# Computing Activities at the Spanish Tier-2s for the ATLAS experiment in the LHC Run3 period and towards the High Luminosity Phase (HL-LHC)

XL Physics Biennial RSEF 2026
20-24 June 2026, Sevilla

**A. Pérez García[1], E. Acción[3], V. Acin[3], C. Acosta-Silva[3], H. Burriel[1], J. Cantero[1], P. Collado[2], J. Del Peso[2], J.E. García Navarro[1], J. Gibson[1], S. González de la Hoz[1], A. Pacheco Pages[3], E. Planas[3], J. Salt[1], J. Sánchez[1], E. Torró[1], R. Uzum[1], M. Villaplana[1]**

On behalf of the Spanish ATLAS Federated Tier-2 community

[1] Instituto de Física Corpuscular (IFIC), University of Valencia and CSIC, Valencia, Spain;
[2] Departamento de Física Teórica y CIAFF, Universidad Autónoma de Madrid, Spain;
[3] Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, Bellaterra (Barcelona), Spain

This work was supported by the Spanish Ministry of Science, Innovation and Universities (MICIU) through the ATLAS Tier-2 computing projects PID2022-136323OB-C21 (IFIC), PID2022-142604OB-C22 (IFAE), PID2022-136323OB-C22 (UAM), project RYC2022-038164-I, and by the Generalitat Valenciana through the GenT Programme (ESGENT-CIESGT2025).

# Spanish ATLAS Tier-1 & Tier-2s

- **Tier-1:**
  - **PIC-Barcelona**
- **A Federated Tier-2 (ATLAS):**
  - **60% IFIC-Valencia**
  - **25% IFAE-Barcelona**
  - **15% UAM-Madrid**

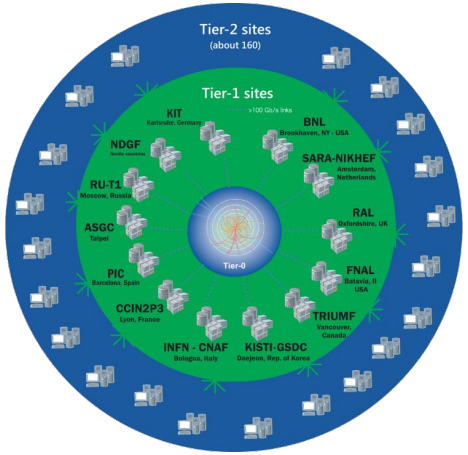


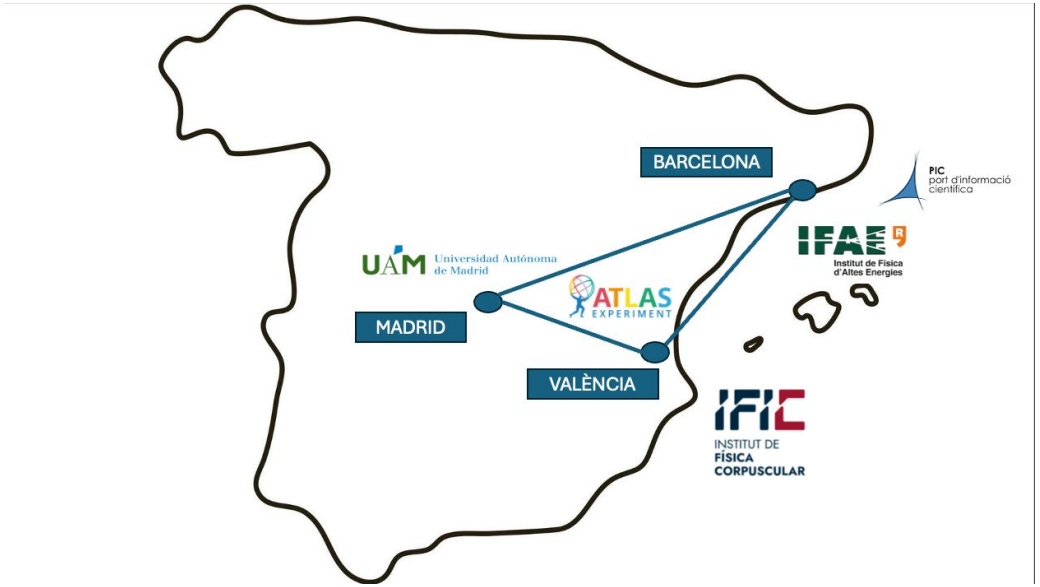


- Integrated in the WLCG project (World Wide LHC Computing GRID) and strictly following the experiment computing model.
- We account for **5% of the total Tier-1 resources and 5% of the total Tier-2 resources**, respectively.
- In 2025, the availability & reliability for the **Tier-1 site exceeded 97%**, while the three **Tier-2 sites** each achieved availability & reliability **above 95%.**


# Spanish ATLAS Tier-1 & Tier-2s (Pledge & Resources)

| Site | Current CPU Pledge 2026 | Comments |
|---|---|---|
| *Barcelona (PIC T1)* | 90,100 HS23 | Usage reached **164%** of the pledge in 2025, driven by BSC usage |
| *Barcelona (IFAE T2)* | 27,525 HS23 | **138%** pledge delivered in 2025 |
| *Madrid (UAM T2)* | 16,515 HS23 | **233%** pledge delivered in 2025 |
| *Valencia (IFIC T2)* | 66,060 HS23 | **265%** pledge delivered in 2025 |

| Site | Current Tape Pledge 2026 | Current Space | Comments |
|---|---|---|---|
| *Barcelona (PIC T1)* | 34,600 TB | 29,800 TB | Currently 2025 pledge +10%. New tape library is being installed. 2026 pledge will be ready soon. |

| Site | Current Disk Pledge 2026 | Current disk (May 2026) | Comments |
|---|---|---|---|
| *Barcelona (PIC T1)* | 9,950 TB | 10,000 TB | **100%** ES-ATLAS pledge |
| *Barcelona (IFAE T2)* | 3,038 TB | 3,050 TB | **25%** ES-ATLAS pledge |
| *Madrid (UAM T2)* | 1,822 TB | 1,973 TB | **15%** ES-ATLAS pledge |
| *Valencia (IFIC T2)* | 7,290 TB | 6,450 TB | **60%** ES-ATLAS pledge. Purchase to be installed in November |

**Current pledges are covered**

# Grid ATLAS Tier-2 Jobs

- Grid jobs running during Run 3 in the Spanish cloud:
  - **Wall-clock time efficiency: 95.1%** vs 93.2% (all sites).
  - **Average CPU Efficiency (successful jobs): 86.2%** vs 87.4% (all sites).
  - **Percent of failures: 7%** vs 7% (all sites).

