEP Analysis Ecosystem - Pere Mato Vila (CERN) 0	HSF-ROOT-20170522.pdf	09:40	My vision - Eduardo Rodrigues (University of Cincinnati (US)) 0	
11:10	Data Formats in HEP Analyses - Jakob Blomer (CERN) 0	hep-file-formats.pdf	10:00	My vision - Gordon Watts (University of Washington (US)) 0	
11:30	Data Analysis in ATLAS - Graeme Stewart (University of Glasgow (GB)) 0	Data Analysis in ATLAS.pdf	10:00	My vision - 2017-05-23 - My Vision.pdf 2017-05-23 - My Vision.pptx	
12:00	Data Analysis in CMS - Gianluca Cerminara (CERN) 0	2017-05-22_DataAnalysisWorkflows_HSF.pdf	10:30	--- Coffee Break ---	
12:30	--- Lunch ---		11:00	The analysis ecosystem vision - Gordon Watts (University of Washington (US)) Brian Paul Bookeman (University of Nebraska-Lincoln (US)) Axel Naumann (CERN) (until 12:30) 0	
13:00	Today's analysis ecosystem landscape and toolset - Atilia Krasznahorkay (CERN) Andrea Rizzi (INFN Sezione di Pisa, Universita' e Scuola Normale Superiore, P) Paul James Laycock (CERN) (until 15:00) 0		11:00	My vision - Gerhard Raven (Nikhef National institute for subatomic physics (NL)) Gerhard Raven (Naturkundig Laboratorium-Vrije Universiteit (VU)-Unknown) 0	
13:30	Data Analysis at the Intensity Frontier - Philippe Canal (Fermi National Accelerator Lab. (US)) 0	IntensityFrontiers.pdf IntensityFrontiers.pptx	11:20	ROOT vision (1) - Pere Mato Vila (CERN) 0	
14:00	Data Analysis in LHCb - Eduardo Rodrigues (University of Cincinnati (US)) 0	EduardoRodrigues-2017-05-22_HSFAmsterdam.pdf EduardoRodrigues-2017-05-22_HSFAmsterdam.pptx	11:40	ROOT vision (2) - Axel Naumann (CERN) 0	
14:30	Data Analysis in Belle II - Martin Ritter 0	2017-05-22-analysis-ecosystem.pdf	12:00	Discussion 0	
15:00	--- Coffee Break ---		12:00	--- Lunch ---	
15:30	Today's analysis ecosystem landscape and toolset - Andrea Rizzi (INFN Sezione di Pisa, Universita' e Scuola Normale Superiore, P) Paul James Laycock (CERN) Atilia Krasznahorkay (CERN) (until 17:00) 0		12:30	Concrete technologies - Benedikt Hagner (CERN) David Lange (Princeton University (US)) (until 15:00) 0	
16:00	Last Steps of the Analysis - Andrew Gilbert (KIT - Karlsruhe Institute of Technology (DE)) 0	Last step of the analysis.pdf	13:30	Survey of data formats, conversion tools - Jim Pivarski (Princeton University) 0	
16:30	Round table - Data Analysis 0	ROOTancestryHistory.pdf roundtable (1).pdf Round table (google/)	13:30	Survey of languages for analysis - Dr Sebastian Binet (IN2P3/LPC) 0	
16:30	--- Reception (Cafe' de Jaren - Monument Room) ---		14:10	Swan - Enric Tejedor Saavedra (CERN) 0	
			14:25	XENON1T - Dr Christopher Tunnell (University of Chicago) Christopher Tunnell (Emico Fermi Institute-University of Chicago-Unknown) 0	
			15:00	--- Coffee Break ---	
			15:30	Concrete technologies (until 17:00) 0	
			15:30	Hydra - Dr Antonio Augusto Alves Junior (University of Cincinnati (US)) 0	
			15:45	Femtoocode - Jim Pivarski (Princeton University) 0	
			16:00	IT analytics tool chains - Dirk Duellmann (CERN) 0	
			16:15	Experience with analysis using Tensorflow and coprocessors - Anton Polektov (University of Warwick (GB)) 0	
			16:30	Dark - Sascha Caron (Nikhef National institute for subatomic physics (NL)) 0	
			16:45	Results of analysis tools in HEP survey 0	
			19:00	--- Workshop Dinner ---	

