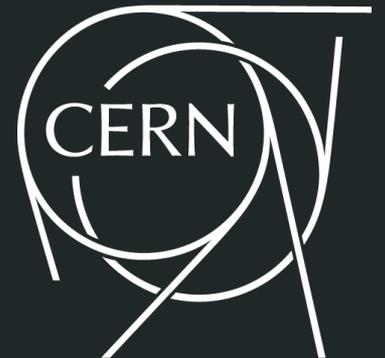




# HEP Software Foundation Community White Paper

Graeme Stewart, CERN EP-SFT

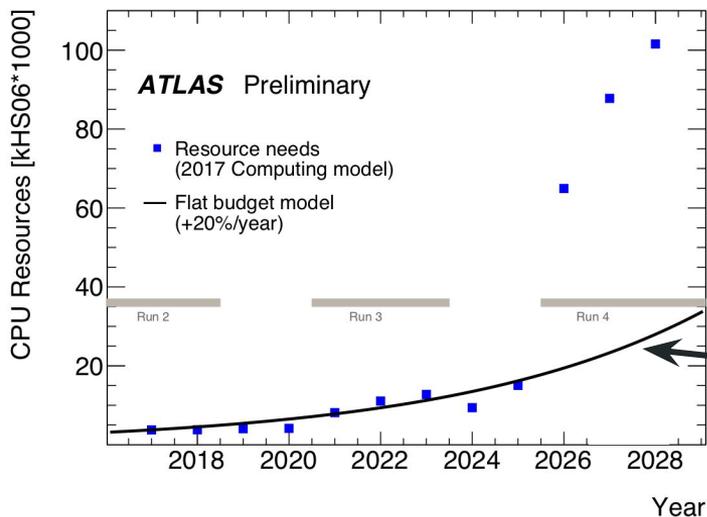


GridPP40, Pitlochry

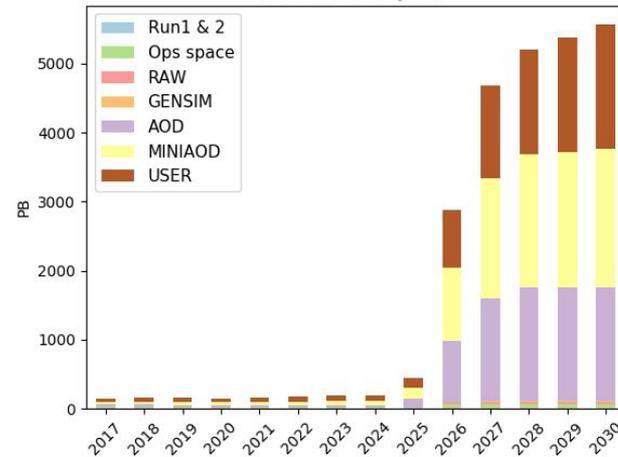


# Software and Computing Challenges

- HL-LHC (only one part of the HEP programme) brings a huge challenge to software and computing
  - Both event rate and complexity rise



