

Computing for HEP: challenges & R&D

José M. Hernández
CIEMAT, Madrid

XLIX International Meeting on Fundamental Physics, Benasque, 6-9 September 2022



Ciemat Centro de Investigaciones
Energéticas, Medioambientales
y Tecnológicas




Outline

- The role of (offline) computing in High Energy Physics experiments
- The Worldwide LHC Computing Grid infrastructure
 - Past, present and future
 - WLCG in Spain
 - The Computing challenge for the High Luminosity phase of the LHC
- HEP Computing R&D
 - Hardware and software
 - HEP Computing Roadmap for the 2020's
- Leverage of LHC Computing infrastructure in other HEP projects
- Outlook

Computing in High Energy Physics

- The computing system (hardware and software) is an essential element in the instrumentation and scientific exploitation of HEP experiments
- The sheer volume and complexity of data in HEP experiments requires complex data acquisition, processing, simulation and analysis



```
patchInt32Value, virtualServiceGVR = schema.GroupVersionResource{
    Group: "networking.k8s.io",
    Version: "v1beta3",
    Resource: "virtualservices",
}
// weight the two routes = 50/50
patchPayload := make([]patchInt32Value, 2)
patchPayload[0].Op = "replace"
patchPayload[0].Path = "/spec/http/0/route/0/weight"
patchPayload[0].Value = 50
patchPayload[1].Op = "replace"
patchPayload[1].Path = "/spec/http/0/route/1/weight"
patchPayload[1].Value = 50
```

analysis

High Energy Physics data-intensive computing

- HEP computing driven by **large-scale data flow + volume**
 - Scale of data from 10s of TB to 100s of PB/year
- Data-intensive applications need **performance, reliability, and low latency**
- Overall balance of compute + I/O + storage + networking needs to be carefully designed
 - Multiple IO requirements, e.g. high I/O workflows: 10-100 Mbit/s/core
- Large variety of workflows
 - Data calibration, reconstruction, simulation
 - Data reduction (skimming, slimming), data analysis
- Data pipelines can be complex and need to be run many times
 - Individual campaigns can last for months
- Most experimental data requires fine-grained analysis
 - Hundreds of analysis users using resources in a “chaotic” way



A key tool for physics

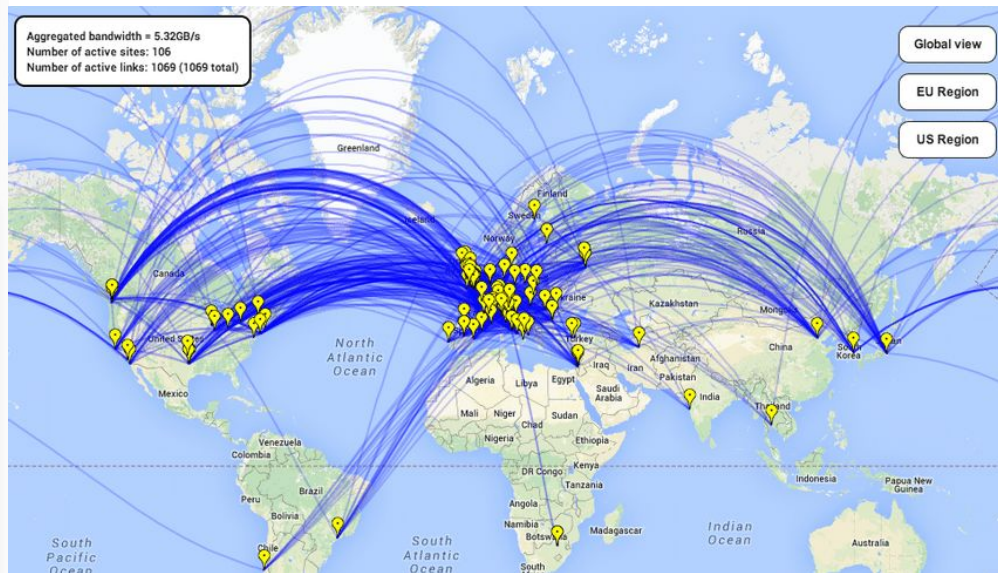
The most sophisticated data-taking & analysis system ever built for science, providing near real-time access to LHC data.



Seamless access

Computing resources which include data storage capacity, processing power, sensors, visualization tools and more.

The Worldwide LHC Computing Grid (WLCG)



Global collaboration

42 countries
170 computing centres
Over 1 million computer cores
2 exabytes of storage



Enabling discovery

WLCG computing enabled physicists to announce the discovery of the Higgs Boson on 4 July 2012.


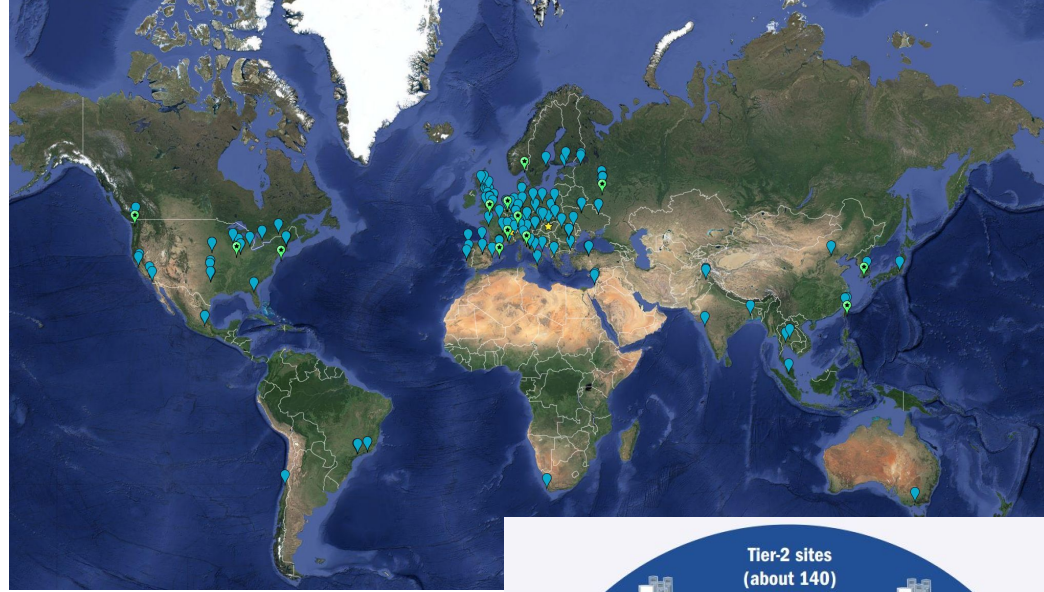
WLCG: design principles and enabling technologies