# AE1 Development Conclusions

The outcomes of AE1 were summarised in an 11 page [report](#)

- Ascendancy of Python
  - Critical to better connect ROOT to Python (PyROOT support) and the data science ecosystem
- Yet C++ will persist
  - We didn't see any serious competitor at that time on the performance front
- Modularity is important
  - Containers help with isolation from underlying OS
- Decouple what you want from how you get it (declarative/functional)
  - And make sure provenance is saved
- Event throughput is the golden metric
  - But constraints of *latency* for interactive and *event file size* for storage costs

# AE1 Tools Conclusions

- ROOT is the foundational layer
  - ROOT I/O is the optimised data format for our field
  - Reflection is an enabling technology
- Connecting tools (data bridges) will enable us to make better use of externals (e.g. ML)
- Notebooks could be connected to larger scale resources
  - Useful interfaces, with scale-up potential
- Automation of workflows is necessary
  - Reproducibility, preservation and validation
- Metadata handling is important - should be easier
  - Common core / API would benefit everyone

# Key Points from the Community White Paper

Further discussions distilled this into the [Community White Paper](#)

- “Time to insight” was the catch-phrase of the CWP
  - Optimisation of physicists’ time to perform useful analysis
  - Extract maximum possible scientific potential
- Not easy to capture *exactly* what that means, but we could identify elements that would help
  - Improve the velocity of data reductions to analysable volumes
    - Mostly offline production data reduction, but Turbo Stream did exist (real-time analysis)
  - Scale up the access to resources, allowing rapid turnaround of analysis ideas
- ROOT valued as an integrated and validated toolkit
- Much interest in new data science tools
  - Great potential from massive industry investments, but not specifically optimised for our field
- Reproducibility of analyses was considered rather shaky...
  - Reinterpretation was very desirable and strongly related

# Key Points from the Community White Paper

- Data formats were  $O(100\text{kB})$
- Analysis trains, to produce reduced formats, were established (ALICE) or in development (ATLAS)
  - Diversity of final “flat ntuple” formats was the rule
- “Virtual data” (produce output on demand) was an idea
- Python was emerging as the language of choice for driving analysis
- Functional and declarative languages were considered interesting
  - Because they separate the “what” from the “why”
  - There was a prototype called `TDataFrame`!
- Tools revolution for analysis should happen
  - Git, code review, continuous integration, ...

# Community White Paper - R&D Programme

- “Enable” new open source tools
  - Interoperability and dynamic plugin, swap in/out components
- Data exchange between different tools
  - ML was highlighted as a case-in-point, where external libraries were gaining ground
- Prototype a low-latency, high-capacity analysis facility
  - With caching and query based interfaces
- Support Python further
  - ROOT Python bindings were to be improved
- Declarative programming model prototypes
- Interpretation gateways
  - Based on analysis preservation and reinterpetation tools

From there to here

---

# ROOT



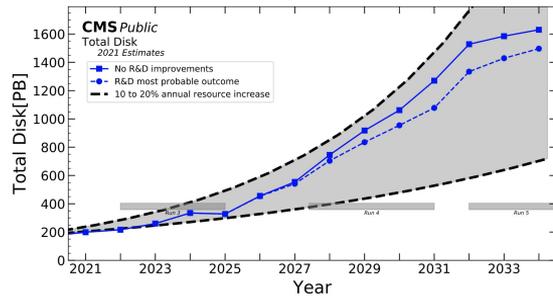
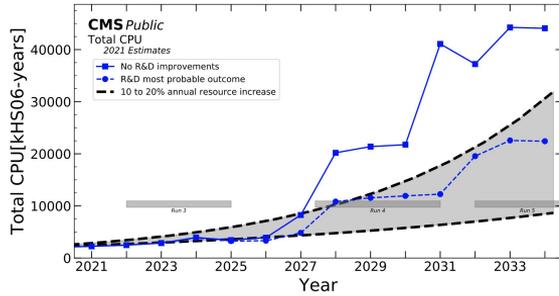
- PyROOT
  - New PyROOT came in 2020: modern, interoperable, Pythonic (RooFit has a proper Python interface)
  - Greatly improved ROOT integration with NumPy and Pandas
- RNTuple
  - Smaller and faster, as well as support for object stores
- RDataFrame
  - Great strides in the development of a well integrated declarative interface
  - Now used in production by a significant fraction of analyses
  - Plus Distributed RDataFrame
    - Proof successor and building block for analysis facility work
- Installation
  - Now much easier than in the past, e.g., distribution via [Conda](#) or Homebrew
- SWAN
  - Connecting storage, software and compute with notebook interfaces (which becomes a form of Analysis Facility)
    - Scaling to 200-250 users daily