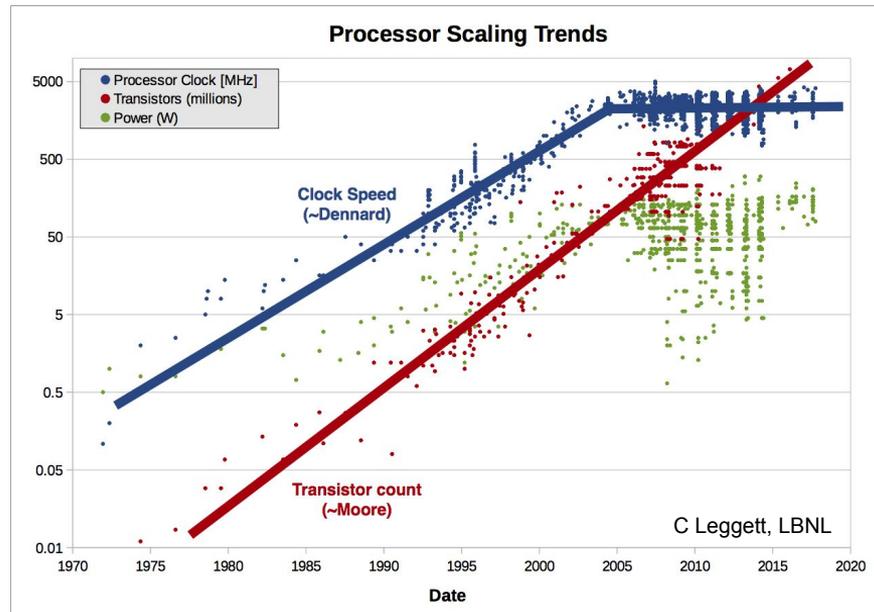
CMS  
Data on disk by tier



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

# Processor evolution

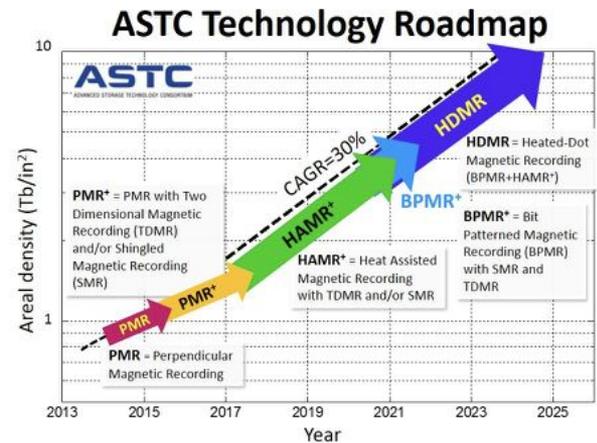
- Moore's Law continues to deliver increases in transistor density
  - Albeit, doubling time is lengthening
  - Much flim-flam spoken about process size
- Clock speed increases stopped around 2006
  - No longer possible to ramp the clock speed as process size shrinks (Dennard scaling failed)
- So we are stuck at ~3GHz clocks from the underlying  $\text{Wm}^{-2}$  limit (power wall)
  - Limits the capabilities of serial processing
  - CPU based concurrency still in development for Run 3



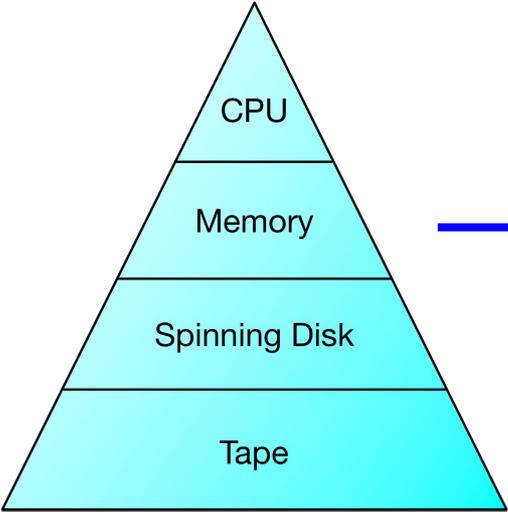
- Meanwhile accelerators continue to deliver huge amounts of FLOPS 'on paper'
  - Much larger fraction of die used for maths
  - But not easy to use (R.I.P. Xeon Phi)
- Our code is not in good shape to take advantage of this at the moment

