HEP Software Foundation: R&D and Activities After Naples

Graeme Stewart, CERN EP-SFT

HEP: Landscape and Frontiers

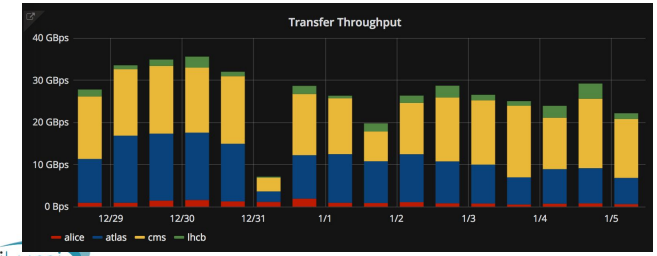

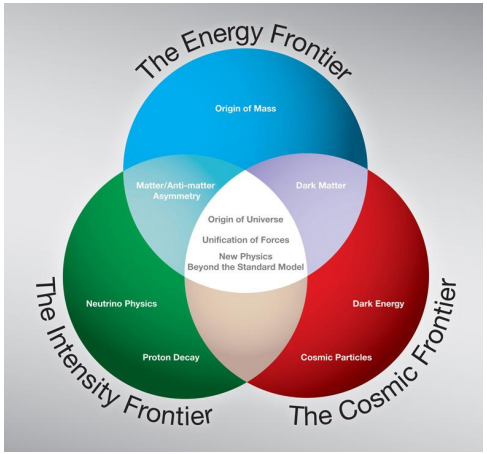

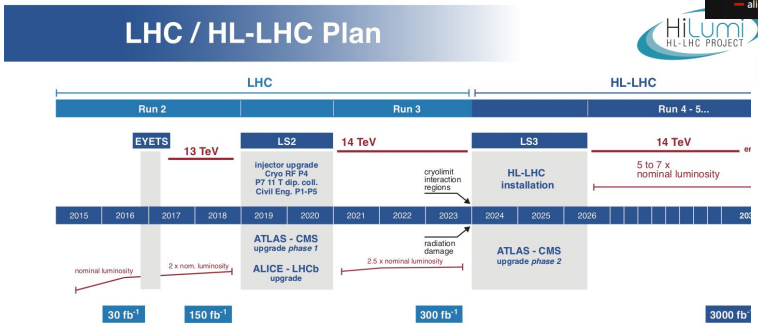

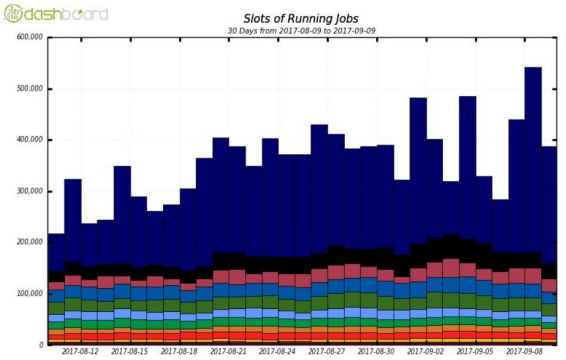

>50M LOC

Fermilab Program Planning
20-Feb-17

LONG-RANGE PLAN: DRAFT Version 7a

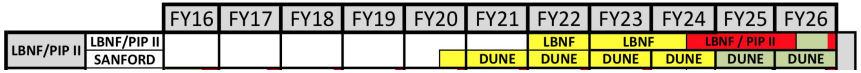

Summer shutdown Construction / commissioning Run

FNAL Intensity Frontier

athena

ATLAS Experiment main repository for Athena code

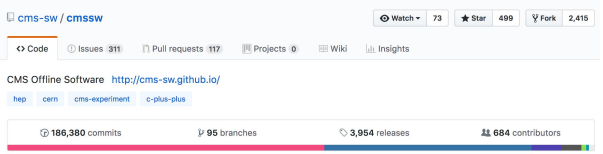

High Luminosity LHC

- Large rise in rate (~10kHz) and complexity (mu~200)

