

R&D in LHC Computing projects in Spain

XVII CPAN Days


19-21 November 2025, Valencia, Spain



Pablo Collado Soto on behalf of the (quite large) Spanish LHC Computing Community

Computing in HEP


The final result



Much more than meets the eye...


- Our **ultimate goal** is to 'do **physics**'.
- **Analysis** results are **succinctly conveyed** by **plots** such as the previous one.
- But **producing** them **encompasses huge and varied efforts**:
 - **Accelerator** development and operation: the **LHC**.
 - **Detector** development and operation: **ATLAS, CMS, LHCb...**
 - **Computing** development and operation: the **WLCG**.
- This **talk focuses** on the **latter** of the **three** and the role the **Spanish community** plays in **computing** in the **HEP** landscape.

The shadow of the HL-LHC I




Last update: November 24



The shadow of the HL-LHC II






Projected resource usage I



ATLAS Software & Computing HL-LHC Roadmap
These plots will be updated in 2026 in an ATLAS TDR



Projected resource usage II



CMS Computing Model Update
These plots will be updated in 2026 in a CMS CDR

Projected resource usage III


**Completing R&D projects
minimises computing needs!**

Can the WLCG tackle the challenge? I



- **Short answer: yes!**
- The **WLCG** has been **operating** since **2004...**
 - ... **always** in **production** even in the face of **large upgrades**.
 - ... **optimising processing** needs **without impacting** the physics programmes.
 - ... **continuously growing** its **capacity** with (at **best**) **stable funding**.
 - ... **including** diverse **resources** such as **cloud** and **HPC** facilities.
- **Growth** has **exceeded** initial **expectations**.
- The **WLCG** has **been** and **is** a **key** enabler for the **LHC** physics program!
- The **LHC** experiments are also **heavily investing** in **evolving** their **computing models**.

Can the WLCG tackle the challenge? II



Key messages from the ESPP Symposium

Future HEP Computing

The goal of the ongoing software and computing development is to facilitate the successful and timely delivery of the best possible physics results


- 'Completing the **HL-LHC** programme is **essential** and must remain a **high priority** for **CERN**':
 - **Computing** is a **pillar** of the **HL-LHC** program.
 - Thus, it is **crucial** to **evolve** and **upgrade computing** models so that they **rise** up to the **challenge** posed by the **HL-LHC**.
- HL-LHC computing is not an easy journey, but it should be feasible, provided an **intense and uninterrupted R&D** is carried out.
- Likewise, for any future project chosen, the required software and computing should be feasible, **provided adequate person power and funding are available**
- **Successful projects and collaborations** such as WLCG, GEANT, ROOT, KEY4HEP, ... and their deliverables and services **need to be sustained**
- Make sure not to miss the opportunities to **adopt new new technologies**, particularly GPUs and AI, as well as Quantum Computing on a longer timeframe.
- We may not have a clear picture what our computing environment will look like 40 years from now, **but we have an excellent record to adapt and use whatever is available to its full potential.**



Spanish R&D in computing

The Spanish computing community

- Beside operating a highly reliable infrastructure, the Spanish computing community contributes to several R&D efforts in computing.
- The rest of the talk introduces these fields of work.



*



The ATLAS, CMS and LHCb Tier-1 and Tier-2 contribute 5% of the experiment's computing resources.

* Be sure to check Álvaro Fernáñez Casani's talk [on computing in LHCb](#) earlier today.

Infrastructure

Data management and infrastructure

- **HEP** is a very **data intensive** field: we need to **buy** and **operate** large **storages**.
- These **storages** are **backed not only** by **HDDs**, but also by **tapes operated** with purposefully-designed **software** such as **CERN Tape Archive** (CTA).
- **Processing** and data **generation** can be largely **externalised**, but **data** should **remain** under every **collaboration's control**!
- **R&D** is being carried out by **PIC** on the **use** of **caches** to **reduce access times** and the **need** for '**cold**' **storage** on sites:
 - **Spain** operates a **federated Tier-2** for ATLAS: a **data-lake model** could be very beneficial!
- **Data challenges** (DC24 and DC27) **stress** our **infrastructure** to find **bottlenecks**.
- **Backbone networks** (GÉANT, ESNet) are a **critical infrastructure** to consider.
- **Operational experience** in the **ES** cloud allow us to **dimension** the **infrastructure** for HL-LHC.