# Other Technology Trends

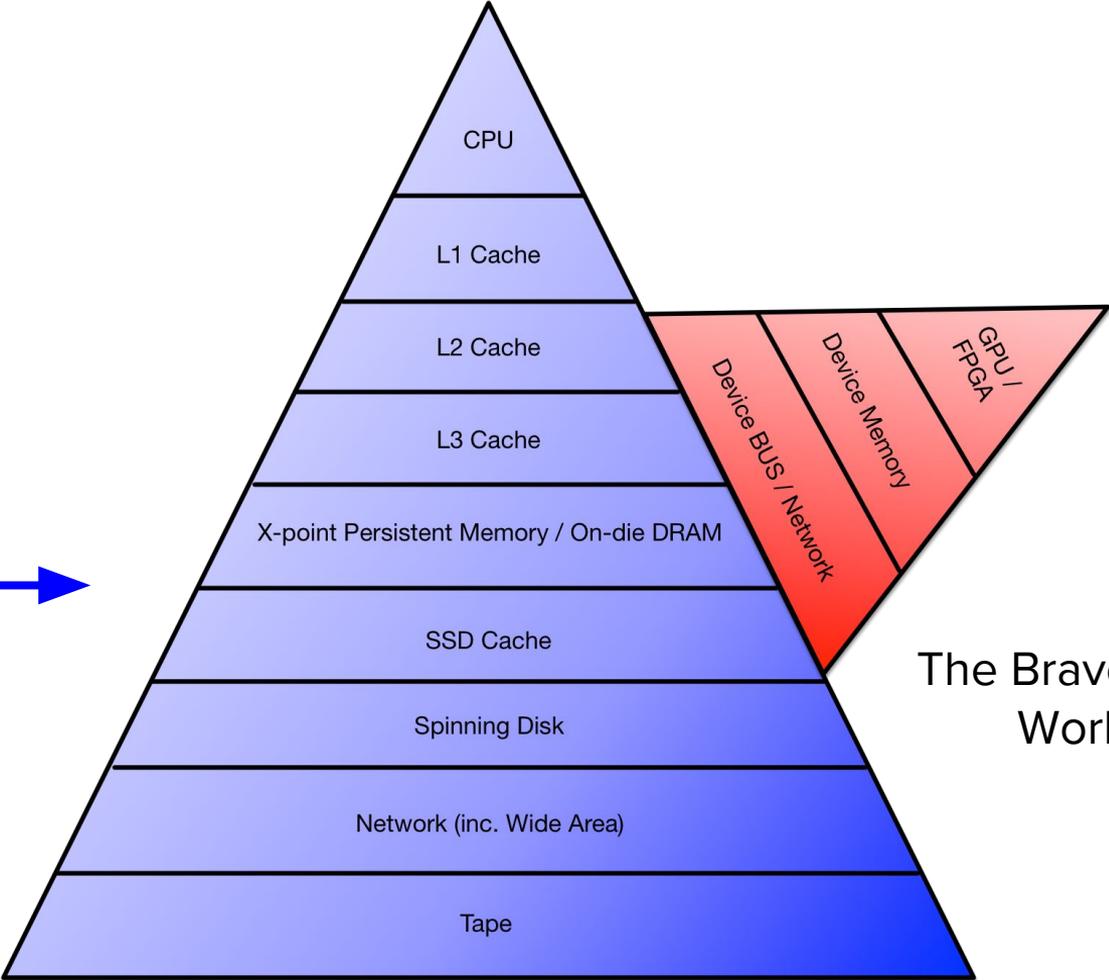
- Memory
  - DRAM improvements now modest
  - Overall, memory 'landscape' becomes more complex
    - Memory/storage boundary blurring
- Storage
  - Spinning disk capacity keeps climbing
    - Time to read and cost improves, but slowly
  - SSDs can read much faster, but price remains too high for bulk storage
  - Tape remains cheap to buy, slow to access with few companies left, O(1)
- Networks
  - Capacity increases expected to continue, latency will not change
  - Next generation networks offer capability to open channels between sites on demand
    - Useful, but an additional complexity



# Meaning...



The Good Old Days



The Brave New World

# HEP Software Foundation (HSF)

- The LHC experiments, Belle II and DUNE face the same challenges
  - HEP software must evolve to meet these challenges
  - Need to exploit all the expertise available, inside and outside our community, for parallelisation
  - 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, ...)
  - We already have to support 50M+ lines of code for current experiments
    - 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*

# 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

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

- 70 page [document](#)
- 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
- 307 authors from 123 institutions
- Signing policy: sign the document if you agree with the main observations and conclusions
  - [hsf-cwp-ghost-writers@googlegroups.com](mailto:hsf-cwp-ghost-writers@googlegroups.com)
- *We really actively encourage you to do this as your name indicates the breadth of support in the community - signing will close very soon in advance of journal publication*