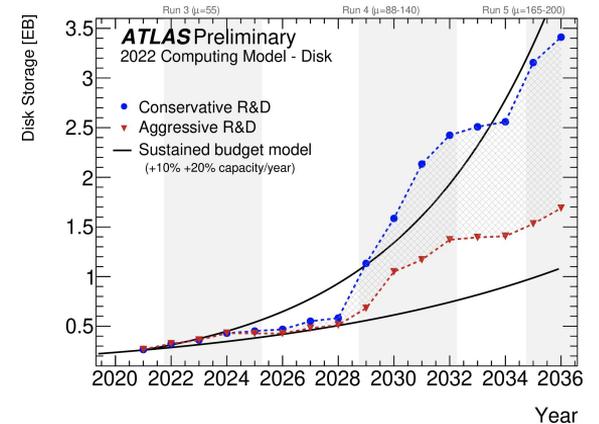
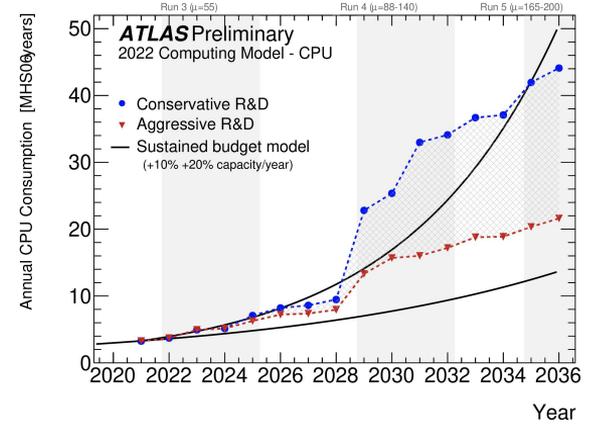
# Scikit-HEP

- Still a newborn at the Amsterdam meeting (DOB Autumn 2016)
- Now the major provider/aggregator of HEP Python packages
  - Awkward Array
  - Boost Histogram / Hist
  - Iminuit
  - zFit (affiliated)
  - Cabinetry
  - Pyhf
  - Uproot
- A lot of packages that ease the integration of HEP data with Data Science toolkits
  - Bringing some order into what could be quite chaotic
- A vibrant ecosystem growing up around Scikit-HEP

# Experiment Progress and Planning



- Key enabler for analysis has been reduced sizes of data formats
  - NanoAOD
  - DAOD\_PHYSLITE
- Use of real-time/trigger analysis in LHCb and ALICE is the cornerstone of Run-3 success
- Common point is reduce the resource footprint per physics useful event
  - More events fit in the resource envelope



# IRIS-HEP, SWIFT-HEP, CERN EP R&D...

## Analysis Systems (WP5)

### Summary

The analysis systems work package 5 focuses on two aspects of HEP data analysis: efficient iteration and portability of code.

### Progress

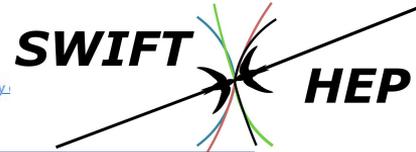
The analysis work package is due to start on the 1st of April 2022. We are currently in the planning and hiring phase.

### Collaboration with other HEP efforts

Planned collaboration with the HSF and IRIS-HEP by contributing to the Analysis Grand Challenges in particular.

