



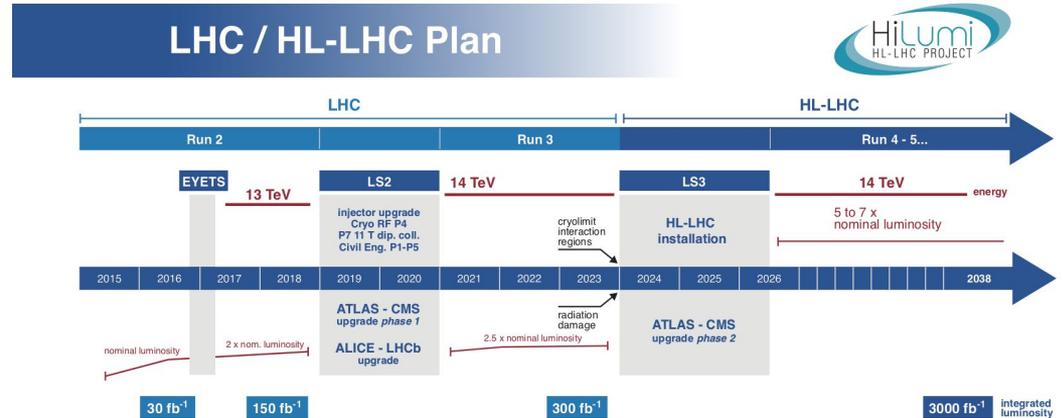
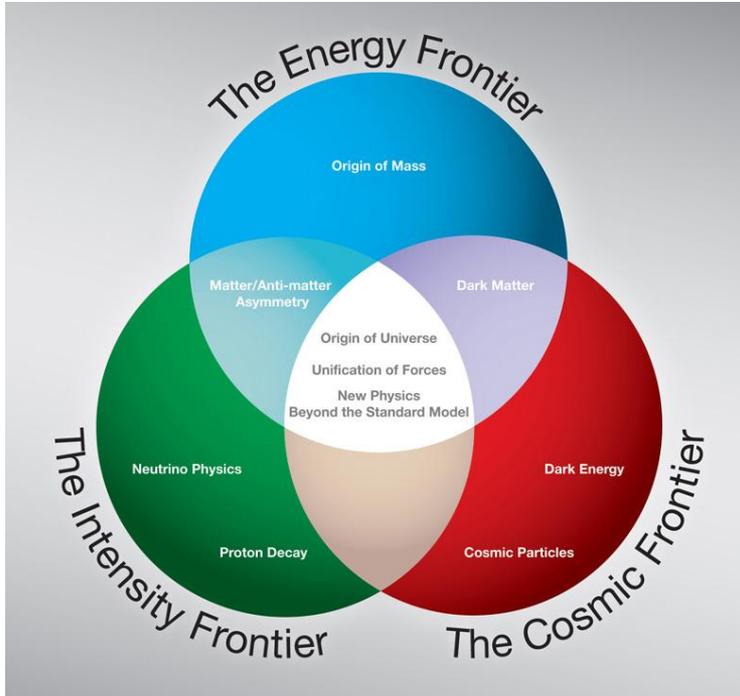
HEP Software Foundation Update