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Software and Computing Challenges</b>	<b>5</b>
<b>3</b>	<b>Programme of Work</b>	<b>11</b>
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
<b>4</b>	<b>Training and Careers</b>	<b>65</b>
4.1	Training Challenges	65
4.2	Possible Directions for Training	66
4.3	Career Support and Recognition	68
<b>5</b>	<b>Conclusions</b>	<b>68</b>
	<b>Appendix A List of Workshops</b>	<b>71</b>
	<b>Appendix B Glossary</b>	<b>73</b>
	<b>References</b>	<b>79</b>

# 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.)*

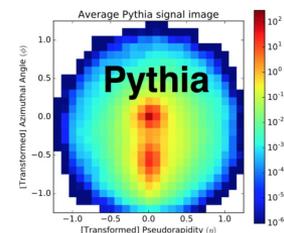
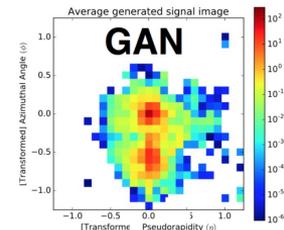
- Attracted considerable attention
- Community White Paper article published in April [CERN Courier](#)
- Presentations to CERN Scientific Computing Forum

# WLCG-HSF Workshop in Naples

- The point of the roadmap was to map out the journey
- Joint WLCG-HSF Workshop in Naples was 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
- 211 Participants
  - Excellent feedback from attendees
- Reports and follow-up are well worth following
  - Last week's HSF meeting
  - Today's GDB

# Detector Simulation

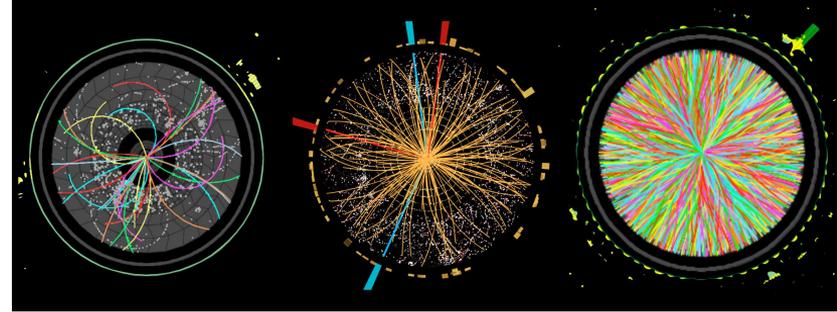
- **Simulating our detectors consumes huge resources today**
  - Remains a vital area for HL-LHC and intensity frontier experiments in particular
- **Main R&D topics**
  - **Improved physics models** for higher precision at higher energies (HL-LHC and then FCC)
    - Hadronic physics in LAr TPCs needs 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)