### Meetings, mailing list and contacts

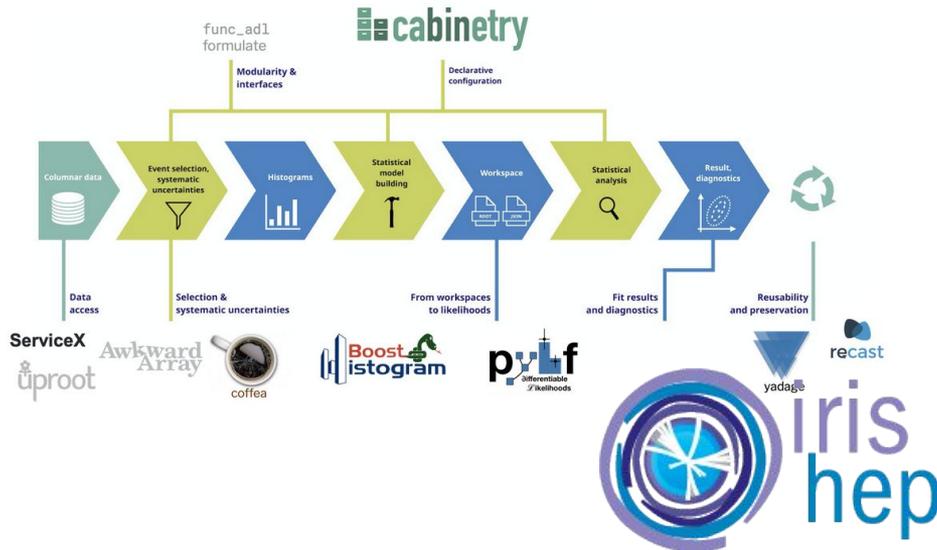
- [Meetings](#)
- [Forum](#) (SWIFT-HEP wide)
- Convenors: Luke Kreczko (CMS/LZ), Eduardo Rodrigues (LHCb) [\[contact by:\]](#)



- Attract new funding to the field...

## Analysis Systems

The goal of the Analysis Systems focus area is to develop sustainable analysis tools to extend the physics reach of the HL-LHC experiments by creating greater functionality, reducing time-to-insight, lowering the barriers for smaller teams, and streamlining analysis preservation, reproducibility, and reuse.



← 7. Software

## 7.4. EFFICIENT ANALYSIS

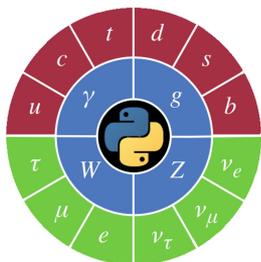
Analysing collision data efficiently is essential for HEP. Since the LHC start-up, the time needed between dataset arrival and the delivery of scientific results was greatly reduced. Analysis software significantly improved and approaches to analysis were streamlined: for example, analysis trains were introduced by experiments to combine analysis steps from different users and execute them together, reducing access to the same input data. Nevertheless, such improvements will not be sufficient in the future.

# Building Communities



- HSF has been playing a role to help strengthen our software development communities
  - There are clearly long established communities of course - **ROOT** is a case in point
- [Data Analysis Working Group](#)
- [PyHEP](#)
- [Adelaide CHEP pre-Workshop](#) on Analysis
- [HSF Analysis Facilities Forum](#)

DESY NAF (15+5)	Christian Voss et al.
15:55 - 16:15	
Break	
16:15 - 16:25	
SWAN over Spark and HTCondor at CERN (15+5)	Enric Tejedor Saavedra
16:25 - 16:45	
US.CMS analysis facilities: coffea-casa AF, EAF and others (15+5)	Lindsey Gray
16:45 - 17:05	
Distributed Disk-based national facility at INFN (15+5)	Mirco Traccoli
17:05 - 17:25	
AF activities in LHCb (15+5)	Donatella Lucchesi
17:25 - 17:45	
Break	
17:45 - 17:55	
AF activities in DOE multi-purpose computing centers (25+5)	Burt Holzman et al.
17:55 - 18:25	
Analysis on Cloud Facilities (15+5)	Fernando Harald Barreiro Megino
18:25 - 18:45	
A kubernetes-based AF at UChicago (15+5)	Fengping Hu et al.
18:45 - 19:05	
Distributing Production-Ready Software (15+5)	Brian Hua Lin
19:05 - 19:25	