- Computing infrastructure geographically **distributed** and **federated**
 - Glued through common software interfaces
 - High availability (24x7 in big centers) and reliability (>95%)
- **Tiered** structure
 - Tier-0: storage of raw data, prompt data processing (calibration, reconstruction)
 - Tier-1: custodial archival of data, organized data processing
 - Tier-2: data simulation and analysis
- Sites interconnected with **low latency and high bandwidth networks**
 - Optical private network (LHCOPN) interconnecting Tier-0 and Tier-1 sites
 - LHC Open Network Environment (LHCONE) linking Tier 2 sites
- **Seamless** access through specific software services
 - **Authentication and authorization** system for secure access to services
 - **Data management**: data transfers, cataloguing, access
 - **Workflow management**: task orchestration (data processing, simulation and analysis), job execution, monitoring

Worldwide LHC Computing Grid

Distributed high-throughput computing infrastructure to store, process & analyze data produced by LHC experiments

- 167 sites, 42 countries, 63 MoU's
- ~1 million CPU cores
- ~750 PB disk storage
- ~1250 PB tape storage
- Optical private network (LHCOPN) and overlay over NRENs (LHCONE) with 10/100 Gbps links
- ~Tbps LAN bandwidth between compute and storage nodes at sites



Major features and capabilities of HEP computing infrastructure

- Networks
 - International and national, private and public
- Data management
 - Key to success, data transfers, storage systems, data management tools and data organization
- Compute
 - Provision of resources and workload scheduling, execution and monitoring
- Authentication and authorization
 - The mechanism of federation, single sign on, etc
- Operations support
 - Security, incidence response, problem tracking, daily operations, upgrade campaigns
- Diverse experiment-specific services and tools, applications

Distributed data-intensive high throughput computing (HTC)

Precursor of Big Data processing and Cloud computing

Hardware technologies in WLCG


CPU: x86 processors
(Intel, AMD)
~1 million CPU cores



Disk: hard disk drives
(up to 20 TB/disk)
~750 petabytes



Tape: cartridges
(up to 18 TB/tape)
~1250 petabytes



Hardware technologies in WLCG

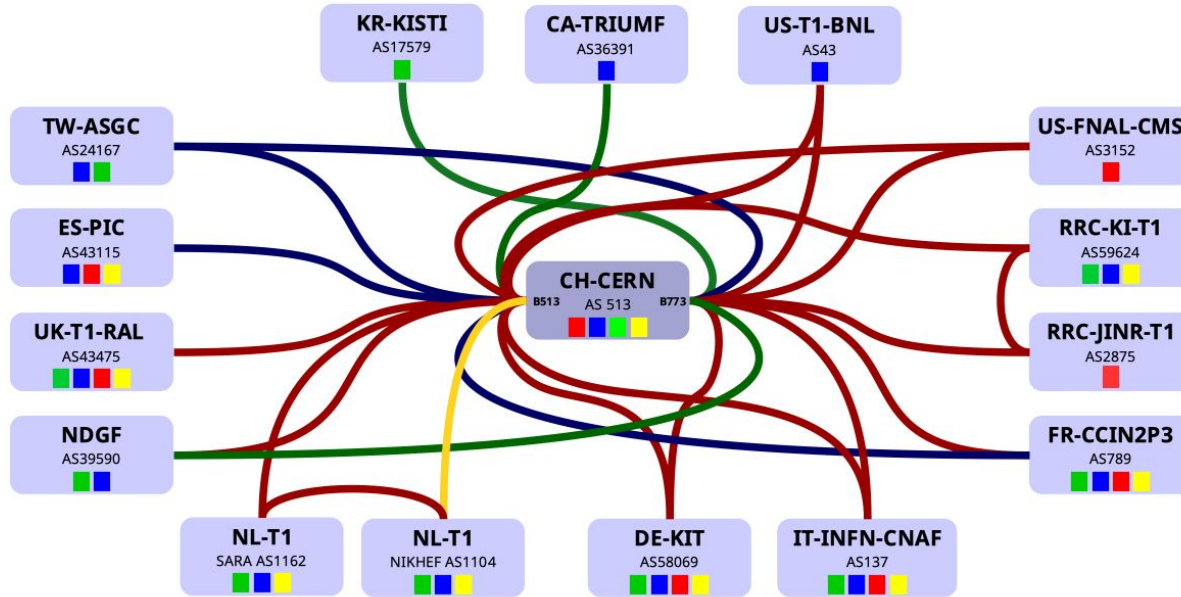
Network:

Ethernet switches interconnecting compute and storage hardware at a site (LAN)

Ethernet routers interconnecting sites (Wide Area Network)



LHCOPN



■ = Alice ■ = Atlas ■ = CMS ■ = LHCb
— 10Gbps
— 20Gbps
— 100Gbps
— 400Gbps
 edoardo.martelli@cern.ch 20220711

Optical private Ethernet networks T0 ↔ T1s

Typically 100 Gbps WAN links

Over infrastructure provided by national research and academic network providers (e.g. Geant, RedIris)

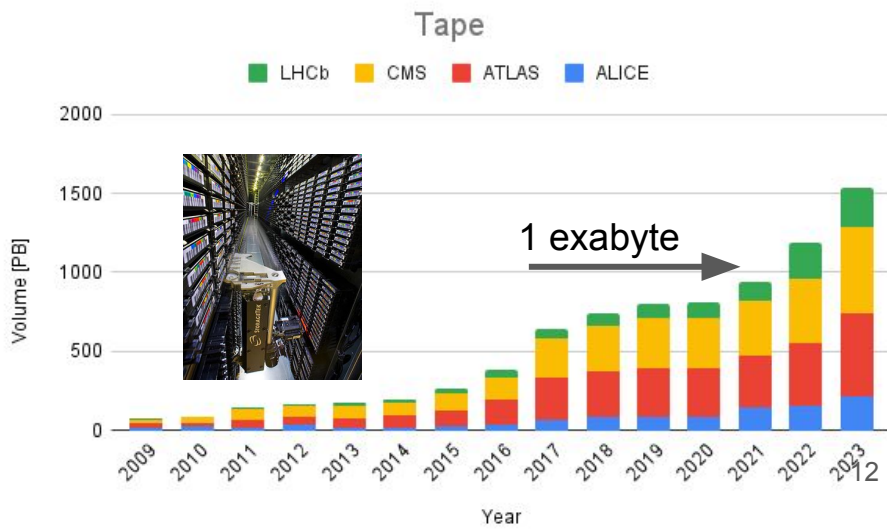


WLCG Computing resource evolution

Countries pledge resources annually according to the experiment needs



Currently:

~1M CPU cores, ~2 exabyte storage

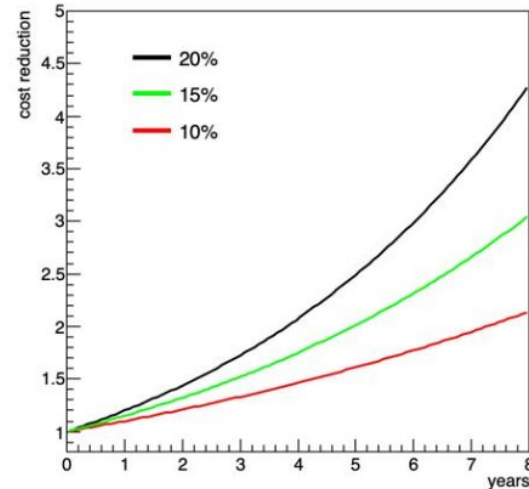
~20% annual growth



Hardware cost evolution



- Growth based on assumption of “flat budget”
 - More hardware with same money profiting from decreasing hardware prices
 - 10-20% price reductions in the past
- Unclear if trend will continue
 - Large impact



WLCG Memorandum of Understanding in Spain

The European Organization for Nuclear Research (CERN)

and

The Spanish Ministry of Education and Science (MEC)

declare that they agree on this Memorandum of Understanding for collaboration in the deployment and exploitation of the Worldwide LHC Computing Grid.

Done in Geneva

on July 31st, 2007

Done in Madrid

on 6 July 2007

in two originals in the English and Spanish languages, each version being equally authentic, it being understood however that in case of contradiction, ambiguity or differences in interpretation, the English text shall prevail.

For CERN



Jos Engelen
Chief Scientific Officer

For MEC



Miguel Angel Quintanilla Fisac
(Real Decreto 567/2006, BOE 108, 06/05/2006)

- Signed by Spain in 2007
- Contribute with 5% of computing resources at the T1 and T2 levels for the ATLAS, CMS and LHCb experiments

Tier1

Centre	Experiments served with priority				Representative to WLCG Collaboration	Funding Agencies
	ALICE	ATLAS	CMS	LHCb		
Canada, TRIUMF		X			M. Vetterli	CFI
France, CC-IN2P3	X	X	X	X	F. Malek (deputy: F. Hernandez)	CNRS/IN2P3 and CEA/DSM/DAPNIA
Germany, FZK-GridKA	X	X	X	X	K.-P. Mickel	BMBF/FZK
Italy, CNAF	X	X	X	X	M. Mazzucato (deputy: L. Dell'Agnello)	INFN
Netherlands LHC/Tier1	X	X		X	J. Templon	NIKHEF
Nordic Data Grid Facility (NDGF)	X	X	X		L. Fisher	NDGF
Spain, PIC		X	X	X	M. Delfino (deputy: G. Merino)	MEC
Taipei, ASGC		X	X		S. Lin	Academia Sinica
UK, RAL	X	X	X	X	N. Geddes	PPARC
USA, BNL		X			M. Ernst (alt: R. Popescu)	DOE
USA, FNAL			X		V. White	DOE

Annex 2. WLCG Tier2 Centres and Federations of Centres that together constitute a Tier2 Centre

Institution	Experiments served with priority				Representative to WLCG Collaboration	Funding Agencies
	ALICE	ATLAS	CMS	LHCb		
Spain, ATLAS Federation - IFAE, Barcelona - IFIC, Valencia - UAM, Madrid		X			J. Salt (alt: A. Pacheco Pages, J. del Peso)	MEC
Spain, CMS Federation - CIEMAT, Madrid - IFCA, Santander			X		F. Matorras (alt: N. Colino)	MEC
Spain, LHCb Federation - UB, Barcelona - USC, Santiago				X	R. Graziani Diaz (alt: J.J. Saborido Silva)	MEC

Spain in WLCG




Spanish contribution:

- ~5% resources T1 & T2 (MoU)
- 1 Tier-1 center (PIC, CIEMAT-IFAE)
- 6 Tier-2 centers
 - CMS federation: CIEMAT-Madrid & IFCA Santander
 - ATLAS federation: IFIC-Valencia, IFAE-Barcelona, UAM-Madrid
 - LHCb federation: USC-Santiago, UB-Barcelona