Slots of running jobs (HS23) Source

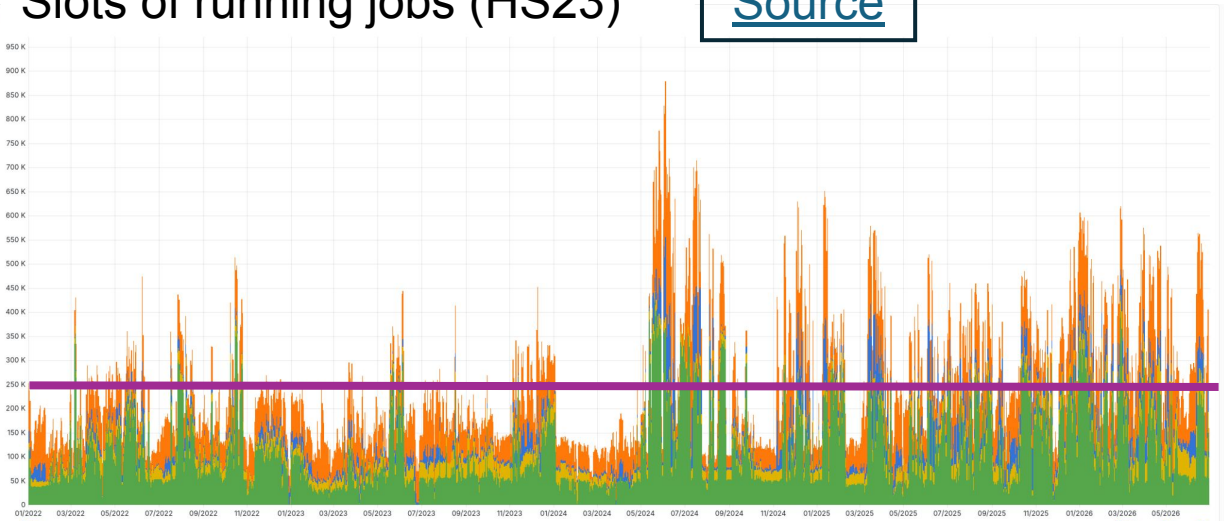

| Name | Mean | Max | Min |
|---|---|---|---|
| pic | 115 K | 420 K | 392 |
| IFIC-LCG2 | 86.3 K | 324 K | 1.43 K |
| UAM-LCG2 | 24.4 K | 166 K | 0 |
| ifae | 23.9 K | 68.5 K | 0 |

**Almost 250k slots of running jobs throughout Run 3.**

Completed jobs (HS23) Source

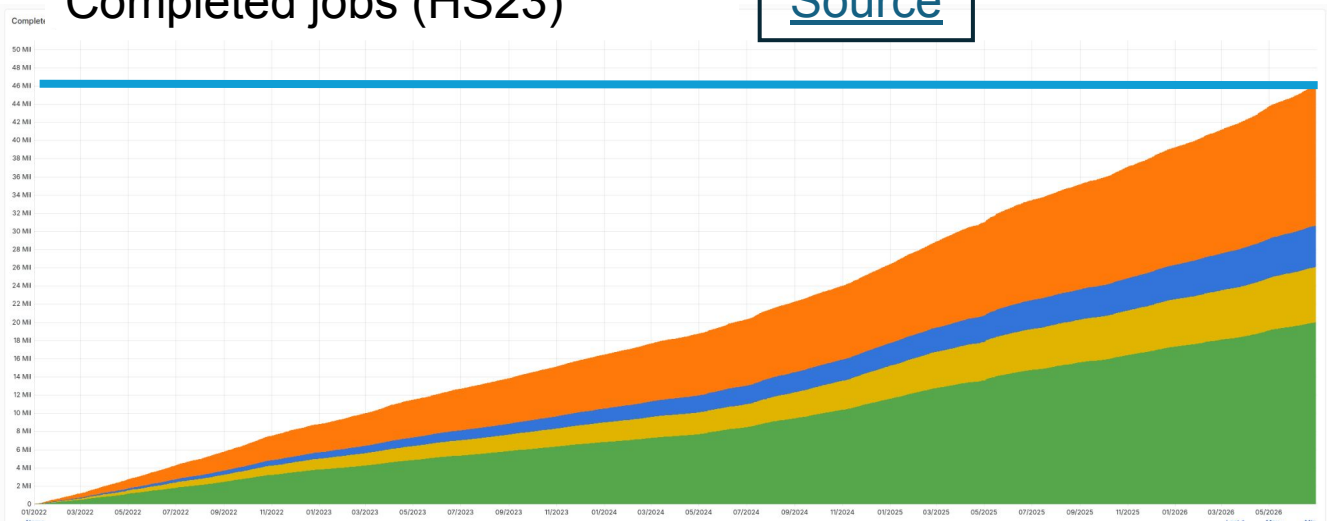

| Name | Last * | Max | Min |
|---|---|---|---|
| pic | 20.0 Mil | 20.0 Mil | 6.39 K |
| IFIC-LCG2 | 15.4 Mil | 15.4 Mil | 7.84 K |
| ifae | 6.10 Mil | 6.10 Mil | 1.00 K |
| UAM-LCG2 | 4.55 Mil | 4.55 Mil | 1.71 K |

**Almost 50M completed jobs throughout Run 3.**

# HPC Jobs

- Use of the **Mare Nostrum5 (HPC)** by ATLAS Tier-1 and Tier-2s:
  - Using **ARC-CE** at PIC, IFIC and UAM to interconnect Mare Nostrum and ATLAS production system.
  - MareNostrum (MN5) accepts only SSH protocol for job submission and data transfer. The worker nodes have **no outbound connection**.
  - Proportion of HS23 (s) provided by GRID resources (blue) and the MN5 HPC (orange) in total contribution to the ATLAS computing by the Spanish cloud.
  - **141 million hours approved at MN4 + MN5 throughout Run 3** by ATLAS through Spanish gateways, which corresponds to **50% of the simulation jobs assigned to Spain.** IFAE is using MN5 GPU for Analysis (Inference).
  - Plans: MN5 HPC and successors are in planning phase. *Currently pushing to request connectivity from MN5 worker nodes.*