ATLAS Event Index I

- **EventIndex**: Catalogue of real and simulated events.
- **Use cases**:
 - Event **picking**
 - **Counts** or **selections**
 - **Overlaps...**
 - Of **triggers** in a dataset
 - Of **events** in derivations
 - Production **checks**
- **AMiLoop is running without issues**:
 - **Datasets** are **attached correctly**.
 - **Search window**: last **3 days**.
 - **Search frequency**: every **12 hours**.
 - **Limits**: **300 datasets/h**; **5000 per day**.
- **Tech stack**:
 - **HBase + Phoenix**
 - **Scala, Spark** and **RDD API**.



ATLAS Event Index II


- **IFIC** is responsible for the **Data Extractor**, **Supervisor** and the **Phoenix Loader Apps**.
- In numbers (10/2025)...
 - **1050 billion** real + **1740 billion** MC data.
 - **121k** real + **581k** MC datasets.



Workload management

- **HEP** is a very **CPU-intensive field** too!
- **Available CPU** power must be **correctly** managed. This is the **task** for a **Workload Management System** such as CMS's **WM**, ATLAS' **PanDA**, LHCb's **DIRAC** and Alice's **AliEn**.
- **Spanish staff** at **PIC** is deeply **involved** in the **redesign** of **CMS' WM** to attain:
 - **Sustainability** and **scalability**: the **Spanish focus** is on the **Submission Infrastructure (SI)**.
 - **Efficiency**: reducing **time to completion** and **operational costs** with automation.
 - **Integration of heterogenous resources**: **abstracting** the **complexity** of the underlying **hardware** is key as we leverage **non-x86** platforms, **GPUs**, **coprocessors**...

Current CMS WM system




This system, currently in production, will be updated for the HL-LHC era.


Network infrastructure



The GÉANT network interconnects Europe's NRENs and provides onward connectivity to regional networks worldwide. Enabling the most important scientific discoveries, it is essential to Europe's research and education communities.



January 2024



PIC is planning upgrading their 200 Gbps symmetric connection to 400 Gbps before 2027 (DC27) in coordination with RedIRIS and CSUC.

UAM recently updated the connection to 100 Gbps.

IFIC uses a symmetric 200 Gbps connection.


Heterogeneous computing

A case study with event generation


- **Event generation** is one of the many **computationally intensive** tasks in HEP.
- As **precision increases**, the **Matrix Element (ME)** computation **becomes** a **bottleneck**.
- We can **leverage** new **architectures** to **speed up** the process in **event generators** such as **MadGraph**: **MadGraph4GPU** already showed nice **progress**!

Simulation, reconstruction and trigger will or are leveraging heterogeneous resources too!

The FPGA performs 19 times faster than the GPU and 51 times faster than the CPU.




Leveraging MN5 for ATLAS




Note: MN5 disallows outbound and inbound **Internet** access...

Leveraging MN5 for CMS



Leveraging MN5 for LHCb




Sustainability

Going greener I

- **HEP** is a very **energy intensive** field: **running** machines uses (a lot of) **energy**!
- **Efforts** are **ongoing** to **reduce** our **footprint**. The **WLCG** is **including sustainability** as a **main** talking point in **workshops**.
- **PIC** is **leading** the **effort** with:
 - **Immersion cooling** with mineral oils.
 - **Processing throttling** in the face of **high** energy **consumption** at the **national level** and/or when the **energy mix** is highly **dependent** on **non-renewable** energy sources.
- **UAM** installed **free-cooling** machines to **reduce** use **AC** in winter.
- **IFIC** and **UAM** are **developing** a **sustainability-oriented** monitoring **dashboard** to **evaluate** energy **efficiency** and environmental **impact** at **ATLAS Tier-2** sites.
- **Training** local **users** on to **optimise** resource **usage** is also **key**!

Going greener II




The final goal is to design a feedback loop to HTCondor at PIC for real-time power modulation based on green energy cycles, with the aim to reduce CO2 emissions and electricity costs.