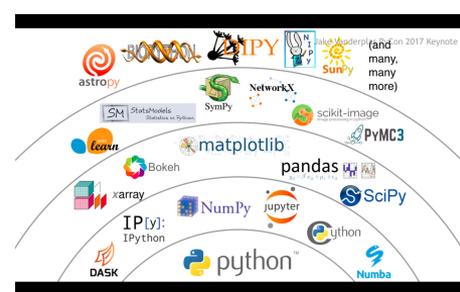
# Data Analysis and Interpretation

- **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

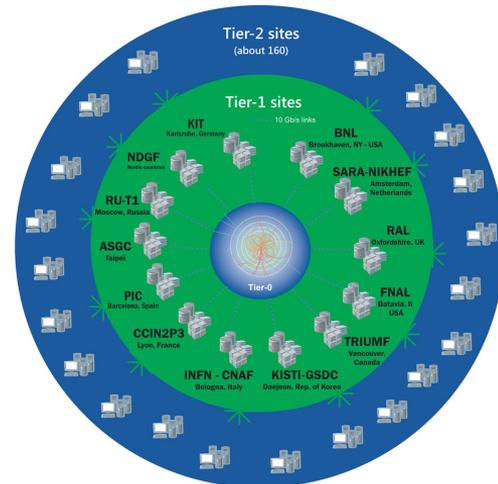
- **Main R&D topics**

- 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

# Facilities and Distributed Computing

- **Storage and computing today are provided overwhelmingly from WLCG resources**
  - Expected to continue for HL-LHC, but to be strongly influenced by developments in commodity infrastructure as a service (IaaS, commercially this is usually Cloud Computing)
- **Main R&D topics**
  - Understand far better the **effective costs** involved in delivering computing for HEP
    - This needs to be sensitive to **regional variations in funding** and **direct and indirect costs**
    - Providing a reasonable gradient analysis for future investment
  - How to reorganise arrangement of storage and compute for improved efficiency
  - Hung off the term **data lake** (buzzword - I would not take that too seriously), key questions are:
    - Should we concentrate storage at fewer endpoints? Ensure experiment/site interaction simplicity and preserve efficiency; reduce human effort more than hardware?
    - How to take advantage of **new network and storage technologies** (software defined networks, object stores or content addressable networks)
  - Strengthen **links to other big data sciences** (SKA) and computing science; how to share network resources



R&D Outlook: Real effort here, needs metrics that work with real world complexity

# Guiding Strategy

- 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
  - As illustrated in CWP chapters, not one single approach for every topic: several paths for moving in this direction are part of the roadmap



# 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 - do please 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