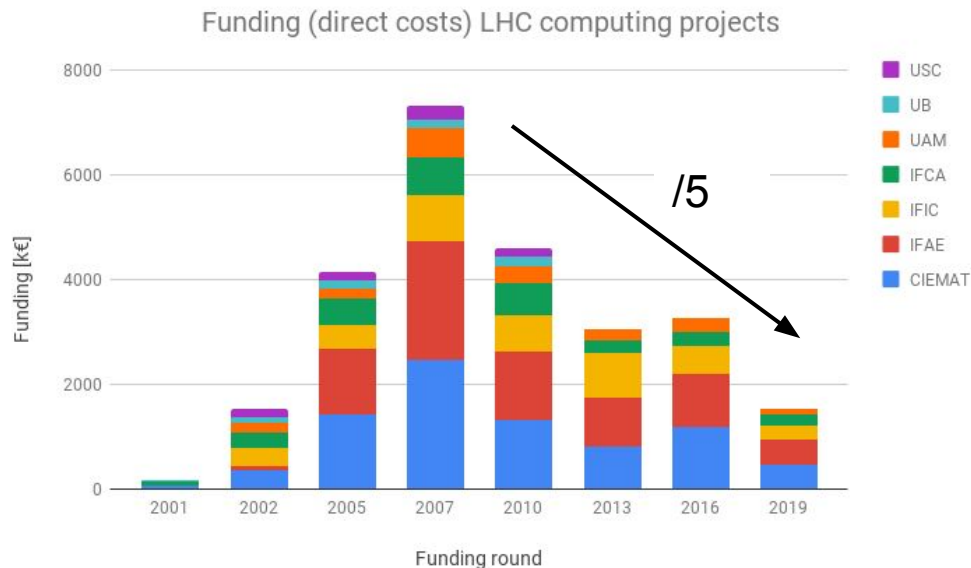
WLCG-Spain: a success story

- ~Two decades contributing to WLCG at high level
 - ~5% WLCG resources, ~1500M CPU hours delivered
 - Providing 1 of the 13 Tier-1 sites worldwide (PIC)
 - Federated Tier-2 sites for ATLAS (IFIC, IFAE, UAM), CMS (CIEMAT, IFCA), LHCb (USC, UB)
 - Among the most reliable sites in WLCG
- A large effort from HEP community and institutions
 - ~24 M€ funding from HEP national plan since 2001
 - Funding from institutions of the same order
 - Funding personnel, electricity, infrastructure
- Large community of experts in distributed high throughput computing
 - Contributions to LHC computing, development, integration, operations, management
- **Big strategic asset for Spanish HEP community!**



WLCG-ES sustainability challenge

- Decreasing funding from HEP national programme
 - From ~2.5M€/year in 2007 to ~0.5 M€/year in 2019
 - Contribution reduced from 5% to 4%
 - Aging equipment (~50% > 5 years)
- Big effort to complement funding and resources
 - From institutions
 - From national/regional scientific infrastructure calls
- Required funding ~1.5M€/year



WLCG's success

- $\sim 50 \cdot 10^9$ CPU-hours delivered
- ~ 2 exabyte of experimental and simulated data stored

Global Effort → Global Success

July 4, 2012


Results today only possible due to extraordinary performance of accelerators – experiments – Grid computing

Observation of a new particle consistent with a Higgs Boson (but which one...?)

Historic Milestone but only the beginning

Global Implications for the future

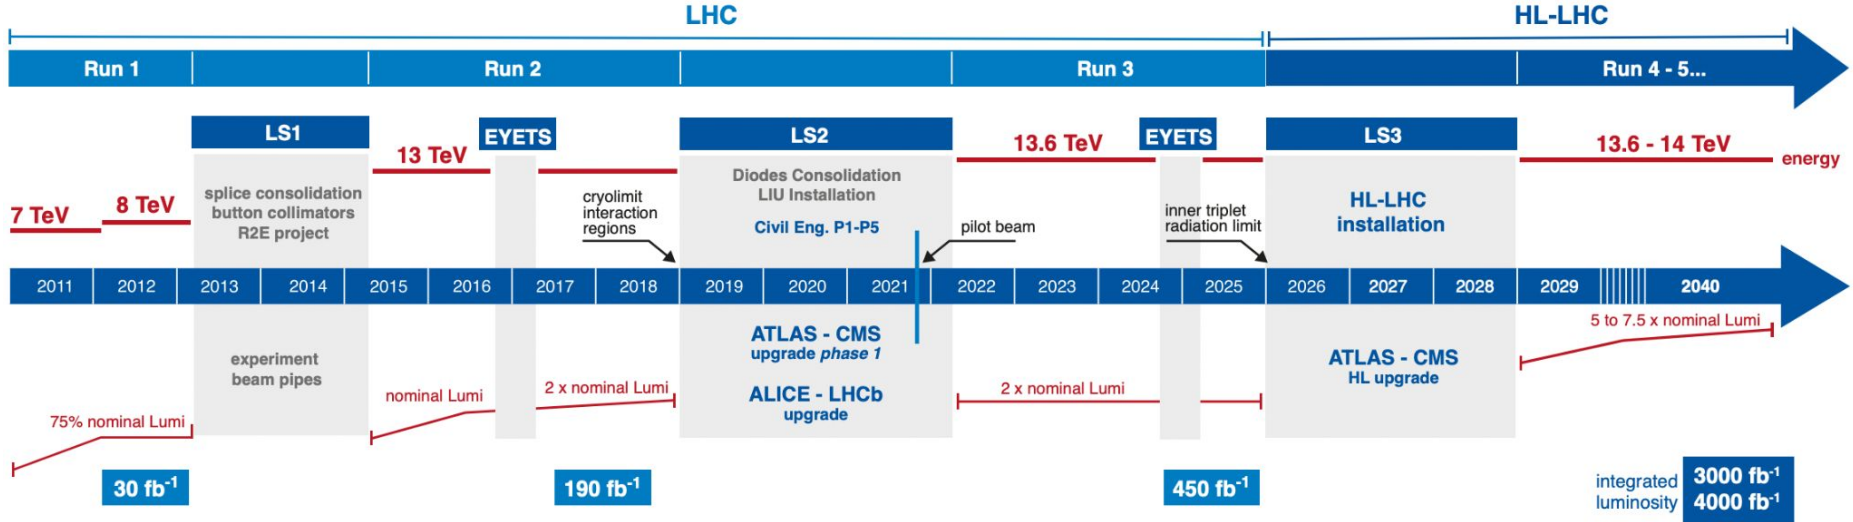
R-D Heuer



WLCG: Lessons learnt

- The network is a fundamental resource and opportunity
 - Rely even more on the network to add flexibility and reduce resource needs
- Distributed data management and storage is expensive
 - hardware and operations
 - Data pre-placement is very complex. Remote data access can be inefficient
- A trusted federated infrastructure is of tremendous value and importance
 - Although X.509 authorization/authentication model difficult to use
- Hardware and cost evolution is becoming a serious concern
 - Only support for x86 CPU architecture is a problem
 - Find and use additional resources outside WLCG
- Scalability and sustainability are key issues
 - Move to industry standards, use common tools