Analysis Facilities

Reducing time-to-insight


Analysis Facility developed at CIEMAT. Other sites like IFIC are working on similar solutions.



Software development


Scitags I

Check the
SciTags site for
more info!




- The **SciTags** framework:
 - Is an **open platform** to be **used** by any **data-intensive science community**. It is **potentially** more **broadly** applicable.
 - Enables the **identification of the owner** (experiment) and **purpose** (activity) of traffic.
 - **Defines standards** for **exchanging information** between scientific **communities, sites and network operators**. This **information** can be **encoded** and **sent** with **different strategies**.
 - **Enables tracking/correlation** with **existing** network flow **monitoring infrastructure** and associated infrastructure already **deployed** by **RENs**.
 - The **technical specification** is readily available **here**.



Scitags II



Scitags III




Scitags IV




Flowd-go is developed
at **UAM**

Scitags V


Total Flows per Exp/Act




Total Flows per Prefix



Total Flows per IP version

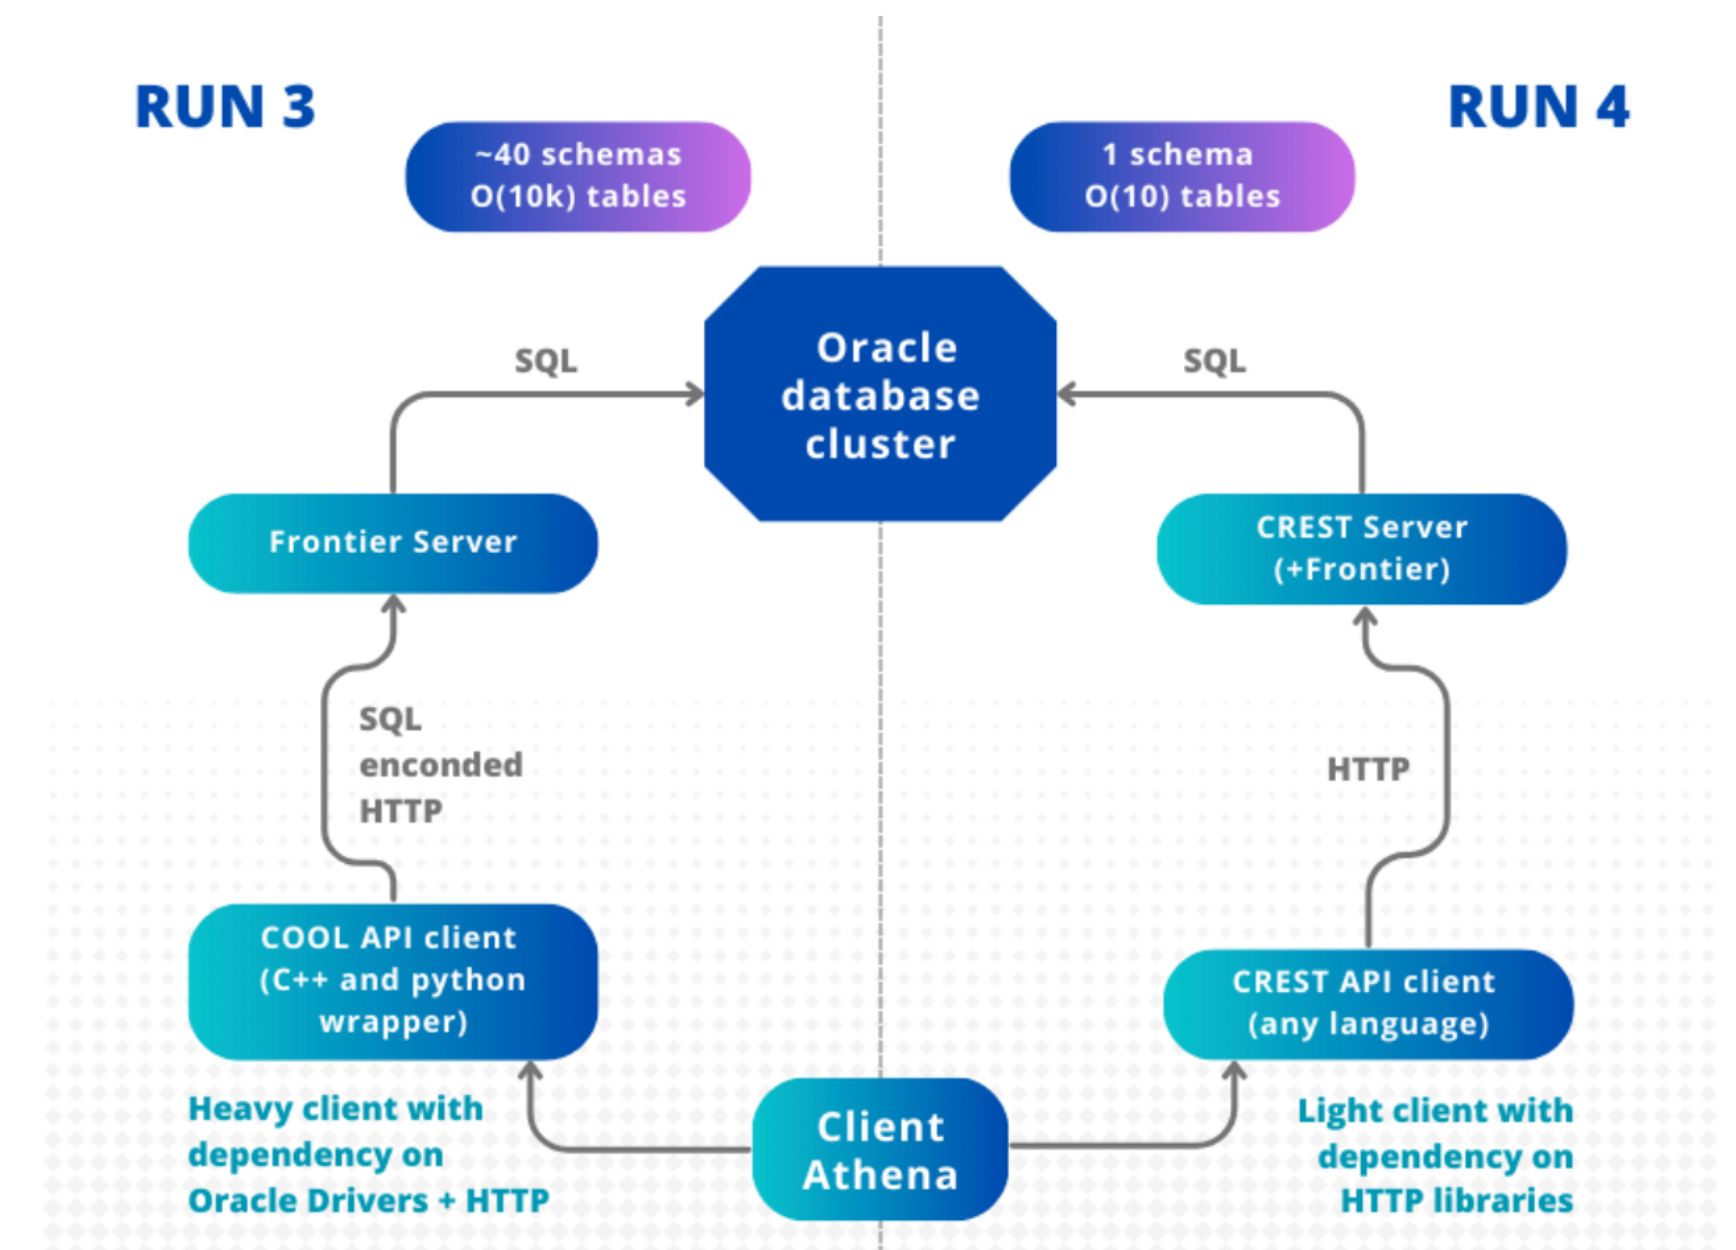


Total Flows per Application (top10)



CREST: Conditions REST

- **Conditions**: data describing the **detector's state** (but not collision data).
- **Conditions change** with **time** and **must** be **known** for event **reconstruction**.
- Large **architectural shift** in **access infrastructure** from **RUN 3** towards **HL-LHC**:
 - **Disentangle** the **client** and **backend** storage solution.
 - **Simplify** data **structures**.
 - **Enhance** **caching** capability.
- **IFIC** is in charge of:
 - **Developing** the **high-level client API**.
 - **Coordinating** the **development** and **deployment**.