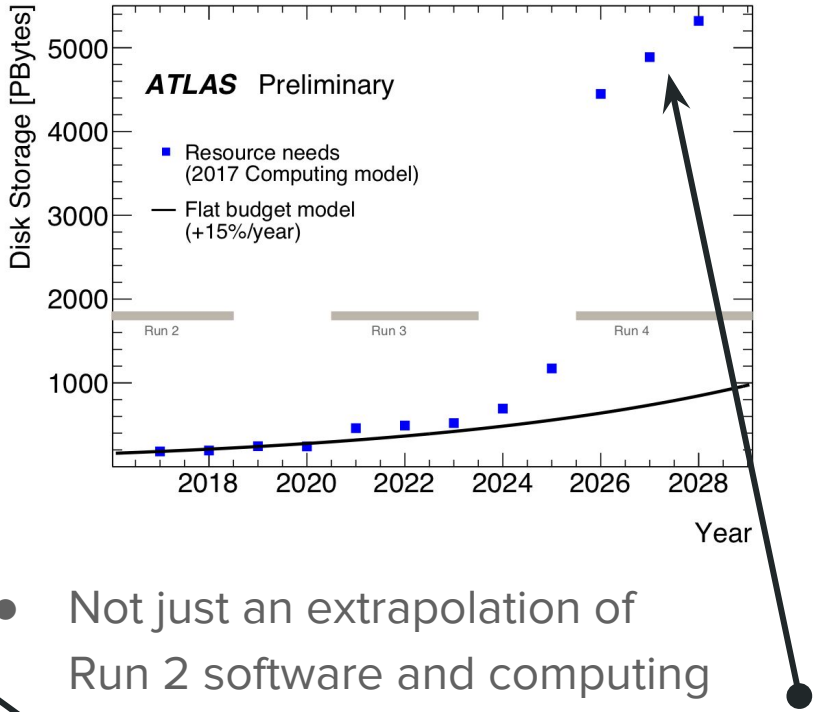

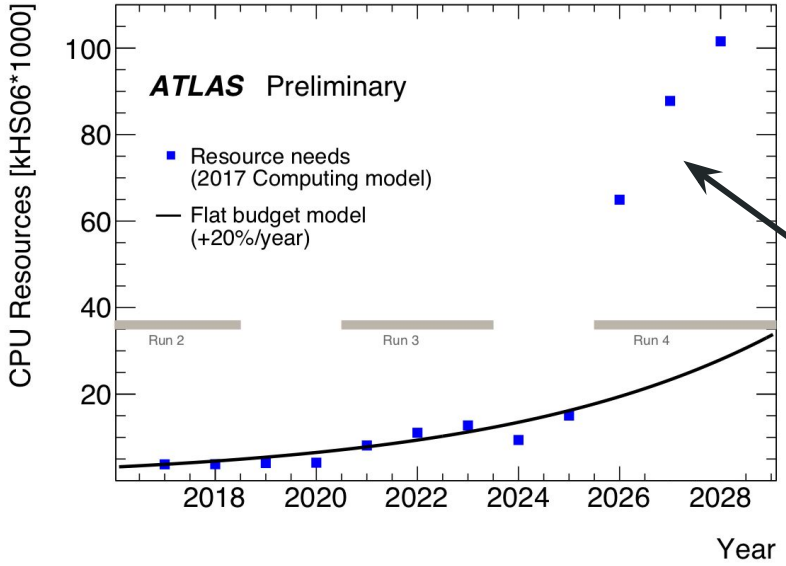

- Not just an extrapolation of Run 2 software and computing
  - Resources needed would hugely exceed those from technology evolution alone