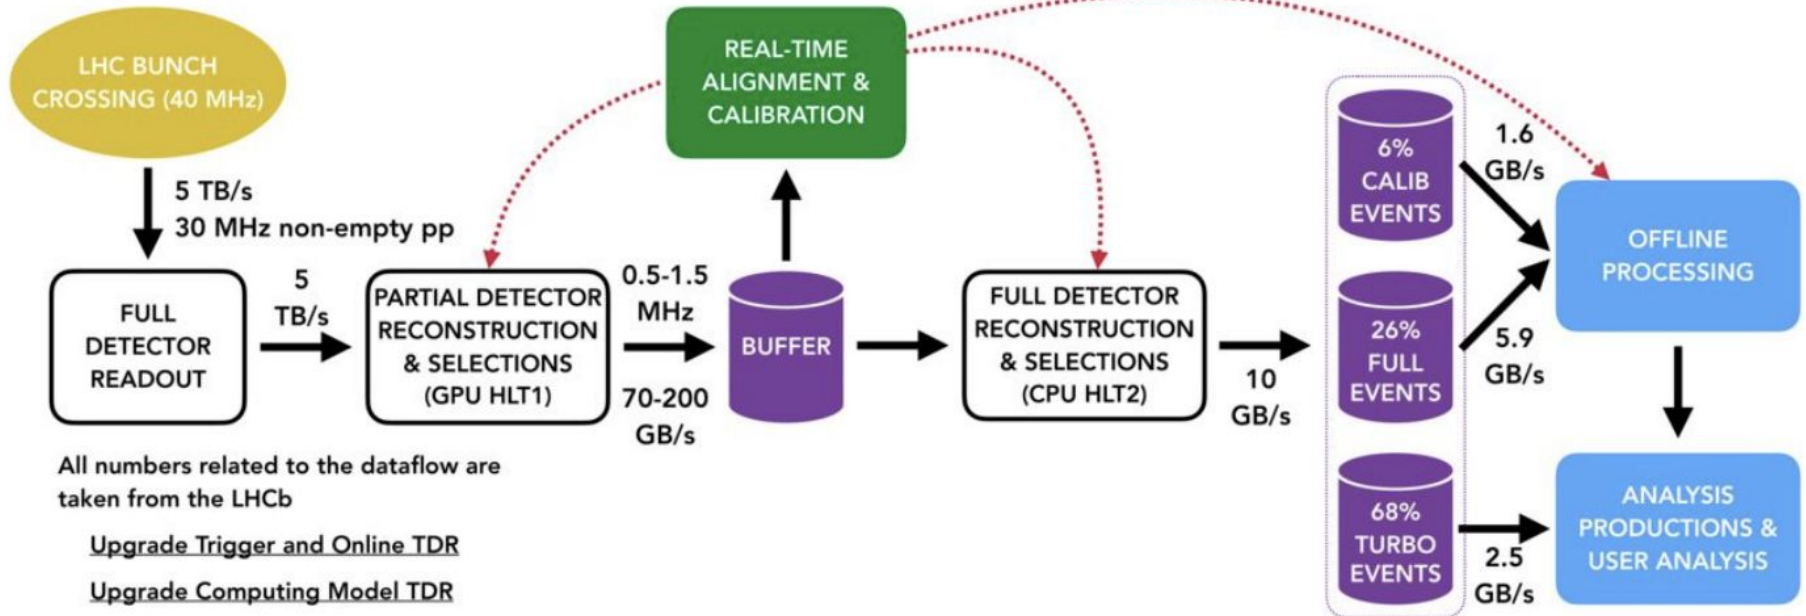
LHC / HL-LHC plan



- Run 3 (2022-2025): ~2x more data. **Evolutionary** changes in computing models
- Run 4 (HL/LHC, 2029+): ~20-30x more data. **Revolutionary** changes required

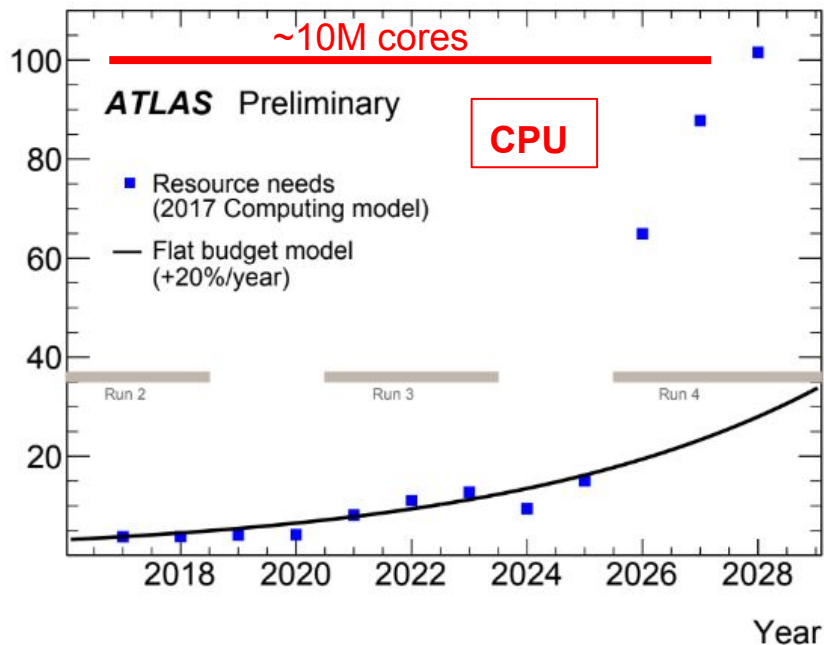
LHCb Run 3

- All-software-trigger @ 30 MHz readout
- >10x increase in output data rate (10 GB/s)
- Discarding raw data for $\sim\frac{2}{3}$ of the events after prompt reconstruction
- Large resource increases required ($\sim 50\%$ /year)




The HL-LHC computing challenge: ATLAS

CPU Resources [kHS06*1000]