# LHCC HL-LHC Review

- Important review of the preparations for HL-LHC
  - Rather positive view of the work from all quarters: “From a technical perspective, all activities and the experiments are showing excellent progress”
  - And encourage that to continue: “increased investments in R&D activities are essential”
- Praise for
  - The role of the HSF in discussions regarding data science tools
  - ROOT, addressing the central needs of the experiments and future challenges
  - Flexibility and agility that comes from the data science tools realm
- Recommendations
  - Even closer interactions between ROOT and PyHEP projects
    - In fact, this exists to a large extent already, with many people working together - this workshop is another case-in-point
  - Increased input from the experiments in driving the evolution of tools
  - Work on benchmarks and performance indicators to make objective assessments
    - Analysis turnaround time, resource requirements, ease-of-use, reusability, ...

# Experiment Analysis Model Overview

Some common themes from the ROOT Workshop's [Analysis Model](#) presentations

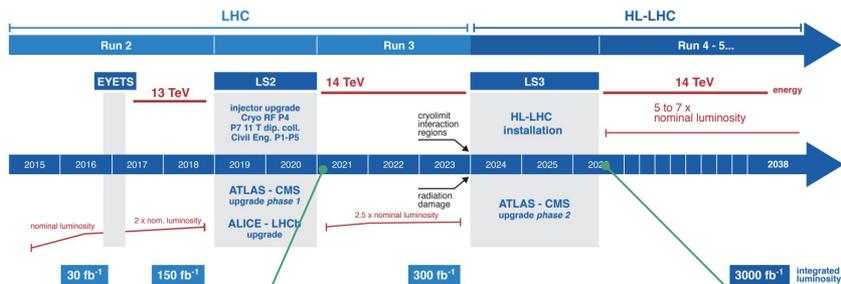
- ROOT is the trusted and used format for our data
  - Some other formats are used during processing (e.g., Arrow, ONNX, ...)
- Central calibrations are generally considered a good thing
  - Smaller data formats, correct application
- Systematic variations and slow grid turn-around are bugbears
- Python and C++ used, different experiments mix differently
  - High level and low level components being used, e.g., Bamboo for a higher level analysis language; DASK for lower level task parallelism
  - Analysis ergonomics are important, to improve user experience and effectiveness
- Use of REANA as a way to capture final analysis workflows seems to be increasing
  - Will this become even more important in the future?

# LHC / HL-LHC Plan

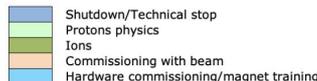


## HL-LHC

- We have more time than we thought in 2017
- On the other hand we are also late with what we thought we could do by 2020!



Last updated: January 2022



# Analysis Ecosystems II

- Lots of progress since AE1
  - Many things that were inchoate ideas have become a reality
- Yet, we are far from finished
- There are lots of new development and new tools
  - R&D is a good thing - it happens in all healthy projects
- HL-LHC and similar data volumes are still scary and exciting
- AE2 comes at a very good time to learn from the past and guide us into the future
- We want to aim for a reasonable consensus on directions to pursue
  - We might not agree on everything
- We have 6 main topics we identified to focus the discussions...
  - But also time to discuss other things: BoFs and Breakouts!

# Topic Questions

---

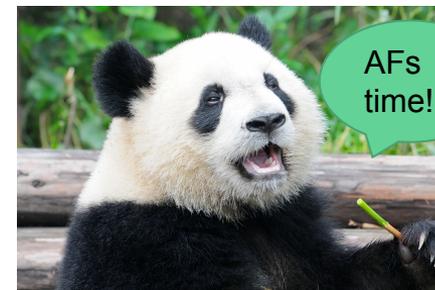
# Analysis Facilities

## Monday morning plenary

- New developments in Analysis Facilities
- Analysis frameworks and Analysis Facilities: user experience
- Impact on Analysis Facilities in the context of DOMA evolution

## Monday afternoon parallel

- Status report about Analysis Facilities in various experiments and CERN
- Integrating different types of resources/"aaS" available on demand: public and private clouds, and Kubernetes (k8s): what will be the best "recipe" for AF?
- "How to marry interactivity with distributed computation?" aspects related to e.g. UI design, resource management policies and debugging tools
- Analysis containerization, feedback from point of view of users, AF architects, resource providers
- Idea of "dream" Analysis Facility, metrics and evaluation of successful AF