Source

Slots of running jobs - simulation (HS23)

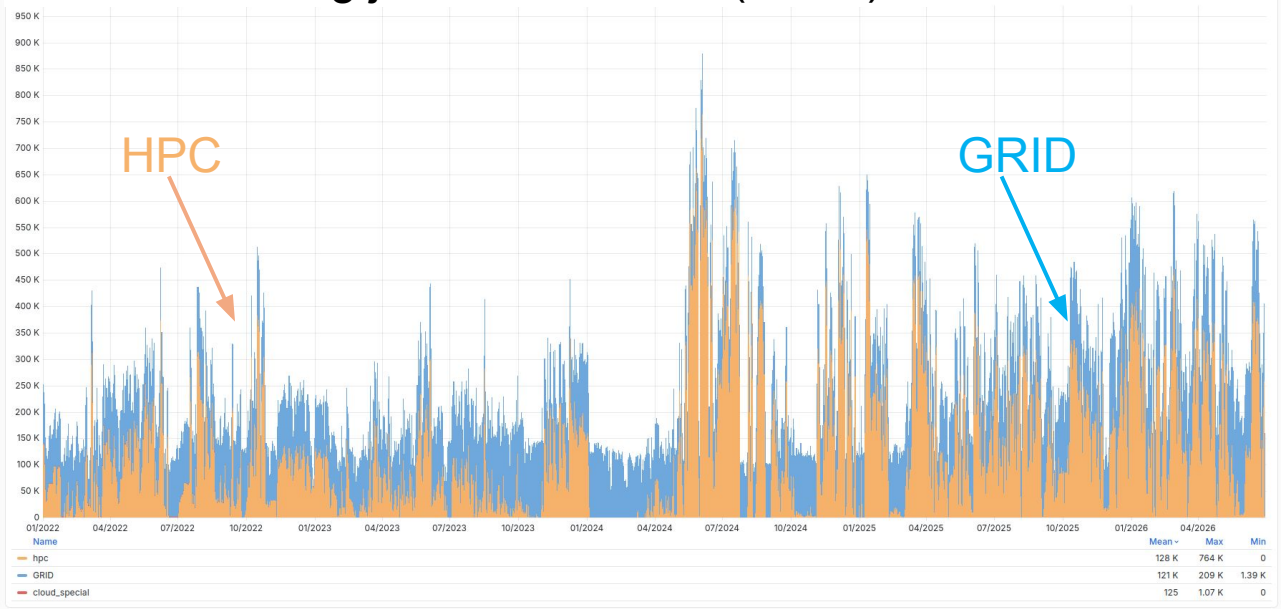


Wallclock time (HS23 seconds)

Source

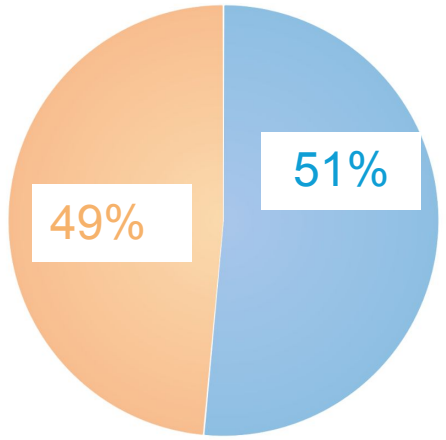


GRID 15.9 Tri 51%
hpc 15.0 Tri 49%

**50% CPU payload on HPC**

# Data Management

- Total transfer volume overview in Run 3:
  - As a source - **192 PB 90% effic.**
  - As destination - **235 PB 90% effic.**

Transfer volume as source

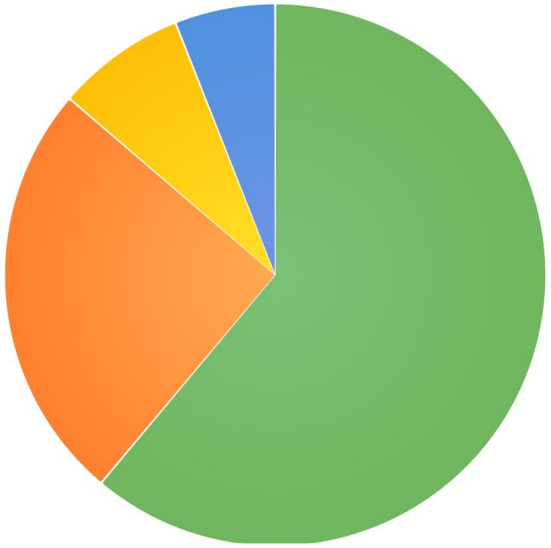

As source (link)

| | | |
|---|---|---|
| pic | 118 PB | 61% |
| IFIC-LCG2 | 48 PB | 25% |
| ifae | 15 PB | 8% |
| UAM-LCG2 | 11 PB | 6% |

Transfer volume as destination

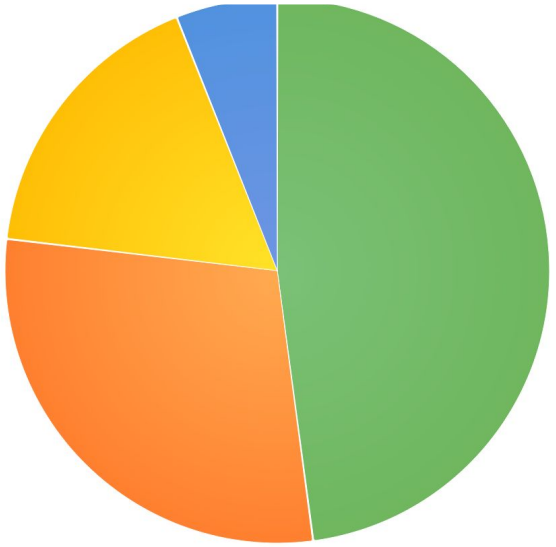

As destination (link)

| | | |
|---|---|---|
| pic | 113 PB | 48% |
| IFIC-LCG2 | 68 PB | 29% |
| ifae | 40 PB | 17% |
| UAM-LCG2 | 14 PB | 6% |

# Network Infrastructure

- Spanish cloud update
  - **LHCONE:** all sites have now successfully migrated to LHCONE.
  - **Transfer rates:** peaks of up to 180 Gbps.
  - **HL-LHC Readiness**: Utilizing **DC24/DC27 to identify bottlenecks** and scale the infrastructure for the future requirements of High-Luminosity LHC era.
  - **Monitoring tools**: recommended tool by WLCG developed at UAM
  - **Mid-term plans:** By Run 4, current goal is to achieve 600Gbps (2031).