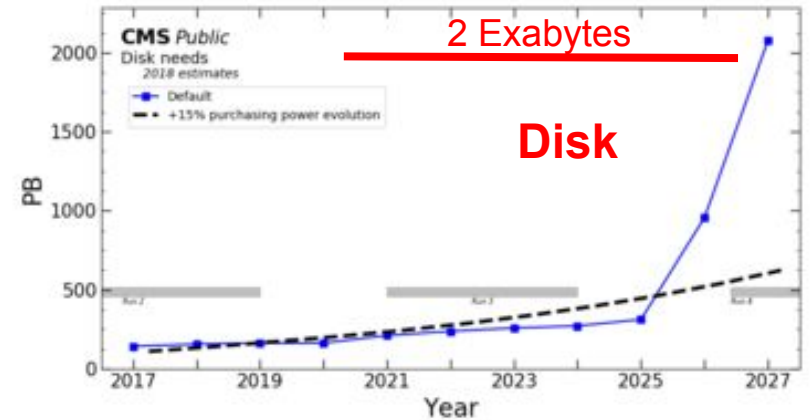
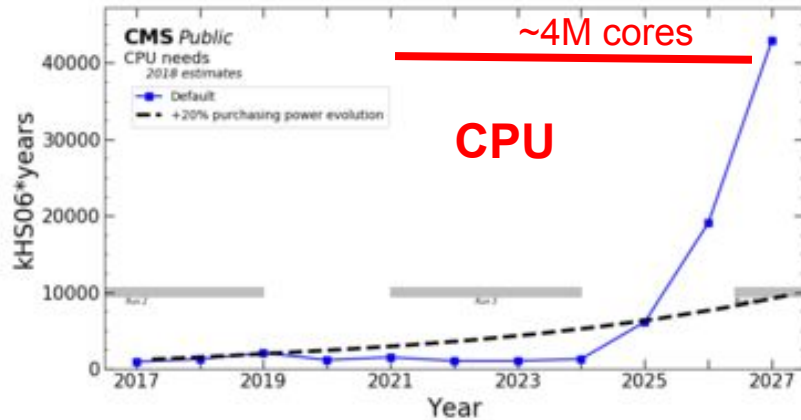


Disk Storage [PBytes]



- 2018 predictions: ~4-5x gap between “flat budget - 20% annual increase” and resource requirements for HL-LHC
- Intense R&D to reduce data and compute resource requirements

The HL-LHC computing challenge: CMS



LHC computing roadmap

arXiv:1712.06982v5 [physics.comp-ph] 19 Dec 2018

HSF-CWP-2017-01
December 15, 2017

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

HEP Software Foundation¹

ABSTRACT: Particle physics has an ambitious and broad experimental programme for the coming decades. This programme requires large investments in detector hardware, either to build new facilities and experiments, or to upgrade existing ones. Similarly, it requires commensurate investment in the R&D of software to acquire, manage, process, and analyse the shear amounts of data to be recorded. In planning for the HL-LHC in particular, it is critical that all of the collaborating stakeholders agree on the software goals and priorities, and that the efforts complement each other. In this spirit, this white paper describes the R&D activities required to prepare for this software upgrade.

¹ Authors are listed at the end of this report.

<https://doi.org/10.1007/s41781-018-0018-8>

WLCG Strategy towards HL-LHC

Executive Summary

The goal of this document is to set out the path towards computing for HL-LHC in 2026/7. Initial estimates of the data volumes and computing requirements show that this will be a major step up from the current needs, even those anticipated at the end of Run 3. There is a strong desire to maximise the physics possibilities with HL-LHC, while at the same time maintaining a realistic and affordable budget envelope. The past 15 years of WLCG operation, from initial prototyping through to the significant requirements of Run 2, show that the community is very capable of building an adaptable and performant service, building on and integrating national and international structures. The WLCG and its stakeholders have continually delivered to the needs of the LHC during that time, such that computing has not been a limiting factor. However, in the HL-LHC era that could be very different unless there are some significant changes that will help to moderate computing and storage needs, while maintaining physics goals. The aim of this document is to point out where we see the main opportunities for improvement and the work that will be necessary to achieve them.

During 2017, the global HEP community has produced a white paper - the Community White Paper (CWP), under the aegis of the HEP Software Foundation (HSF). The CWP is a ground-up gathering of input from the HEP community on opportunities for improving computing models, computing and storage infrastructures, software, and technologies. It covers the entire spectrum of activities that are part of HEP computing. While not specific to LHC, the WLCG gave a charge to the CWP activity to address the needs for HL-LHC along the lines noted above. The CWP is a compendium of ideas that can help to address the concerns for HL-LHC, but by construction the directions set out are not all mutually consistent, nor are they prioritised. That is the role of the present document - to prioritise a program of work from the WLCG point of view, with a focus on HL-LHC, building on all of the background work provided in the CWP, and the experience of the past.

At a high level there are a few areas that clearly must be addressed, that we believe will improve the performance and cost effectiveness of the WLCG and experiments:

- **Software:** With today's code the performance is often very far from what modern CPUs can deliver. This is due to a number of factors, ranging from the construction of the code, not being able to use vector or other hardware units, layout of data in memory, and end-end I/O performance. With some level of code re-engineering, it might be expected to gain a moderate factor (x2) in overall performance. This activity was the driver behind setting up the HSF, and remains one of the highest priority activities. It also requires the appropriate support and tools, for example to satisfy the need to fully automate the ability to often perform physics validation of software. This is essential as we must be adaptable to many hardware types and frequent changes and optimisations to make the best use of opportunities. It also requires that the community develops a level of understanding of how to best write code for performance, again a function of the HSF.

1

<https://cds.cern.ch/record/2621698>

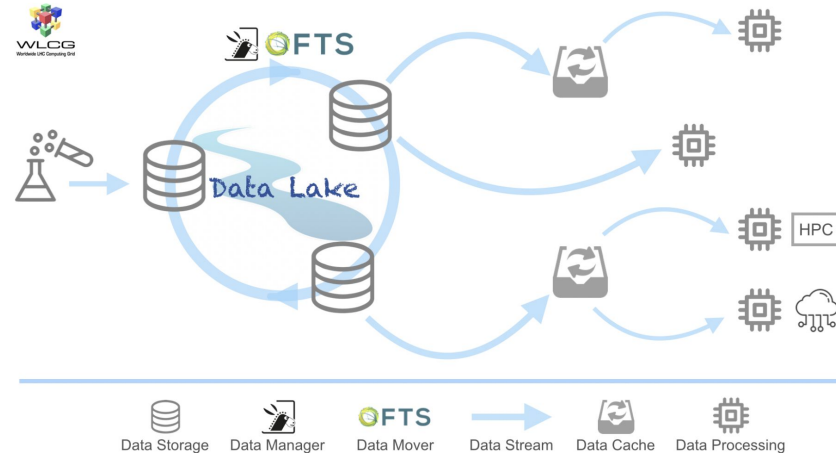
LHC computing roadmap: storage

Reduce amount of data

- Current data processing model is based on data replication and local access
 - Local access to data is more efficient (bandwidth, latency, reliability)
 - A single site has not enough CPU to fulfil the processing demands
 - Datasets are replicated several times at different sites for performance reasons
 - Lots of sites (150+) with managed storage