Appointments

Management positions

- **Jose Enrique García Navarro** (IFIC) is an ATLAS Database And Metadata (**ADAM**) coordinator.
- **Antonio Pérez-Calero Yzquierdo** (PIC) is the **manager** for the **CMS Computing Submission Infrastructure**.
- **Jose Salt Cairols** (IFIC) is:
 - An ATLAS Distributed Computing Operation Shift (**ADCOS**) expert.
 - The **Spanish contact** for the **Computing Resources Review Board** (C-RRB).
- **José Flix** (PIC) is:
 - The **manager** for the **CMS Resource Management** Office.
 - The **chair** of the worldwide **HEP*i*X** forum.
 - A **member** of the **WLCG Technical Coordination** Board.
 - A **member** of the **WLCG Collaboration** Board representing PIC's **Tier-1**.
 - A **member** of the **WLCG Management** and **Overview** Boards representing **Spain**.
 - Was the **chair** of the **WLCG Deployment** Board from **2020 to 2024**.
- **Santiago González de la Hoz** (IFIC) is:
 - A **member** of the **WLCG** international project **Collaboration Board** (CB) in **representation** for **Tier-2 Spain ATLAS** Federation.
 - The **ATLAS Event Index** data **production coordinator**.
- **Andreu Pacheco Pages** (IFAE/PIC) is the **ATLAS Distributed Computing (ADC) Fabrics Coordinator**.
- **Jose María Hernández Calama** (CIEMAT):
 - **Chair** of the **CMS Computing Resources Board** (2020-2025)
 - **Chair** of the **WLCG Computing Resources Scrutiny Group** (2025-)

Summing up...



- **Computing** is a **pillar** of the **LHC** physics **program**.
- **Upgrades** are **necessary** to **meet** the stringent **demands** of the **HL-LHC** era.
- Pursuing **R&D efforts** is **key** to **lowering** computing **needs**.
- The **Spanish community** is:
 - **Operating** a highly **reliable** infrastructure...
 - ... and extensively **contributing** to **R&D projects**!
- **Personpower** is **needed** to **maintain** **progress** on course.
- **Efforts** are also being **made** to **enhance** computing **sustainability**.

ACKs

- Special thanks to everybody involved in putting the slides together:

L. Fiorini¹, A. Fernández Casani¹, J. Flix Molina^{2,3}, J. García Navarro¹, S. González de la Hoz¹, J. Hernández Calama², A. Pacheco Pages^{3,4}, A. Pérez-Calero Yzquierdo^{2,3}, J. Del Peso⁵, J. Salt¹, X. Vilasís Cardona⁶, M. Villaplana¹

¹Instituto de Física Corpuscular (IFIC), University of Valencia and CSIC, Valencia

²Centre for Energy, Environmental and Technological Research (CIEMAT), Madrid

³Port d'Informació Científica, UAB, Bellaterra

⁴Institut de Física d'Altes Energies (IFAE), UAB, Bellaterra

⁵Department of Theoretical Physics and CIAFF, UAM, Madrid

⁶La Salle-Universidad Ramón Llull, Barcelona.

Any questions?

Backup

Challenge overview

- **Challenges for LHC computing from 2025 and into the HL-LHC era include:**
 1. **Agreeing** on how to **increase** the **infrastructures** to ready them for the expected **large** amount of **data**. This should be **ready** by **2030!**
 2. Incorporating **software engineering** practices to **optimise computing** resource usage.
 3. **Controlling** and **minimising data** volumes.
 4. **Applying AI** in different layers and **leveraging** it within the **LHC computing model**.

Open source contributions



- Operating GRID infrastructure offers development opportunities:
 - At UAM we have developed a **SLURM plugin** for Telegraf.
 - At UAM we host a public **Almalinux mirror**.

Flowd-go architecture