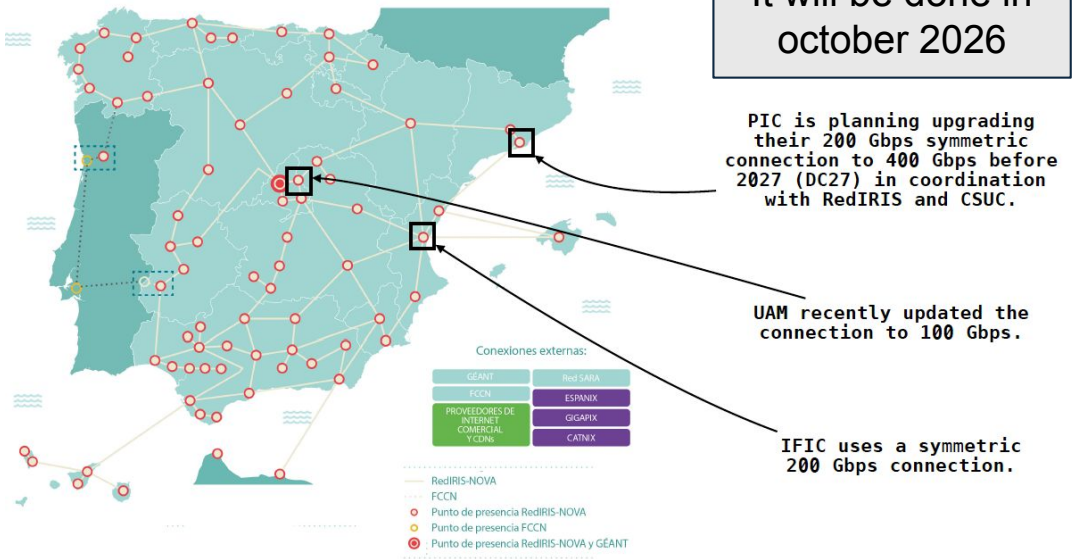


It will be done in october 2026

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.

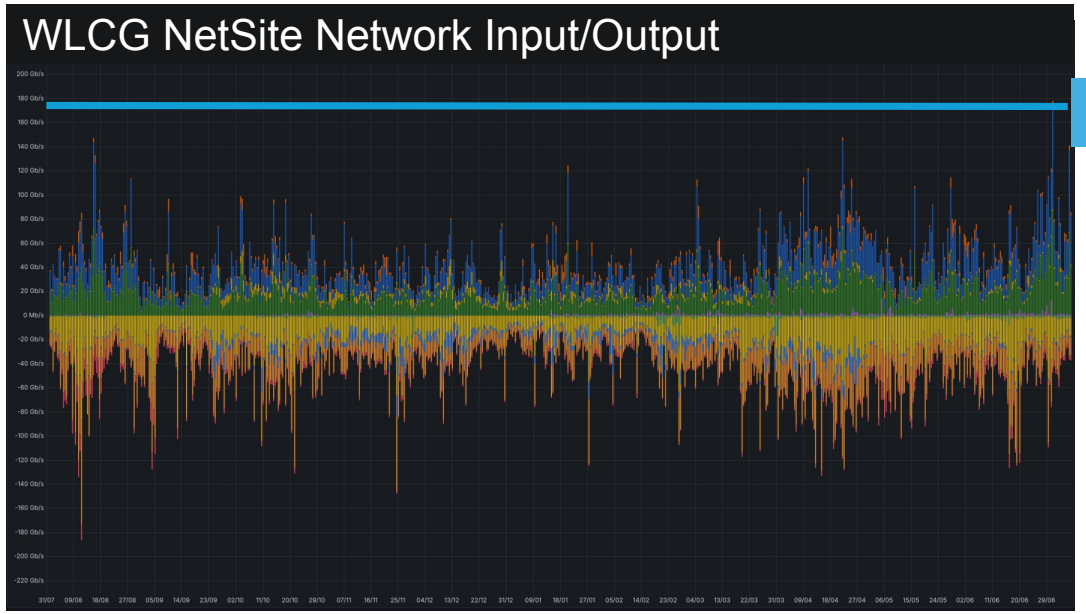


180 Gbps

| | | |
|---|---|---|
| IN: ES-UAM | Mean: 1.19 Gb/s | Max: 25.4 Gb/s |
| IN: ES-PIC | Mean: 20.0 Gb/s | Max: 85.8 Gb/s |
| IN: ES-IFIC | Mean: 4.63 Gb/s | Max: 27.7 Gb/s |
| IN: ES-IFAE | Mean: 20.0 Gb/s | Max: 85.8 Gb/s |
| IN: ES-CIEMAT | Mean: 3.15 Gb/s | Max: 15.9 Gb/s |
| OUT: ES-UAM | Mean: 775 Mb/s | Max: 14.1 Gb/s |
| OUT: ES-PIC | Mean: 18.8 Gb/s | Max: 86.2 Gb/s |
| OUT: ES-IFIC | Mean: 2.73 Gb/s | Max: 19.0 Gb/s |
| OUT: ES-IFAE | Mean: 18.8 Gb/s | Max: 86.2 Gb/s |
| OUT: ES-CIEMAT | Mean: 2.93 Gb/s | Max: 18.6 Gb/s |

ES Cloud A traffic snapshot (source)