# ML tools and differentiable computing workflows

- General ML Questions
  - In industry there is a very large trend to **very large pretrained “foundation models”** (similar to our “reco”, but diffable), that can then be fine-tuned, will we see the same in HEP?
  - What notions of **Uncertainty Quantification** are meaningful in HEP?
  - Would we expect **large-scale self-supervised representation learning** to find representation that meaningfully differ from our hand-engineered hierarchical features (tracks -> clusters -> particles -> jets ...)
  - How can we **integrate ML** more into “standard analysis” (away from “extract then reintegrate”)
  - First **Generative Models in production**. Do we expect a meaningful extension of this in the next 5 years? (Full detector GMs?)
  - What **common infrastructure** is needed to speed up ML development ? Common Datasets, Special ML tools/APIs in AF (e.g. KubeFlow)
  - How large can neural blocks become (e.g. single-shot all-detector pFlow, end-to-end GNN tracking?) What are indispensable structural components / inductive biases? Will Rich Sutton’s [“Bitter Lesson”](#) kick in?
- Differentiable Programming
  - What are examples where **downstream feedback would meaningfully change upstream tasks?** I.e. where **does end-end learning beat hierarchical per-task training?** Is the gain worth it?
  - How do we assess, what’s useful to rewrite vs effort to integrate AD into existing code?
  - Will we need to pass gradients across language layers?
  - What’s the relationship between neural surrogates vs making the actual code differentiable?
  - We use ML for both acceleration and learning. More DiffProg less Fixed Layers → less benefits from HW acceleration

# “Real-time” online/trigger-level analysis

- What are the AE challenges that come from doing Real Time & Trigger Analysis?
  - For heterogeneous hardware in triggers, how do we guarantee that trigger software and offline software produce the same (indistinguishable) results?
    - What level of difference is tolerable, and what tests do we need to demonstrate that we can tolerate the differences?
  - Analysis moving towards smaller and smaller data formats. How to be sufficiently flexible and accommodate
    - The need for some re-calibration prior/during analysis workflow
    - Understanding anomalous events requiring larger data formats
  - How does RTTA interact with MC in terms of Analysis Ecosystems?
  - Can analysis workflows be instantiated in the trigger software, e.g. make histograms, Dalitz plots?
  - Can studies performed in real time help to better understand both detector performance and particle physics questions?
  - Could ML instantiated in accelerators help? If so, would WLCG need accelerators?

# Analysis User Experience and Declarative Languages

- **Analysis user experience**
  - do we understand “average” analyzer experience, or only know from a particular group (bias)
    - Should we design a survey?
  - specific requirements for **small / new experiments**?
  - what fraction of analyzers use “new” approaches we may recommend?
  - what might be missing in “modern” envisioned workflows?
- **New programming languages and analysis description languages (ADLs)**
  - what is the way forward here? Several examples, limited (?) adoption so far
  - what is the future of Julia and other new languages in HEP analysis?
- **Benchmarks, documentation and examples**
  - how to benchmark full pipelines that implement cutting-edge analyses?
  - how to document libraries / frameworks with relevant and up-to-date examples that go beyond minimal: also relation to things like funding / dedicated positions, etc.
- **Interoperability** (e.g. between ROOT and Python/data science world)
  - where do we stand and what needs to be improved?

# Analysis on reduced formats or specialized inputs

- How are reduced formats designed, and how are they used in different experiments? (skims or common ntuple formats such as NanoAOD or PhysLite)
  - Addressed in plenary session tomorrow!
- What about analyses that cannot use the reduced format? (eg. searches for long-lived particles that need low-level detector information)
  - Addressed in the parallel session tomorrow!
- Looking ahead, what R&D is being done to make this problem easier?
  - Can we make it easy to add new variables to the existing format, eg using ROOT friend trees?
  - Is it feasible to use object stores to get rid of data tiers and ntuples altogether?
  - Addressed in the parallel session tomorrow!