LHC computing roadmap: storage

- Build powerful data repositories (**data lakes**) and serve data to remote CPU resources
 - Reduce operational cost: deploy fewer (larger & federated) storage services
 - Global redundancy, economy of scale
 - Efficient data streaming needed: **Content delivery service, data caching**
 - Caching layer to hide network latency (read-ahead) and reduce data transfers over the network (cache hits)
 - Data cache is unmanaged and requires small capacity



Reduce amount of data

- Less data ➔ less storage, less processing and analysis compute needs
 - Reduce **trigger** output rate (HL-LHC planned 7.5 kHz ➔ ?)
 - Reduce data **formats**
- **Impact** in physics?
- NanoAOD format in CMS
 - ~1 kB/event
 - Goal: to be used by 50% of physics analyses
 - Reduces by 4x CMS storage needs for HL-LHC


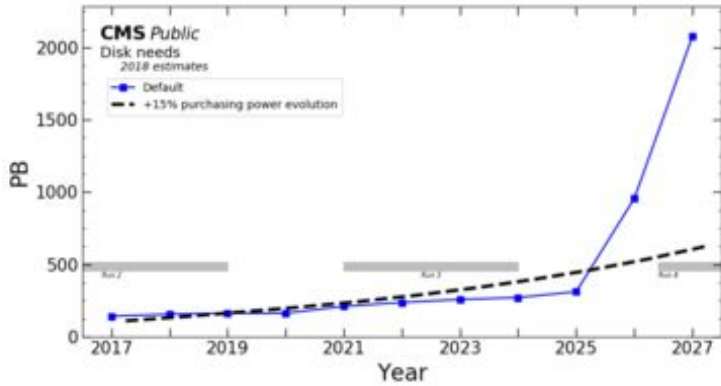


Data Tier	Size (kB)
RAW	1000
GEN	< 50
SIM	1000
DIGI	3000
RECO(SIM)	3000
AOD(SIM)	400 (8x reduction)
MINIAOD(SIM)	50 (8x reduction)
NANOAOD(SIM)	1 (50x reduction)


Analysis data formats

Storage reduction by using smaller data format

2018





2021



- Intense R&D programme to reduce computing resource needs
- Use of nanoAOD data format reduces disk space needs by 4x!
- Flat budget should be ~enough!!! But sustained funding required!

ATLAS storage requirements reduction



LHC computing roadmap: CPU

Reduce CPU needs or get access to new resources

- Use external resources
 - Supercomputers?
- Use new CPU architectures
 - Accelerators or co-processors (GPU)
- Make the software more efficient
 - New of faster algorithms (parallelization, machine learning)
- All that requires significant investment in software, infrastructure and service adaptations

Exploiting supercomputers for LHC

- Lot of public **funding** worldwide in supercomputer (HPC) facilities
 - Defined roadmap towards ExaFlop machines
 - e.g. EuroHPC B€ funding: 2 ~200 PFlop machines by 2022, 2 exaFlop by 2025
 - Funding agencies pushing us to use those resources
- Data intensive computing with HPC facilities is a **challenge**
 - Limited/no network connectivity in compute nodes
 - Limited storage for caching input/output event data files
- Our applications are not really suited for HPC
 - No large parallelization (no use of fast node interconnects)
 - No substantial use of accelerators (GPU)
- Substantial **integration** work to make HPC work for HTC
 - No one-fit-all solution: each facility is different
 - Little effort available in the LHC experiments
- Not suitable resource **allocation** model
 - We would need a guaranteed share of resources rather than apply for allocations

Barcelona supercomputing center & LHC computing

- The Barcelona Supercomputing Center (BSC) is the largest supercomputing center in Spain
 - MareNostrum4 (150k CPU cores); MareNostrum5 10x larger (expected from 2023)
- BSC - WCLG-ES agreement
 - LHC computing designated as a BSC “strategic project”
 - Access to dedicated resources (up to 7% of MareNostrum4)
 - Providing CPU for LHC simulation (~50M hours/year, ~50% of WCLG-ES CPU)




Hardware evolution

- WLCG compute resources based on x86 architecture (Intel, AMD)
- Resources outside WLCG (e.g. HPCs) available in other CPU architectures
 - IBM Power9
 - ARM (low energy consumption with lightweight cores)
- Dramatic development of massively parallel architectures
 - Graphics Processing Units (GPU)
 - Field Programmable Gate Arrays (FPGA)
- New HPC machines will bring a lot of these cards

Use of compute accelerator cards


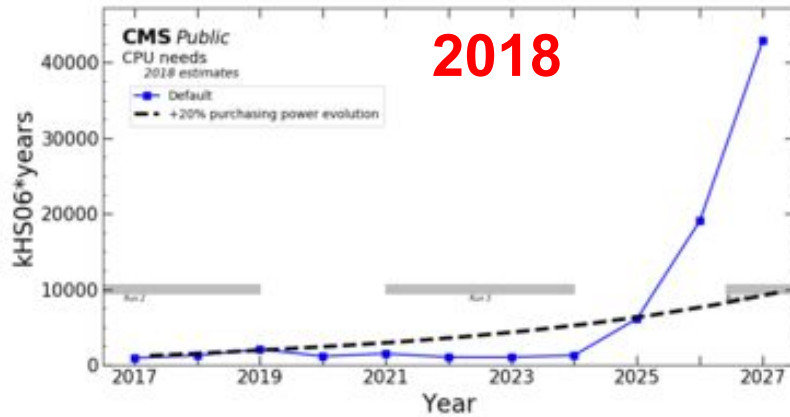
- Potential large speed improvement from hardware accelerated coprocessors
 - Larger performance/€ and smaller electric consumption/performance
- **Difficult to use**
 - Need to re-engineer HEP codes to a massively parallel environment
 - Data ingestion can be a limiting factor
- Very suitable for certain applications
 - E.g., excellent at training deep neural networks