# ATLAS R2R4 (Road to Run4) R&D Topics where ES-Cloud sites contribute

| Site | R2R4 interests | Comments |
|---|---|---|
| **Barcelona (PIC T1) & (IFAE T2)** | DC4, DC5, DC12 | Preparation Data Challenge 2027, HPC evolution, Sustainability in ATLAS Computing |
| **Madrid (UAM T2)** | DC4, DC5 | Scitags, HPC evolution |
| **Valencia (IFIC T2)** | DC5, DB2, DB4, DC12 | HPC evolution, CREST, Event Index, Sustainability in ATLAS Computing |

R2R4 stands for "Road to Run 4" - ATLAS Computing programme towards Run 4 (HL-LHC).
DC - Distributed Computing. DB - Database infrastructure, conditions and metadata (ADAM)

ATLAS R2R4 R&D list:
https://docs.google.com/spreadsheets/d/1p6RNWyrgEgV4_V7CI2YmmNz4TUMhShUIM7_jNMqhHV0/edit?gid=0#gid=0

# ATLAS R2R4 (Road to Run4) R&D Topics where ES-Cloud sites contribute

## R&D at IFIC: The ATLAS Event Index

- **IFIC** is **responsible** for the **Data Extractor**, **Supervisor** and the **Phoenix Loader Apps**.
- In numbers…
  - **+1306 billion events** processed.
  - **+614K datasets**
  - **+48 million files**

In a single HBase Table!