Shifting landscape for end-to-end computing

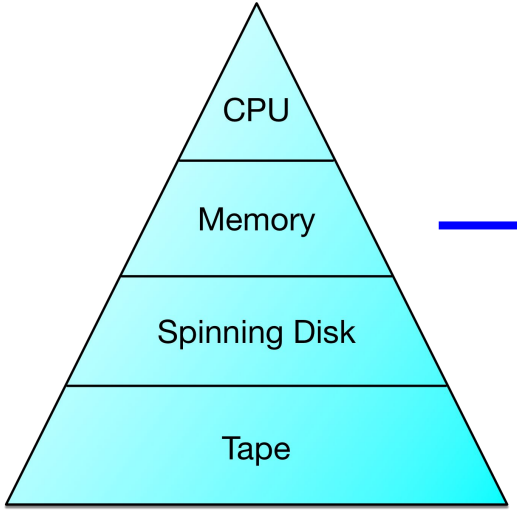

The Good Old Days

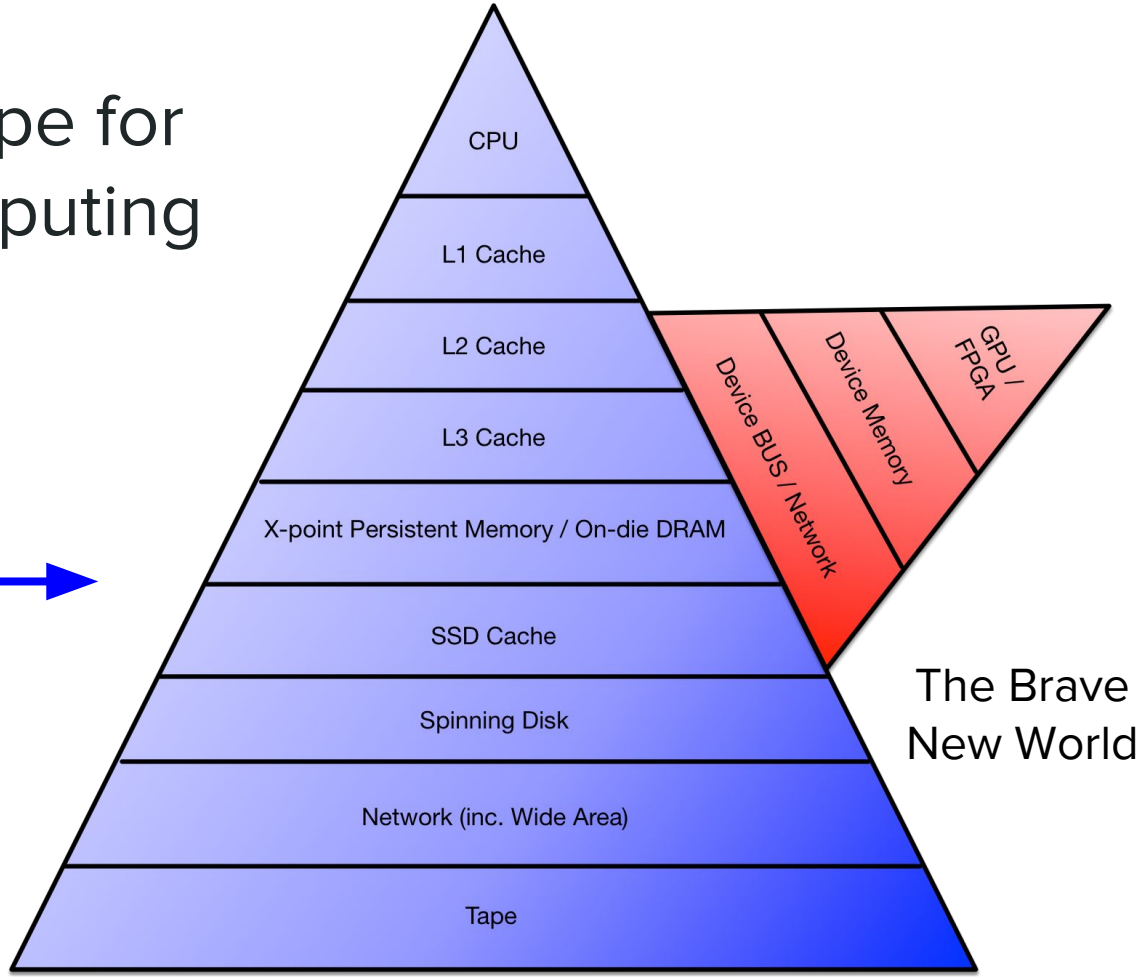