Graeme Stewart, CERN EP-SFT

HSF in the last 12 months

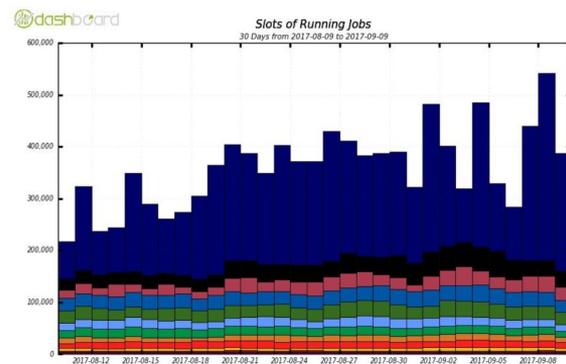
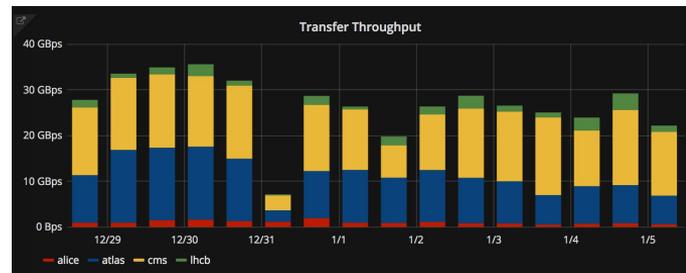
- Many thanks for the opportunity to speak to the ESC consortium
- This talk is a follow on from [the talk](#) from Torre Wenaus gave last year
 - Mostly an update, though I will remind you again of some of the main points
- It was an important year for the HSF
 - Followed through and defined the Community White Paper Roadmap
 - Engagement with many strategic bodies at labs and in funding agencies
 - First formal FTE commitment to HSF from CERN
 - Promoting ground-up activities
 - Following through with workshops and community led meetings
- But let's start with the *whys* before the *whats* and *hows*...

Experimental Particle Physics, Frontiers and Plans



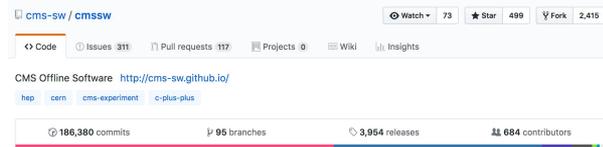
HEP Software and Computing

- High Energy Physics has a vast investment in software
 - Estimated to be around 50M lines of C++
 - Which would cost more than 500M\$ to develop commercially
- It is a critical part of our physics production pipeline, from triggering all the way to analysis and final plots as well as simulation
- LHC experiments use about 600k CPU cores every hour of every day and have around 400PB of data stored on disk and 600PB on tape
 - We are in the exabyte era already
- This is a huge and *ongoing* cost in hardware and human effort
- With significant challenges ahead of us to support our ongoing physics programme