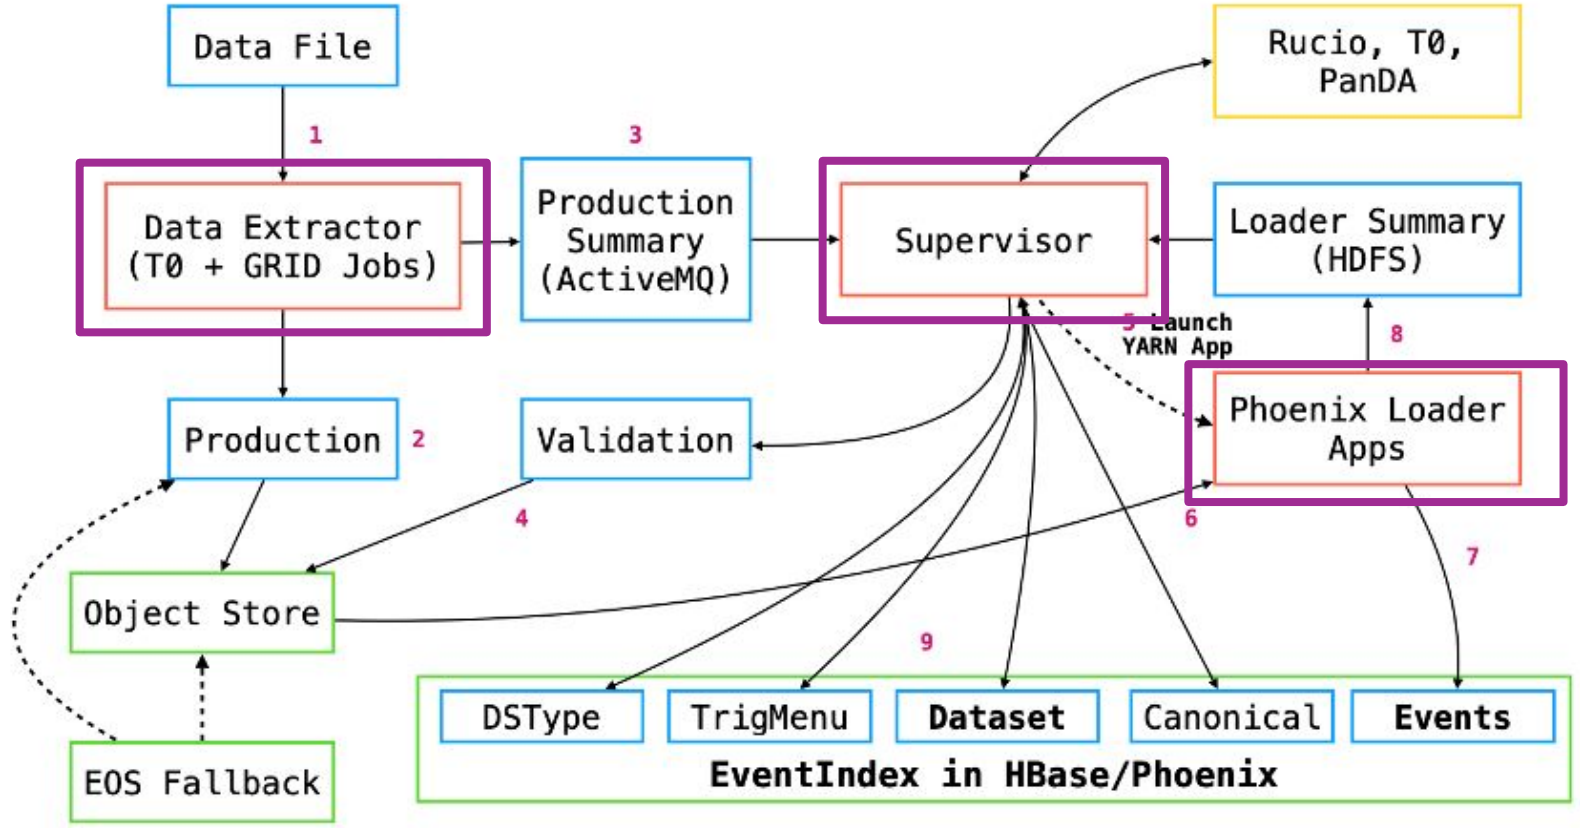


See Dario's presentation at CHEP 2026

# ATLAS R2R4 (Road to Run4) R&D Topics where ES-Cloud sites contribute

## UAM-LCG2: Scitags

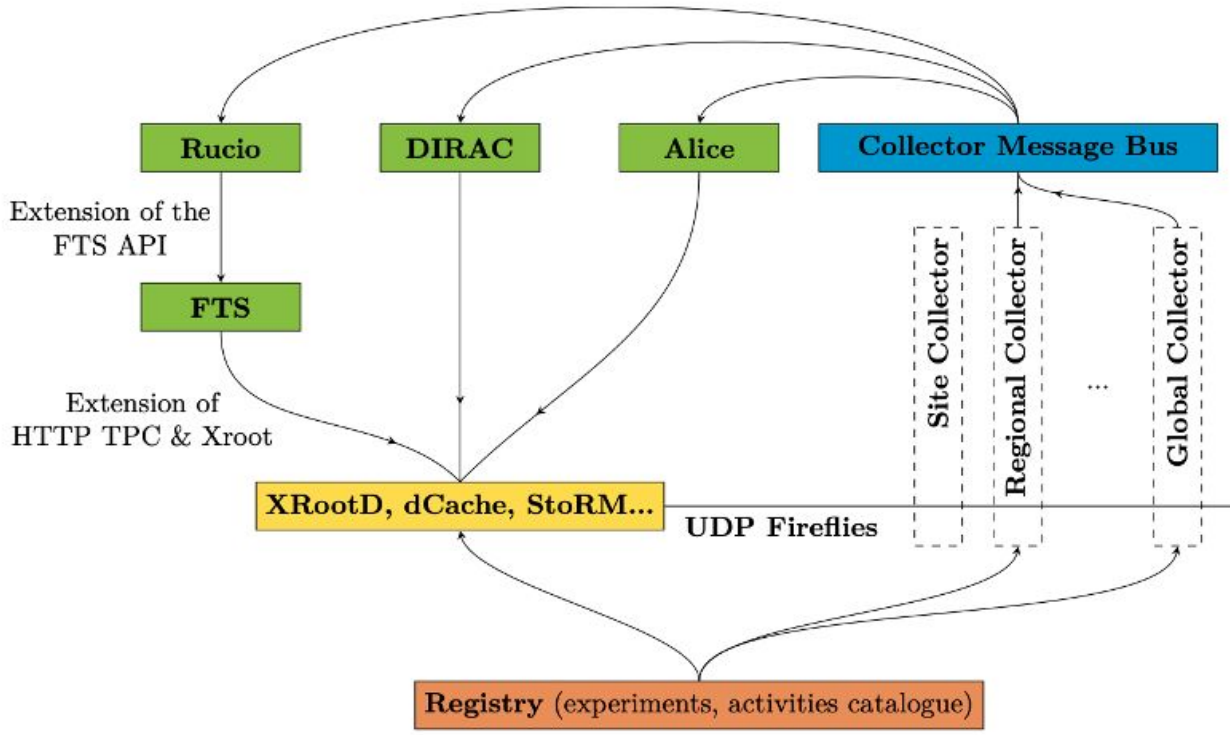


**SciTags** is an open framework that **enables the identification of data owners and the specific purpose of network traffic through standardized information** exchange between scientific communities and network operators. The standard implementation underpinning the SciTags framework, flowd-go, has been developed at **UAM.**

DC4 - Network Infrastructure

See Marian Babik's presentation at CHEP 2026.