The Brave New World

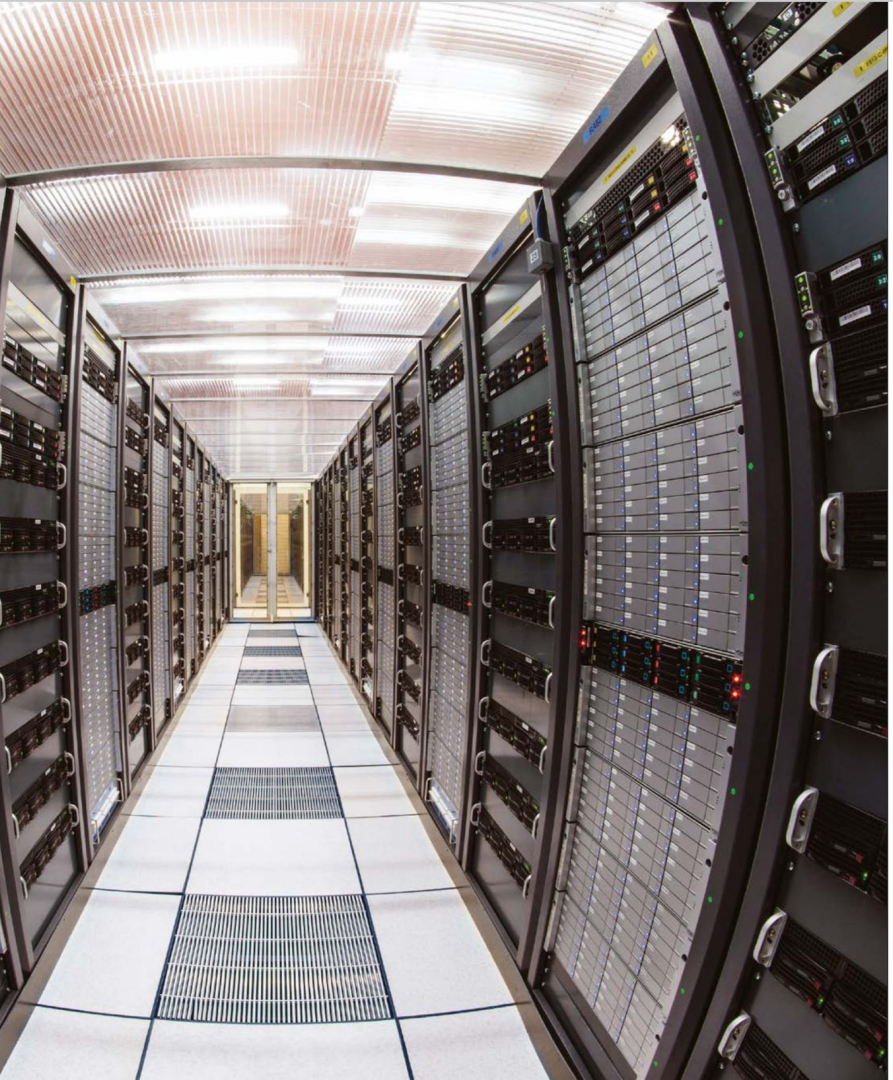

Software and computing

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.)