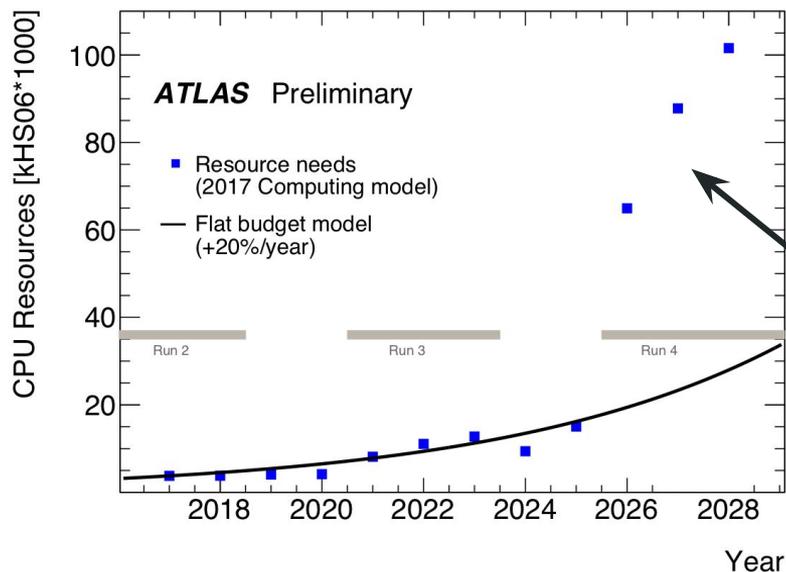
athena

ATLAS Experiment main repository for Athena code

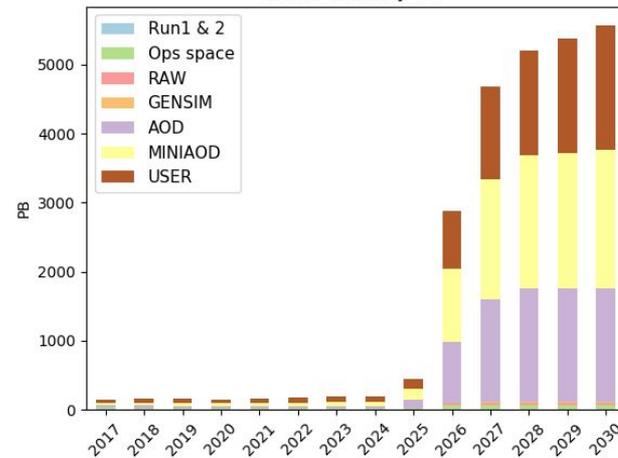


Challenges for the Next Decade

- HL-LHC brings a huge challenge to software and computing
 - Both rate and complexity rise



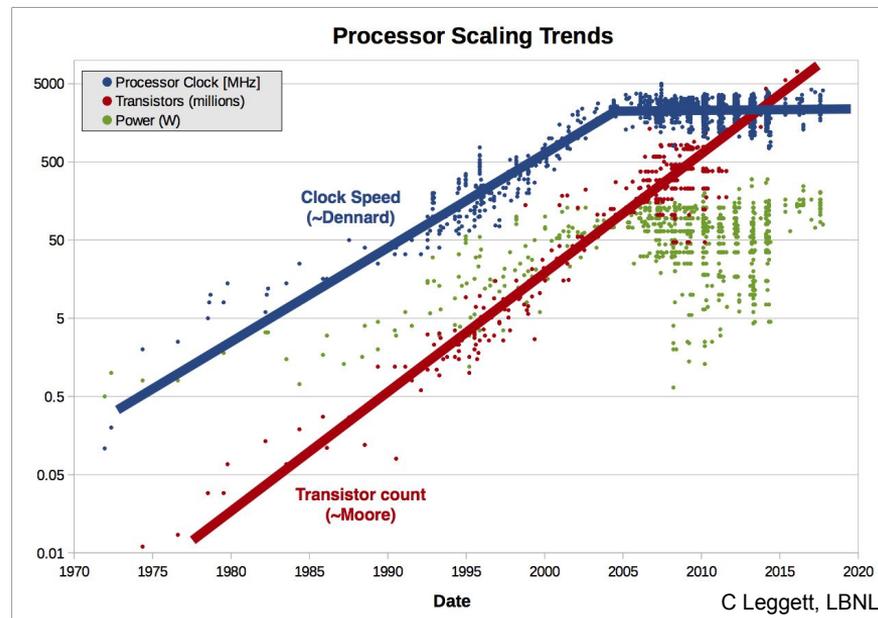
CMS
Data on disk by tier



- Not just a simple extrapolation of Run 2 software and computing
 - Resources needed would hugely exceed those from technology evolution alone

Processor evolution

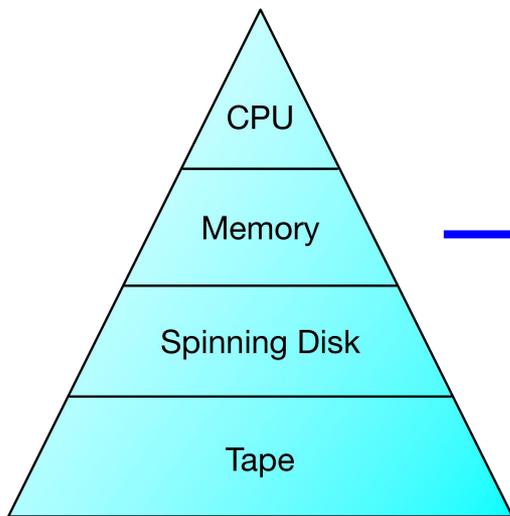
- Moore's Law continues to deliver increases in transistor density
 - Doubling time is clearly lengthening
- Clock speed increases stopped around 2006
 - No longer possible to ramp the clock speed as process size shrinks (Dennard scaling failed)
- So we are basically stuck at $\sim 3\text{GHz}$ clocks from the underlying Wm^{-2} limit
 - This is the *Power Wall*
 - Limits the capabilities of serial processing
 - CPU based concurrency still in development for LHC Run 3