# Metadata and systematics handling

Metadata infrastructure for analysis has historically been lower priority

- Definition a little murky, solutions *ad hoc* and divergent
- In 2021, HSF DAWG summarised current status and a wishlist for future analysis metadata systems: <https://arxiv.org/abs/2203.00463>
- Today: technical panel review, this week follow up towards concrete proposal

Systematic uncertainties are crucial, and a major contribution to analysis CPU

- What are the major challenges that are common across experiments?
  - “Philosophical” differences between conceptually similar ATLAS/CMS:  
How to distribute systematics in code? Degree of dependence on EDM
- How can we collaborate on solutions to the systematics workload?
- How does systematics processing intersect with new trends in analysis computing?

# Practical Points and Outcomes

---

# Introduce Yourself

- Let's get to know one another...
  - Please fill in a few words on your background and interests



**Graeme  
Stewart**

*I am a staff scientist at CERN, working in the Software for Experiments group (EP-SFT)*

*My research: I help to run the HSF and make it work as a community organisation; I also look after software work packages in the CERN EP-R&D programme and the AIDAInnova project*



**My expertise is:** Distributed computing, software development, wrangling HEP software projects

**A software and computing problem I'm grappling with:** Ensuring that our field will train the next generation and support HEP software developers with meaningful careers

**I've got my eyes on:** Software development communities, differentiable programming and compute accelerators

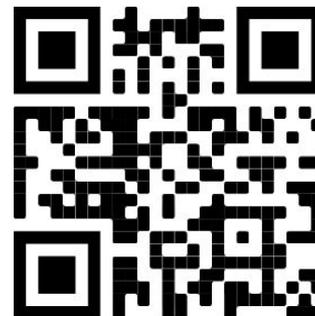
**I want to know more about:** How we build effective software ecosystems, with the stability to do useful things, but the dynamism to try out new stuff (and ensure it still works in 30 years time!)



**Oksana  
Shadura**

*Software engineer at University Nebraska-Lincoln, IRIS-HEP Analysis Grand Challenge Co-coordinator*

**My research:** Development of a prototype of Analysis Facility for HL-LHC



**My expertise is:** Analysis Facilities and Data Organization, Management and Access (DOMA)

**A software and computing problem I'm grappling with:** development of new cyberinfrastructure for data analysis, integration of data delivery and data organisation & management solutions

**I've got my eyes on:** software challenges for HL-LHC

**I want to know more about:** to understand better community requirements or how we are planning to organise more efficient data analysis for HL-LHC

- <https://bit.ly/3yM4a6J>



# Practical Matters

- Dinner
  - If you are coming to the dinner on Tuesday and you have a dietary requirement that you did not yet tell us about in the survey, please *let Michel know immediately*
- Live notes
  - In the best traditions of the HSF we have a [Google Doc](#) for live notes
    - Take care to put your comments in the correct place
    - Parallel sessions have parallel documents - follow the links in the master document
  - You may also ask questions here
- Birds of a Feather and additional discussion topics
  - We cannot anticipate everything, even with our 6 well crafted topics!
  - In addition to informal discussions we have the option to organise some Birds of a Feather or Additional Discussions
    - Please write down what you would like to discuss [here](#)
    - And email [Michel](#) and [Antoine](#) who will find you a room (be sure to send to both)
  - We have Tuesday PM for these chats

# Topics and Outcomes

- Each of the topics we focus on has
  - Some plenary session time
    - Big picture, introduction, general discussion
  - Some parallel session time
    - Drill down to the details
- On Wednesday we will hear short summaries from the convenors about the discussions and outcomes
- Then we will split into groups to start the process of writing up succinct conclusions
- The final aim is to produce a concise outcomes document, with R&D topics and ongoing activities, within a few weeks

# Invasion of Ukraine by the Russian Federation

The war in Ukraine has disrupted many things and education & research is of course one of them.

We have some funding to support Ukrainian students for research activities this summer and perhaps also this coming academic year, and are connecting with people to put together a list of interested students. We may succeed in getting some of them to CERN, otherwise they would be remote (either in Ukraine or elsewhere in Europe) and projects would be similar to GSoC.

If you have still-open projects appropriate for summer students and/or bandwidth for additional mentoring activities for students, please contact Peter Elmer and Oksana Shadura.