CWP Roadmap

- 13 chapters, 310 authors, 123 institutes, 1712.06982
  - Thank you!
- Article published in April CERN Courier
- Notable presentations:
  - CERN Scientific Computing Forum
  - LHCC
  - European Committee on Future Accelerators
  - CHEP Plenary

WLCG-HSF Workshop in Naples

- The point of the roadmap was to map out the journey
- At the joint WLCG-HSF Workshop in Naples we were able to progress more concretely in how to do that
  - Started with an overview of the HEP Science Goals (Lix Sexton-Kennedy)
  - Reviewed the whole of the Community White Paper (Eduardo Rodrigues)
  - Followed by many plenary and parallel sessions
    - Technology Watch
    - Frameworks and Concurrency
    - Analysis Facilities
    - Simulation
    - Software Development
- 211 Participants
  - Excellent feedback from attendees

For WLCG specific topics, wait for Simone's talk on Wednesday

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)
  - Adapting to new computing architectures
    - Can a vectorised transport engine actually work in a realistic prototype (GeantV early releases)? How painful would evolution be (re-integration into Geant4)?
  - Faster simulation - develop a common toolkit for tuning and validation of fast simulation
    - How can we best use Machine Learning profitably here? Multi-level approach, from processes to entire events
  - Geometry modelling
    - Easier modelling of complex detectors, targeting new computing architectures

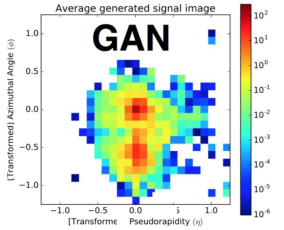

Machine learning simulated calorimiter

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

Software Trigger and Event Reconstruction

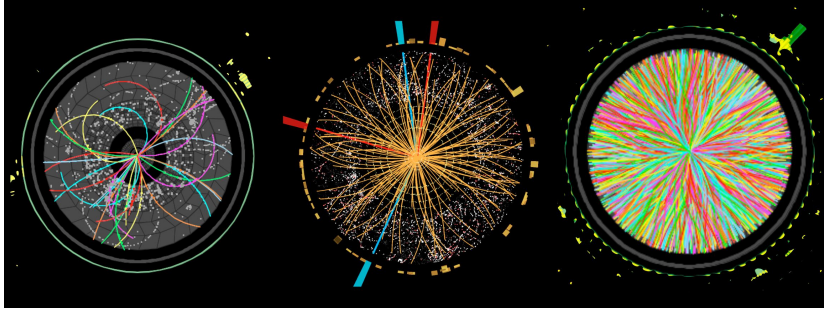

- **Move to software triggers is already a key part of the program for LHCb and ALICE already in Run 3**
  - ‘Real time analysis’ increases signal rates and can make computing more efficient (storage and CPU)
- **Main R&D topics**
  - Controlling charged particle tracking resource consumption and maintaining 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

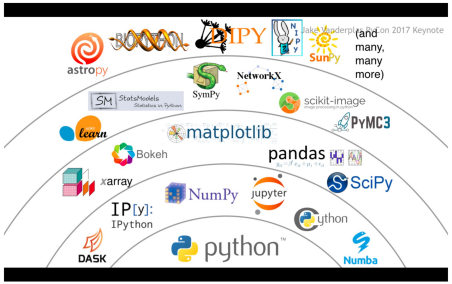

- **Today we are dominated by many cycles of data reduction**
  - Aim is to reduce the input to an analysis down to a manageable quantity that can be cycled over quickly on ~laptop scale resources
  - Key metric is ‘time to insight’
- **Main R&D topics**
  - How to use the latest techniques in data analysis 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

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

New Working Groups

- Intention to form working groups in this three key areas of HEP software:
  - Simulation
  - Reconstruction
  - Analysis
- Building on work in the CWP
- Raise awareness of work being done in these areas
  - Not all projects are as known as they should be
  - New projects can begin with a broad scope and common goals
- These will be areas reviewed by the LHCC next year
  - These groups will able to answer the charge of whether we really have learned to work together or not
- *Please let us know if you're interested!*