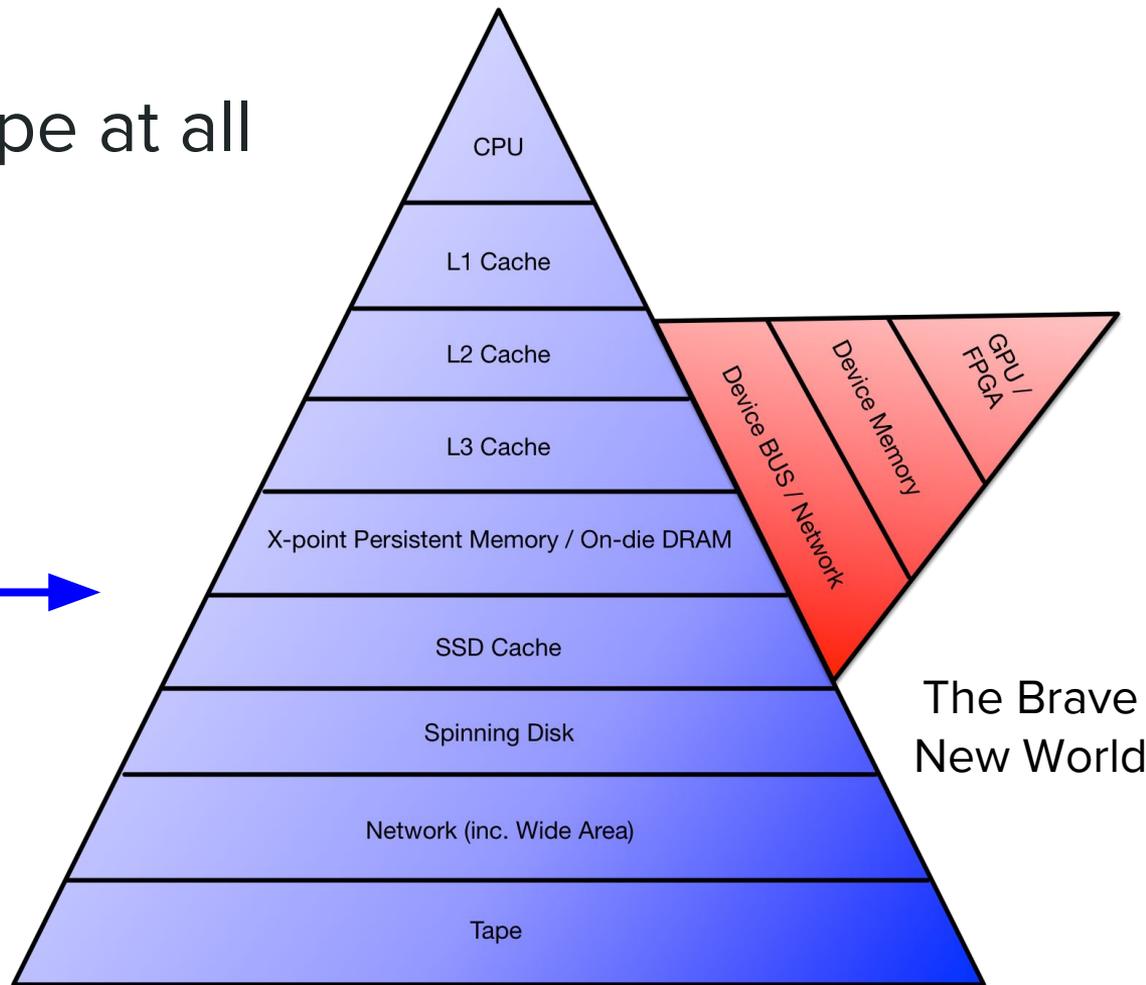
Accelerated computing devices (GPUs, FPGAs) offer a different model