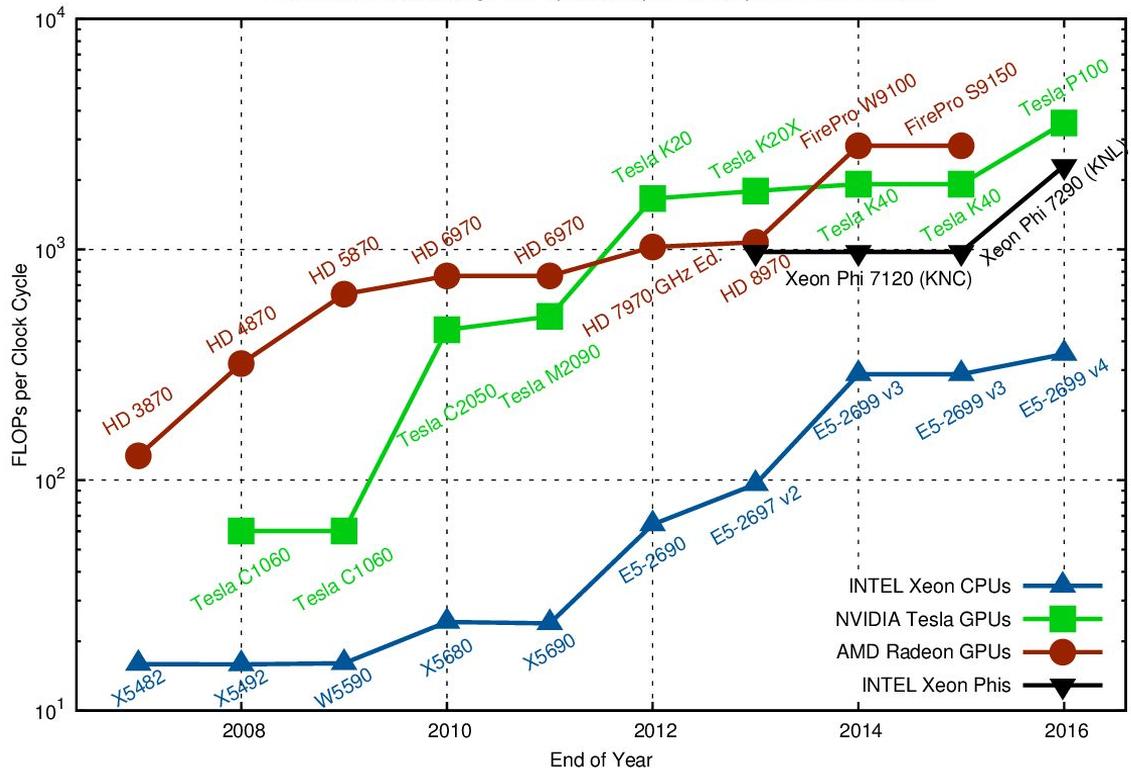
# Advancing from here

- **Community White Paper** process is **concluded** has been a **success**
  - Engaged more than 300 people and produced more than 300 pages of detailed description in many areas
  - Summary roadmap lays out a path forward and identifies the main areas we need to **invest** in for the future for our **software upgrade**
  - Supporting the HL-LHC Computing TDRs, NSF S2I2 strategic plan, CERN EP R&D, etc.
- **Naples workshop** helped define the programme of R&D which follows this
  - A lot was advanced, but **not evenly** or completely in all areas - **follow up required**
- **We need to marshal the R&D efforts in the community, refocusing our current effort and helping to attract new investment in critical areas**
  - The challenges are formidable, working together will be the most efficacious way to succeed
  - HSF will play a vital role in **spreading knowledge** of new initiatives, **encouraging collaboration** and **monitoring progress**
    - Working with partner organisations like WLCG

# Backup

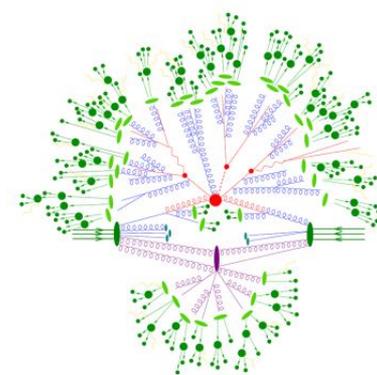
# Accelerator performance

Theoretical Peak Floating Point Operations per Clock Cycle, Double Precision



This is not a potential we even reach

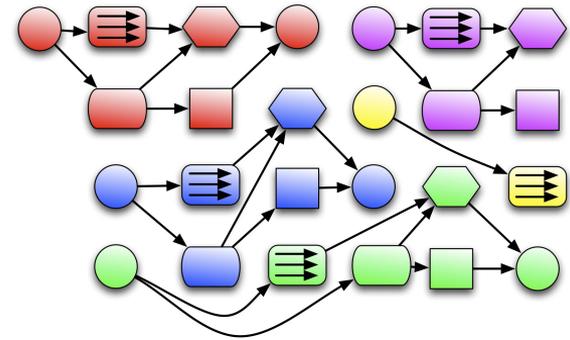
# Physics Event Generators



- Physics event generation starts our simulation chain to enable comparisons with detector events
  - Depending on the precision requested, CPU for event generation ranges from modest to huge
  - At Next-to-Leading Order (NLO) precision used today, CPU consumption can become important
  - Study of rare processes at the HL-LHC will require the more demanding NNLO for more analyses
- Generators are written by the theory community
  - Need expert help and long-term associations to achieve code optimisation
  - Even basic multi-thread safety is problematic for many older, but still heavily used, generators
  - Ongoing maintenance of tools like HepMC, LHAPDF, Rivet is required and needs rewarded
- Writing this section was the result of intense contacts between HEP experts and the main players in the generator community
  - Several actions identified to foster the collaboration

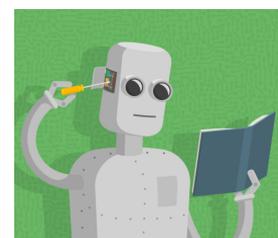
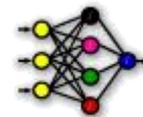
# Data Processing Frameworks

- Experiment **software frameworks provide the scaffolding for algorithmic code**
  - Currently there are many implementations of frameworks, with some (limited) sharing between experiments (e.g. ATLAS/LHCb Gaudi)
  - Ongoing efforts in all these frameworks to support concurrency
- Main R&D topics
  - **Adaptation to new hardware**, optimising efficiency and throughput
  - Incorporation of external **(co)processing resources**, such as GPGPUs
  - **Interface with workload management** to deal with the inhomogeneity of processing resources
- General agreement that it is an area for consolidation in the future
  - Reasons for so many frameworks are not really related to experiment specificities...
  - But also the hardest component to change in the experiment SW stack: need to identify the best approaches to promote commonalities, e.g., in underlying components