Software Forum

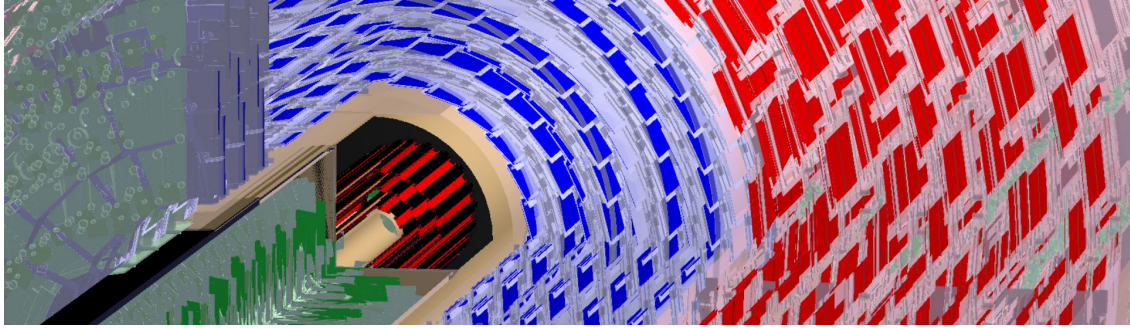

- HSF has relaunched the Software Forum
  - Meetings that can
    - Showcase common software projects
    - Introduce tools that help us face challenges like concurrency or vectorisation
    - Open dialogue with other like-minded communities
- First meeting looked at the DD4hep geometry modeling package
  - Adopted by CLIC, FCC and now CMS; LHCb very interested
- Next meeting (18 July) will look at VecCore and SOAContainer
  - Common libraries addressing vectorisation and data layout
- We would be very happy to hear suggestions for other topics
  - Please use this as a way to connect to the wider HEP software community and different experiments

Copyright and Licensing

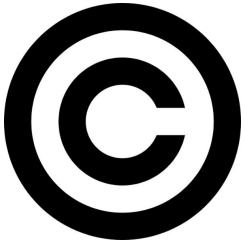

- We continue to work in this much neglected area in HEP software
  - Much code exists with no clear copyright or licence
  - The issues of large and deep stacks of experiments' software and license combinations were often neglected up to now
    - *Does impact on our ability to collaborate*
- Experiments continuing to be more open with their software (LHC, Belle II)
  - Goal is to maximise our useful user base and interactions with others
- 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 (Fastjet in particular)

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
- 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
    - We should be independent of site installed base OS
- R&D Projects looking at some of the directions for the future
  - Nix - pure functional package manager, build everything (really, even `libc`)
  - Portage - from the Gentoo Linux distribution, prefix distribution isolated from system, consistent harmony with `RPATH`
  - Spack - from LLNL, widely used scientific build orchestrator, very multi-version friendly

Other Working Groups

Software Tools Working Group

- Meeting on performance analysis software and how to share data
  - Common work on warehousing and visualisation possible
- Will also look at static analysers and grid tools

Visualisation

- Ed Moyse's WebGL event display now an HSF project (Phoenix)

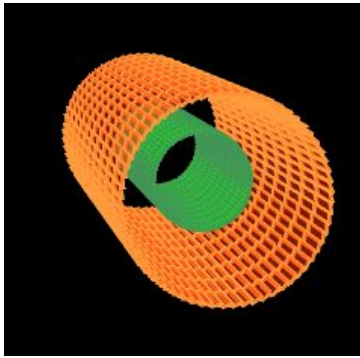

Training

- Recognition of training 'pyramid', from core skills to expert
- Organising a federation of training schools
- Working on a curated set of training materials

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
- Event Generators Re-engineering Workshop
  - Tackle critical issues of software design and performance in event generators
  - Very much needed for HL-LHC
  - Still in planning stages - volunteers welcome!

Conclusions

- CWP and Naples workshop brought us a long way forwards in understanding the problems of the next decade
  - And the areas where we can work together profitably
- HSF continues to act as a focal point for common software efforts
  - Continued work on important technical matters: licensing, packaging, software tools
  - Inventory of software projects and tools; advice on publication and dissemination; training
  - Communication channels (hsf-forum, hsf-tech-forum lists) are vital
- New working groups will form nuclei of solving the grand challenges for HL-LHC
- New software projects will come
  - The should be agile and cooperative from the outset

There are many opportunities to be involved and shape our common work in the field