- *Potentially much greater throughput*
- *Still many unresolved issues for legacy code and complexity of heterogeneous processing*

Shifting landscape at all levels



The Good Old Days



The Brave New World

HEP Software Foundation (HSF)

- Founded in 2014, explicitly to address current and future software challenges in common
 - Our software must evolve to meet the challenges of future detectors and technology evolution
 - Need to exploit all the expertise available, inside and outside our community
 - New approaches needed to overcome limitations in today's code
- Cannot afford any more duplicated efforts
 - Each experiment has its own solution for almost everything (framework, reconstruction algorithms, geometry, data model, ...)
 - Much of it not in good shape already
- The goal of the HSF is to facilitate coordination and common efforts in software and computing across HEP in general
 - Our philosophy is bottom up, a.k.a. *do-ocracy*

Compelling case to also work in common with like minded organisations in other science disciplines

- *Square Kilometre Array Telescope, Research Software Engineers, etc.*
- *ESC seems also to be a natural fit*

Community White Paper



- Charge from the WLCG in July 2016
 - Anticipate a "software upgrade" in preparation for HL-LHC
 - Identify and prioritize the software research and development investments
 - i. to achieve improvements in software efficiency, scalability and performance and to make use of the advances in CPU, storage and network technologies
 - ii. to enable new approaches to computing and software that could radically extend the physics reach of the detectors
 - iii. to ensure the long term sustainability of the software through the lifetime of the HL-LHC
- Workshops in 2017
 - San Diego (January) to Annecy (June)
 - Many active working groups around key topic areas producing white paper inputs
- Editorial board digested this and produced a final roadmap paper at the end of 2017

A Roadmap for HEP Software and Computing R&D for the 2020s

- 70 page [document](#) on arXiv
- 13 sections summarising R&D in a variety of technical areas for HEP Software and Computing
 - Almost all major domains of HEP Software and Computing are covered
- 1 section on Training and Careers
- Large support for the document from the community
 - 310 authors from 124 institutions
- Journal submission to [Computing and Software for Big Science](#) is underway

Contents

1	Introduction	2
2	Software and Computing Challenges	5
3	Programme of Work	11
3.1	Physics Generators	11
3.2	Detector Simulation	15
3.3	Software Trigger and Event Reconstruction	23
3.4	Data Analysis and Interpretation	27
3.5	Machine Learning	31
3.6	Data Organisation, Management and Access	36
3.7	Facilities and Distributed Computing	41
3.8	Data-Flow Processing Framework	44
3.9	Conditions Data	47
3.10	Visualisation	50
3.11	Software Development, Deployment, Validation and Verification	53
3.12	Data and Software Preservation	57
3.13	Security	60
4	Training and Careers	65
4.1	Training Challenges	65
4.2	Possible Directions for Training	66
4.3	Career Support and Recognition	68
5	Conclusions	68
	Appendix A List of Workshops	71
	Appendix B Glossary	73
	References	79



Time to adapt for big data

Radical changes in computing and software are required to ensure the success of the LHC and other high-energy physics experiments into the 2020s, argues a new report.

It would be impossible for anyone to conceive of carrying out a particle-physics experiment today without the use of computers and software. Since the 1960s, high-energy physicists have pioneered the use of computers for data acquisition, simulation and analysis. This hasn't just accelerated progress in the field, but driven computing technology generally – from the development of the World Wide Web at CERN to the massive distributed resources of the Worldwide LHC Computing Grid (WLCG) that supports the LHC experiments. For many years these developments and the increasing complexity of data analysis rode a wave of hardware improvements that saw computers get faster every year. However, those blissful days of relying on Moore's law are now well behind us (see panel overleaf), and this has major ramifications for our field.