Software optimization


- Recent initiatives
 - HEP Software Foundation (coordinate software R&D for LHC)
 - Institute for Research & Innovation in Software for HEP (IRIS-HEP); 25M\$, 5 years
 - Proposal a EU scientific software institute
 - COMCHA forum in Spain
- Exploit **new hardware architectures**
 - **High level parallelism**, new instruction sets, non x86 processors
 - Support in software frameworks for **heterogeneous** hardware
 - Support for multi-threading, vectorisation, CPU/GPU orchestration
- **Innovative algorithms**
 - **Machine/deep learning**
 - Recast physics problem as machine learning problem vs re-rewrite physics algorithms for new hardware

CPU reduction by improving algorithms




- Intense R&D programme to reduce computing resource needs
- Improve simulation and reconstruction algorithms
- Offload part of the processing work to GPUs (~3x better cost/performance)

CMS Public
Total CPU HL-LHC (2029/No R&D Improvements) fractions
2021 Estimates



ATLAS CPU needs reduction by using fastsim/fastreco

Faster physics algorithms:
exploit more broadly fast
simulation & reconstruction




HL-LHC computing challenge status

- Intense R&D has already paid off in a drastic reduction of computing resource needs for HL-LHC
 - Attainable with constant funding (~20% annual increase with cheaper hardware)
- Data challenge can be met extensively using smaller data formats
 - ~1 kB/event for analysis
- Compute challenge can be met through software improvements, speeding up algorithms and using new hardware architectures
 - GPUs and accessing supercomputing resources
- **But keeping the current level of funding is needed!!!**
 - High electricity prices is a serious concern!

Machine learning




Training



Inference

- Learn (training) and predict (inference - classification or regression)
- Mainly based on (deep) multilayer neural networks (DNN)
- Flourished due to wealth of data and huge processing capacity
 - New hardware architectures (GPU, TPU)
 - Developments of AI algorithms



Use of ML in HEP

- Started to be used in 1990s and 2000s; explosion of applications in 2010s
- Used in almost all aspects of the HEP experiments with large impact
 - Data analysis
 - Event classification, anomaly detection
 - Online event selection (trigger)
 - Data quality monitoring
 - Anomaly detection
 - Object reconstruction, identification and calibration
 - Jet substructure, b-tagging, etc
 - Event fast simulation
 - Electromagnetic calorimeter showering
- [Machine learning in HEP community white paper](#)
 - Outline of R&D for the next decade

Data analysis ecosystem

- Challenges in scaling up HEP analysis to meet the needs of HL-LHC
 - [2nd Analysis ecosystem workshop, 2022](#)
- Development of highly performant data analysis systems that **reduce "time-to-insight"** and maximize physics
 - **Specialized infrastructure** and services that provide integrated data, software and computational resources to execute analysis workflows
 - Innovate in existing community tools like ROOT and incorporate new cutting-edge python data science tools
 - **Fast analysis turnaround** is key
 - Fast access to input data, data processing parallelism
 - New **data formats** for performance reasons (columnar analysis, RNTuple)
 - **Declarative interfaces** (RDataFrame, Coffea)

Quantum computing

- Quantum technology is an emerging field of physics and engineering with the potential to revolutionise science and society
 - Quantum effects, such as **superposition and entanglement**, are used to speed up certain classes of computational problems beyond the limits achievable with classical systems based on logical bits
- Large investment in development of quantum technologies for computing
- CERN Quantum Technology Initiative
 - CERN ambitions to be at the forefront of this revolution
 - Foster innovative ideas in the field of high-energy physics and act as a hub for innovation and knowledge creation and sharing
 - Collaboration with industry and academia
 - Ongoing [R&D projects](#)
 - Computing and algorithms, quantum sensors, communication and networks
 - [Strategy and roadmap](#)

Quantum computing in Spain


- Quantum Computing Technology [Group at IFAE](#)
 - Building quantum processors out of superconducting quantum circuits
- Quantic group at BSC
 - Leading the Quantum Spain project to create a national quantum computing ecosystem for Artificial Intelligence
 - Funded with 20 M€ together with other members of Spanish supercomputing network
 - Build a quantum computer with superconducting qubits
 - Create a cloud-based remote access service to the processor, to enable industry and the public sector to experiment with new quantum algorithms
 - Develop useful quantum algorithms

Leveraging WLCG infrastructure and services

- Other future HEP projects beyond LHC are largely increasing their experimental data volumes and processing requirements
 - Neutrinos (DUNE), high energy gamma rays (CTA), gravitational waves (ET), radio astronomy (SKA), etc
- Create economies of scale, through the adoption of common approaches for data management
- Share existing infrastructure
 - Sites supporting several projects using the same tools
- Use existing services
 - Large scale data management (data transfers, data streaming, caching, etc)
 - Compute resource provisioning and scheduling
 - Access to opportunistic resources using existing interfaces (HPCs, Clouds)
 - Authentication and Authorization infrastructure

Support of HEP experiments at PIC

- Tier-1 for ATLAS, CMS, LHC
- Tier-0 for MAGIC and PAUS
- Science Data Center for EUCLID
- Data center for CTA
- Support for VIRGO/LIGO, T2K, DUNE



The ESCAPE project

- ESCAPE (European Science Cluster of Astronomy & Particle physics ESFRI research infrastructures) brings together the astronomy, astroparticle and particle physics communities with aligned challenges of data-driven research
- EU funded project, 2019-2023, 16 M€
- Build a link between ESFRI projects and e-infrastructure providers
- Provide access to a scalable federated data infrastructure

Outlook

- Management of exabyte-scale science data in HEP
 - HL-LHC and other experiments reaching similar scales
- Building on solid ground
 - Distributed high throughput computing infrastructure developed over the past two decades for LHC computing
 - Intense R&D program ongoing
- Changing landscape in resources, architecture and technologies
 - Heterogeneous facilities (HPCs, dedicated HTC data infrastructure) and resources (GPU)
 - New technologies, software algorithms (machine learning, quantum computing)
 - Federated, network-centric computing is even more important for future

Outlook

■ Sustainability

- Infrastructures & centres likely to be **common** between HEP & Astronomy, Astroparticle, Gravitational Waves, etc.
 - Leverage infrastructure, effort, technologies
- HEP should keep at the forefront
 - Share our experience
 - Synergies and collaborations across disciplines and domains is important and positive