Flowd-go repository

# ATLAS R2R4 (Road to Run4) R&D Topics where ES-Cloud sites contribute

DC12 - Going Greener

See Alejandro's presentation tomorrow

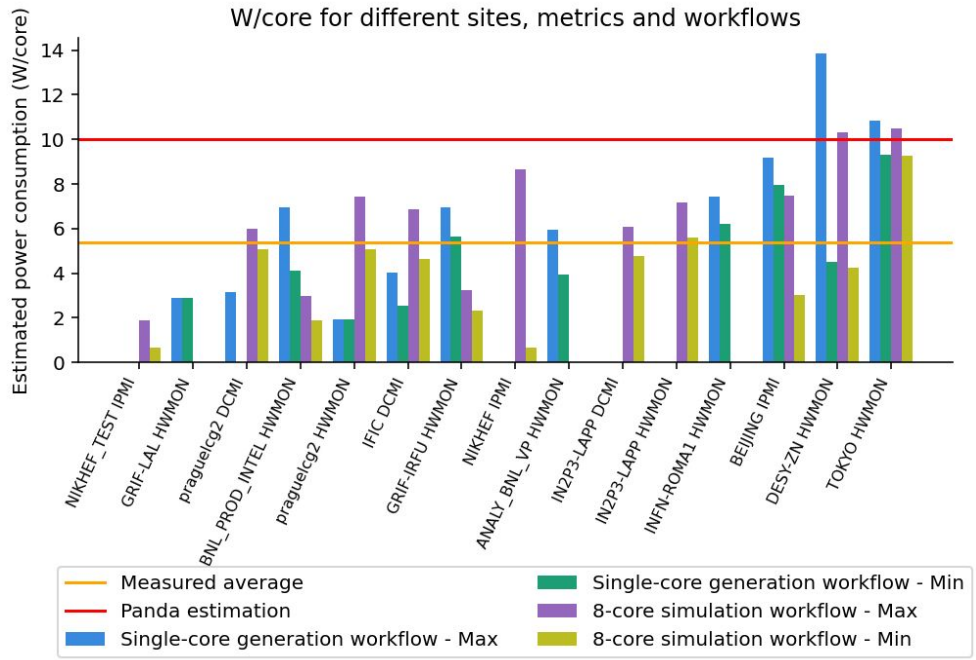


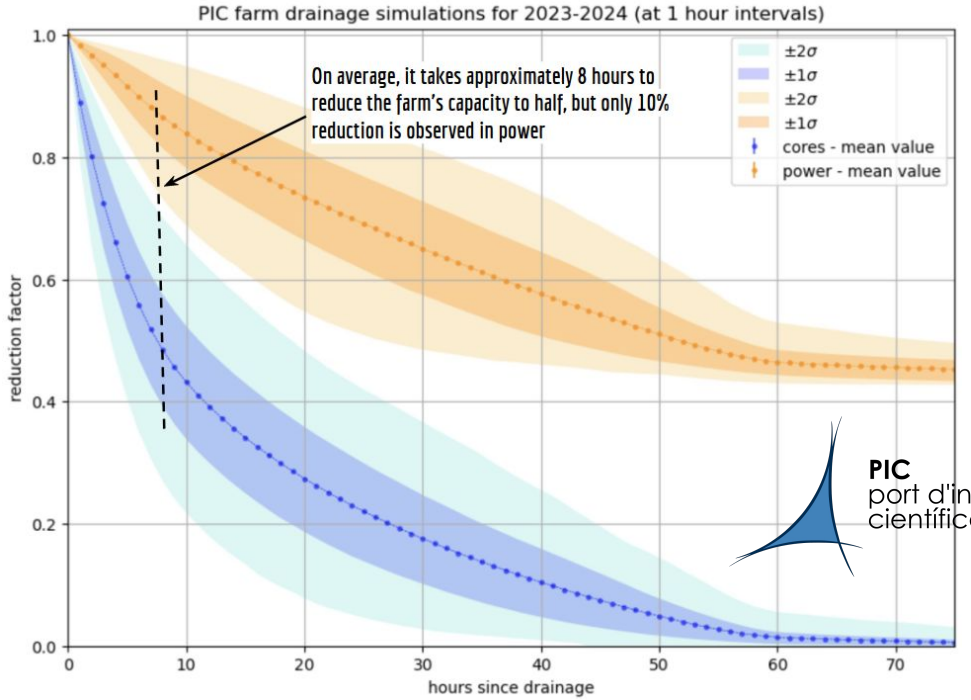


PIC port d'informació científica

An updated version of prmon together with Prometheus exporters installed on each site would notably improve current power consumption estimation.

These are simulations of how computing resources (active cores) at PIC behave during a drainage process, aimed at optimizing power consumption based on Artificial Intelligence predictions.

# Facility for Interactive Data Analysis (FIDA)

- **Guiding principle**: Help physicists minimize time-to-insight, enabling **iterative exploration** of the data.
  - In the future, when processing 10x lumi/evts, avoid physicists spending 10x time!!
  - Boost productivity and competitiveness of our physics communities.
- **Key features**:
  - **Local access** to the reduced data samples (e.g. DAOD_PHYSLITE) with low latency from compute.
  - **Processing resources** available with capacity and priority (CPUs, GPUs, Memory).
  - Efficient and **elastic infrastructure** (e.g scheduling based on HTCondor) with dynamic expansion to HPC/Cloud to absorb peaks.
  - **Enhanced support** of software tools (ROOT + Python ecosystem).
  - Enable and encourage **use of common repositories of code for analysis routines/workflows (e.g. analysis tree).**
  - **Expert data/code manager**: critical liaison role (not a final user, nor an infrastructure expert, however facilitating technology/access).

# Facility for Interactive Data Analysis (FIDA)

- **Typically consists of**:
  - Dedicated storage resources (on the order of several 100s of TB).
  - CPU resources used interactively and/or via a batch system (mostly HTCondor).
    - Future User Interface (UI) to be designed in a flexible way, with user-friendly interfaces that do not discourage users.
    - IFIC/IFAE: **Jupyter Notebook** instances can be spawned via a dedicated portal (kore.ific.uv.es and jupyter.pic.es).
  - SW delivery mostly via **CVMFS** with increasing presence of containers.
  - GPU resources available, but often not dedicated.
    - IFIC: **ARTEMISA** infrastructure (http://artemisa.ific.uv.es, artemisahub.ific.uv.es).
  - Network:
    - LAN of multiple 10 Gbps to support the intense data throughput.
    - WAN of 100 Gbps connectivity with the WLCG dedicated network for data lake access.
  - **Tier-3 estimate 10-20% of Tier-2 size resources + GPU usage.**

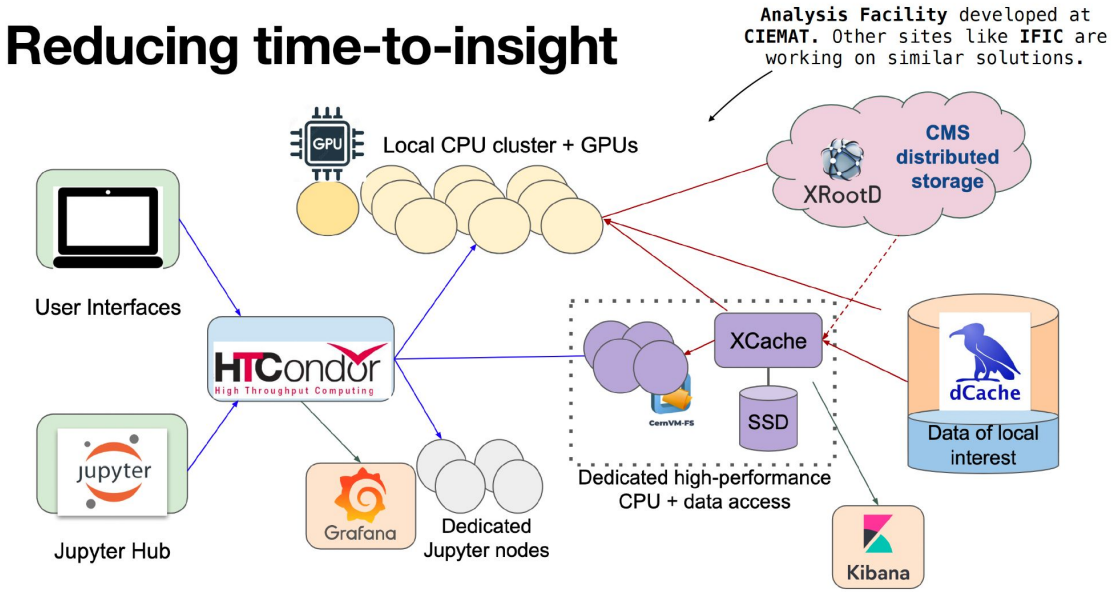


- **IFAE Tier-3**: Utilized 36M HS23 hours (4158 HS23) in 2025, 400 TB disk, and 200 TB tape.
- **IFIC Tier-3**: Local computing provided via Artemisa (GPU) and Gluon (CPU) clusters.
- External Resources: Significant use of CERN lxplus, Chicago Analysis Facility, and Mare Nostrum 5 (no official accounting available).