The high-luminosity upgrade of the LHC (HL-LHC), due to enter operation in the mid-2020s, will push the frontiers of accelerator and detector technology, bringing enormous challenges to software and computing (*CERN Courier* October 2017 p5). The scale of the HL-LHC data challenge is staggering: the machine will collect almost 25 times more data than the LHC has produced up to now, and the total LHC dataset (which already stands at almost 1 exabyte) will grow many times larger. If the LHC's ATLAS and CMS experiments project their current computing models to Run 4 of the LHC in 2026, the CPU and disk space required will jump by between a factor of 20 to 40 (figures 1 and 2).

Even with optimistic projections of technological improvements there would be a huge shortfall in computing resources. The WLCG hardware budget is already around 100 million Swiss francs per year and, given the changing nature of computing hardware and slowing technological gains, it is out of the question to simply throw

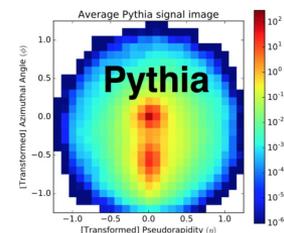
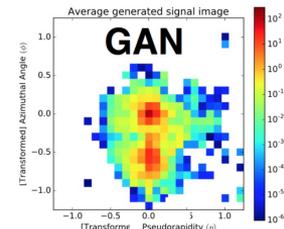
*Inside the CERN computer centre in 2017.
(Image credit: J Ordan/CERN.)*

- Attracted considerable attention
- Community White Paper article published in April [CERN Courier](#)
- Notable presentations:
 - CERN Scientific Computing Forum
 - European Committee on Future Accelerators
 - CHEP Plenary

Let's pick a few highlights...

Detector Simulation

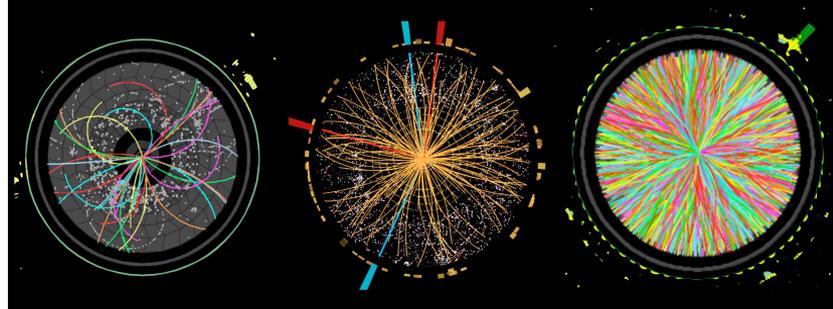
- **Simulating our detectors consumes huge resources today**
 - Remains a vital area for HL-LHC and intensity frontier experiments in particular
- **Main R&D topics**
 - **Improved physics models** for higher precision at higher energies (HL-LHC and then FCC)
 - Hadronic physics in LAr TPCs needs to be redeveloped
 - Adapting to **new computing architectures**
 - Can a vectorised transport engine be demonstrated to work in a realistic prototype (GeantV early releases)? How painful would evolution be (re-integration into Geant4)?
 - **Fast simulation** - develop a common toolkit for tuning and validation of fast simulation
 - How can we best use Machine Learning profitably here?
 - **Geometry modelling**
 - Easier modelling of complex detectors, targeting new computing architectures



Machine learning simulated calorimeter

R&D Outlook: Community is well organised and actively pursuing many lines

Software Trigger and Event Reconstruction



- **Move to software triggers is already a key part of the program for LHCb and ALICE already in Run 3**
 - ‘Real time analysis’ increases signal rates and can make computing more efficient (storage and CPU)
- **Main R&D topics**
 - Controlling charged **particle tracking resource consumption** and maintaining performance
 - Do current algorithms’ physics output hold up at pile-up of 200 (or 1000)
 - Can tracking maintain low p_T sensitivity within budget?
 - Detector design itself has a big impact (e.g., timing detectors, track triggers)
 - Improved use of **new computing architectures**
 - Multi-threaded and vectorised CPU code, use of GPGPUs and possibly FPGAs
 - Robust **validation** techniques when information will be discarded
 - Using modern continuous integration, multiple architectures with reasonable turnaround times
 - **Reconstruction toolkits** can help adapt to experiment specificities: ACTS, TrickTrack, Matriplex