Event processing framework schematic (colours are events, boxes algorithms)

# Machine Learning



- Not a challenge for HEP, per se, rather a tool set providing possible solutions
  - ML techniques applicable to many different areas
- Neural networks and Boosted Decision Trees have been used in HEP for a long time
  - e.g., particle identification algorithms
- The field has been significantly enhanced by new techniques (Deep Neural Networks), enhanced training methods, and community-supported (Python) packages
  - Very good at dealing with noisy data and huge parameter spaces
  - A lot of interest from our community in these new techniques, in multiple fields
- Main R&D topics
  - **Speeding up** computationally intensive pieces of our workflows (fast simulation, tracking)
  - **Enhancing physics reach** by classifying better than our current techniques
  - Improving **data compression** by learning and retaining only salient features
  - **Anomaly detection** for detector and computing operations
- **Good links with** the broader **Machine Learning and Data Science communities** required
  - Required to match the efforts and expertises needed to make effective use of these techniques

# Other technical areas of work

## Conditions Data

- Growth of alignment and calibration data is usually linear in time
  - Per se, this does not represent a major problem for the HL-LHC
- Opportunities to use modern distributed techniques to solve this problem efficiently and scalably
  - Cacheable blobs accessed via REST
  - CVMFS + Files
  - Git

## Visualisation

- Many software products developed for event visualisation
  - Part of the framework, with full access to event and geometry data
  - Standalone as a lightweight solution
- New technologies for rendering displays exist, e.g., WebGL from within a browser

- These areas are examples of where we can refocus current effort towards common software solutions
- This should improve quality, economise overall effort and help us to adapt to new circumstances

# Security Matters



- We have a large infrastructure that is an important resource for us
  - Protecting it is necessary for both our work and for our reputation
- **Trust and policy**
  - Evolve away from “in HEP house” to modern data exchange
  - Support integration of commercial resources and hybrid clouds
  - Need to comply with new legislation, e.g., EU Data Protection
- **Operational Security**
  - Better intelligence sharing for threat monitoring and response
  - Broaden regional participation
  - Coordinate with other communities, R&E Forum for incident response
- **Authentication and Authorisation**
  - Generic authentication services (e.g., eduGAIN) help users and are easier than X.509
  - Authorisation still needs to be in HEP control (VO management)

# Data, software and analysis preservation



- We seem to be doing well compared to other fields
- Challenge is both to physically preserve bits and to preserve knowledge
  - [DPHEP](#) has looked into both
- Knowledge preservation is very challenging
  - Experiment production workflows vary in significant details
  - Variety of different steps are undertaken at the analysis stage, even within experiments
- Need a workflow that can capture this complexity
  - Technology developments that can help are, e.g., containers
- CERN [Analysis Preservation Portal](#) forms a good basis for further work
  - Needs to have a low barrier for entry for analysts
  - Can provide an immediate benefit in knowledge transmission within an experiment

# Software development, training and careers

- **Experiments have modernised their software development models a lot recently**
  - Moving to git and CMake as standard components
  - Using social coding sites (gitlab, github) coupled to Continuous Integration
  - **Additional tools would benefit the community:** Static analysis of code, refactoring code, performance measures, **re-establishing the development forum** (*a la* Concurrency Forum)
- **Using new tools requires investing in training for the community**
  - The more commonality in the tools and techniques, the more training we can share
    - ALICE and LHCb recently did this in practice using the StarterKit material
  - This provides preservation and propagation of knowledge
- **Our environment is becoming more complex; we require input from physicists *whose concerns are not primarily in software***
  - **Sustainability** of these contributions is extremely important
- **Recognition of the contribution of our specialists in their careers is extremely important**
  - We should become better at publication and citation of work to help this (and use new tools like [Zenodo](#))