# Conclusions and Perspectives

- **Consistent and Excellent Performance:** For more than 21 years, the Spanish cloud has provided high-reliability computing, contributing **5% of global ATLAS resources** with an operational efficiency above 95%.
- **Robust Infrastructure:** Successful operation of the Spanish **Tier-1** and **Federated Nucleus Tier-2** and integration of **HPC resources (PIC, IFIC, UAM)** for large-scale simulations.
- **Network & Scalability:** Active ongoing improvements, deployment of **Analysis Facilities** to meet the high-bandwidth demands of the HL-LHC era.
- **Driving Innovation:** Continuous R&D contributions to optimize computing models and significantly reduce future resource requirements.

# Thanks to our national research projects


- **Funded by MICIU under contracts:**
  - **IFIC's Tier-2 project PID2022-136323OB-C21 of the MICIU**
  - **IFAE's Tier-2 project PID2022-142604OB-C22 of the MICIU**
  - **UAM's Tier-2 Project PID2022-136323OB-C22, of the MICIU**


# Acknowledgements


- Co-funded by the Recovery and Resilence Facility (Spain) with funding from the European Union NextGenerationEU (PRTR-C17.I1)
- Also funded by project RYC2022-038164-I




# Thanks

# Sustainability R&D Topics ES-Cloud contributions

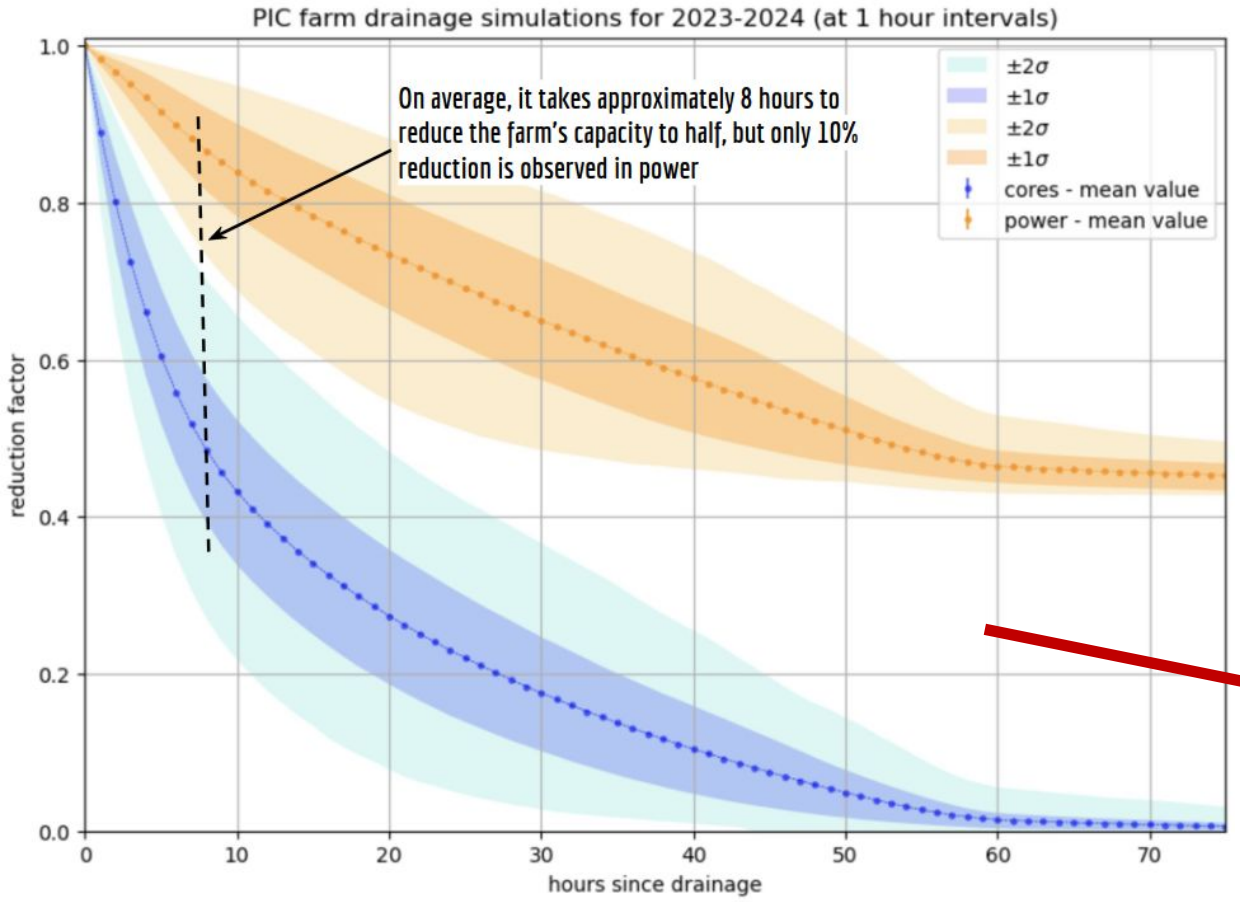


DC12 - Going Greener

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.

These are simulations of how computing resources (active cores) at PIC behave during a drainage process, aimed at optimizing power consumption based on Artificial Intelligence predictions.

# Sustainability R&D Topics ES-Cloud contributions


1. The availability of each of these metrics is dependent on implementation of the server manufacturer
2. Currently, these stats take a lot of time to plot due to high dimensionality (each status per core per node).
    1. We implemented another metric so that Prometheus sums per node automatically.

# About HS23

| CPU | Ncores | Score | Score/Ncores | Spread | HS23/W |
|---|---|---|---|---|---|
| filter | filter | filter | filter | filter | filter |
| AMD EPYC 9965 192-Core Processor | 768 | 14722.9 | 19.2 | 1.288 | 7.8 |
| AMD EPYC 9965 192-Core Processor | 768 | 13320.2 | 17.3 | 3.295 | 9.6 |
| AMD EPYC 9755 128-Core Processor | 512 | 11247.4 | 22.0 | 2.753 | 7.4 |
| AMD EPYC 9755 128-Core Processor | 512 | 11051.1 | 21.6 | 2.426 | 7.2 |
| AMD EPYC 9755 128-Core Processor | 512 | 11003.3 | 21.5 | 2.201 | 7.1 |
| AMD EPYC 9825 144-Core Processor | 576 | 9731.0 | 16.9 | 2.947 | 8.1 |

- Reference server:

  Intel® Xeon® Gold 6326 CPU @ 2.90 GHz (HT=On)