R&D Outlook: A lot of projects in healthy states - keep up level of cooperation and sharing
(Connecting the Dots; Tracking Kaggle Challenge)

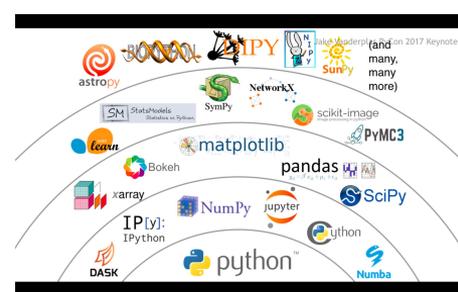
Data Analysis and Interpretation

- **Today we are dominated by many cycles of data reduction**

- Aim is to reduce the input to an analysis down to a manageable quantity that can be cycled over quickly on ~laptop scale resources
- Key metric is ‘time to insight’

- **Main R&D topics**

- How to **use the latest techniques** in data analysis that come from outside HEP?
 - Particularly from the Machine Learning and Data Science domains
 - Need ways to seamlessly interoperate between their data formats and ROOT
 - Python is the *lingua franca* here, thus guaranteeing our python/C++ bindings is critical
- New Analysis Facilities
 - Skimming/slimming cycles consume large resources and can be inefficient
 - Can **interactive data analysis clusters** be set up? SWAN, Spark, Dask interesting
 - Characterised by rapid column-wise access reads, with writes of new columns



R&D Outlook: Many potential directions, no clear overall structure, certainly needs good exchange of information

Data Management and Organisation



- **Data storage costs are a major driver for LHC physics today**

- HL-LHC will bring a step change in the quantity of data being acquired by ATLAS and CMS
- Notwithstanding **serious reductions** in the data stored by the experiments we need to optimise management and access

- **Main R&D topics**

- Adapt to new needs driven by changing algorithms and data processing needs, e.g.,
 - The need for **fast access to training datasets** for Machine Learning
 - Supporting **high granularity access** to event data
 - **Rapid high throughput** access for a future analysis facility (previous slide)
 - Efficient processing at sites with **small** amounts of **storage** (pre-stage buffer better than cache)
- **Consolidate** storage access interfaces and protocols
- Support **efficient hierarchical access** to data, from high latency tape and medium latency network

R&D Outlook: One part of the solution for many different areas,
needs coordinated work

Guiding Strategy for the Roadmap

- HEP faced many computing challenges before other communities and has developed over the decades a lot of community-specific solutions
 - Mainly for good reasons!
 - Several HEP-tools adopted by some other communities, e.g. GEANT4 and ROOT, and WLCG itself is a model/driver for large-scale computing adopted by some other disciplines
- But the world changed: other scientific communities and industry facing some similar challenges and HEP must be able to benefit from them
 - Machine learning, distributed analysis, distributed infrastructure
- Does not mean that we have drop-in replacements for our solutions
 - Challenge: find the proper integration between our community tools and the available technologies outside, maintain the necessary backward compatibility/continuity and **long-term sustainability**
 - As illustrated in CWP chapters, not one single approach for every topic: several paths for moving in this direction are part of the roadmap



HSF Activity Highlights

- Of course the HSF has not just been doing the CWP in the last year...



Copyright and Licensing

- We continue to work in this much neglected area in HEP software
 - Much code exists with **no clear copyright or licence**
 - The issues of large and deep stacks of experiments' software were often neglected in documents, advice and discussions up to now
 - It's sort of combinatorial problem, considering multiple license interactions
- LHC experiments continuing to be more open with their software
 - Goal is to maximise our useful user base and interactions with others
 - CMS and ALICE went open source a few years ago, ATLAS and LHCb are now almost there too
 - Good interactions with Belle II, who are also trying to resolve this issue
- GPL licenses have become disfavoured as they place obligations on any users can can inhibit collaboration (e.g., industrial)
 - ATLAS, CMS and ALICE **want non-GPL licenses**
 - Matches shifts at CERN, e.g., Indico moving from GPL to MIT
 - We made **significant progress** in moving packages like HepMC and DD4hep to *LGPL*
 - Widespread **use of GPL by theory community** still affects us greatly

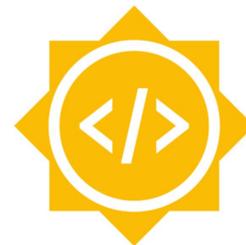
Packaging



- Packaging is one of the de facto areas of common interest between experiments
- Building and deploying our software is a significant task and there is much duplicated effort between experiments
- HSF Packaging Group decided to formalise the problem we are trying to solve
 - Write down the actual use cases we have
 - Recognise that CVMFS and Containers simplified the problem a lot for us
 - Use cases can be enabled or become redundant as technology develops
 - Now trying to test drive packaging tools to see how well they perform
- Active candidates we're looking at
 - Nix - pure functional package manager
 - Portage - from the Gentoo Linux distribution
 - Spack - from LLNL

Training

- Equipping our community with the right set of skills to do its work is a big challenge
 - Especially given the evolution of technology today
 - Goal of software training is to maximise our physics productivity
- Useful discussions from the Community White Paper and from our Naples workshop in the spring have reinvigorated work here
 - Recognition of training 'pyramid', from core skills (git, python, ...) to expert (OpenCL, low level optimisation)
 - Healthy interactions with many of the computing schools (Bertinoro, CERN School, GridKA) and with experiments (LHCb StarterKit)
- One near term goal is to offer a training portal to a set of curated and validated materials



Google Summer of Code

- Google sponsored students working on open source projects
 - Running since 2005
 - CERN EP-SFT active since 2011
- HSF acts as an umbrella organisation for many HEP institutes and projects
 - Makes it much easier for our institutes and software projects to get involved
 - Much lighter weight than going to Google directly
 - This year had 29 students accepted for projects (amazing result, >2% of GSoC total)
 - ROOT, Geant4, Rucio, CVMFS, SixTrack, GoHEP, Falcon, YAMPL and many more
- Coding period has just begun

Workshops and events on their way...

- PyHEP workshop at CHEP
 - Review of Python in our community, right from the origins
 - Interactions between ROOT and Python ecosystems
 - Critical point recognised from Amsterdam review and CWP
 - Migration to Python 3
- Frameworks Birds of a Feather
 - Frameworks are like religions, of course!
 - But concurrency and accelerators lay down challenges that are not solved
 - We believe there is the chance, therefore, to work together more closely
- HSF Software Forum (new!)
 - Chance to engage the community in current problems and solution strategies
 - Present new and ongoing software products that should be more widely known
 - E.g., DD4hep geometry package

Working more closely with the EIC Software Consortium and the Nuclear Physics Community would be very welcome!

The HSF Role

- Improve communication in the community
 - We can't work together if we don't know what's going on
 - We have the [hsf-forum](#) mailing list - you are welcome to use it
 - Along with other specialist lists, e.g. the [technical forum list](#)
- Description and inventory of community activities and projects
 - **Everyone** can [put this](#) onto the HSF website
 - New working groups are really welcome
 - HSF can help marshal interested people in the community
- Organise [Community Reviews](#)
 - Great way to get feedback from a dedicated group (GeantV, Analysis Ecosystem)
- Help improve education, training and recognition of software
 - Publish your software, cite the software of others (DOIs); common training initiatives
- New projects should be building cooperation into their core
 - Visit, discuss, **collaborate** - funding agencies really want this
 - HSF can help to advise on how to do this