PIC
port d'informació
científica

Computing for Gravitational Waves in Spain

ET Spain meeting, Nov 5th 2026

G. Merino, PIC

IFAE EXCELENCIA SEVERO OCHOA

GOBIERNO DE ESPAÑA MINISTERIO DE CIENCIA E INNOVACIÓN

Ciemat Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas

ICTS INFRAESTRUCTURAS CIENTÍFICAS Y TÉCNICAS SINGULARES

Computing for GW - at a glance

IGWN (LVK)

- **Resources**
  - PIC & CIEMAT contributing to IGWN computing infrastructure since 2019
    - PIC: IGWN Grid site providing opportunistic CPU and GPU, with a local disk Cache for efficient data access.
    - CIEMAT: former Virgo low-latency cluster (500 cores), now reinstalled as an offline Grid site.
  - No dedicated project funding — contributions come from institutional resources.
  - Discussions ongoing with Virgo to formalize these contributions in an MoU.

- **Personnel**
  - IFAE providing 0.5 FTE for IGWN Computing Operations (Jam Sadiq)

Einstein Telescope

- PIC & BSC-CNS members of the ET e-Infrastructure Board (EIB)
  - PIC has led the Computing Model & Resources EIB Division 3 (2022-2026).
    - Plan to continue active work in Div3. Focus on actual infrastructure deployment.
  - BSC-CNS leads the computing WP in the ET-PP EU project (WP8).

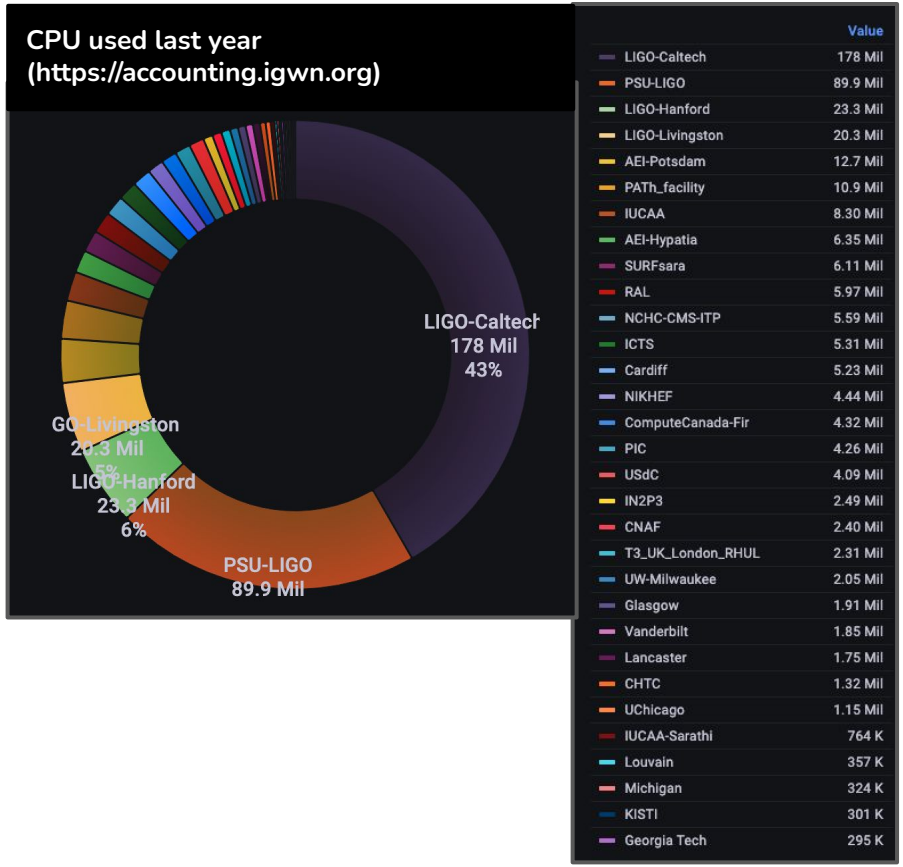
PIC in the IGWN Grid - CPU

PIC is providing useful CPU & GPU cycles to IGWN from day one, even if access is just opportunistic.

Leveraging the common software stack (and expertise) with other major experiments: LHC, CTAO, DUNE …

Consistently providing over 4 million core-hours per year.

- Comparable to NIKHEF, and ahead of CC-IN2P3 and CNAF.
  - These Virgo centres probably do work outside IGWN accounting.



PIC in the IGWN Grid - GPU

In the last 12 months, PIC stands out as the #1 GPU provider for IGWN.

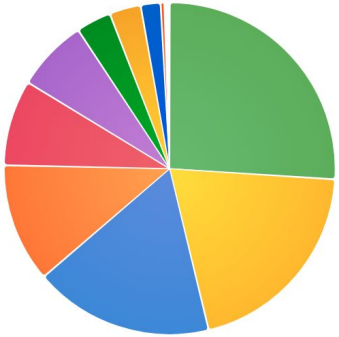
GW is providing a constant stream of GPU workloads which are able to efficiently use idle GPU cycles.

Over 150,000 GPU hours in the last year.

Pelican Cache serving h(t) at >3 Gbps yearly average

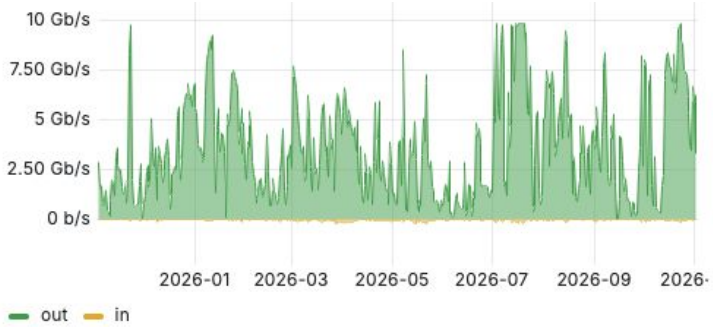
Over 10PB/yr of data served

GPU used last year (from https://gracc.opensciencegrid.org)



pic	157 K
ND-CAML-GPU-CE1	122 K
SURFsara	106 K
SDSC-PRP-OSPool-Provisioner	70 K
ComputeCanada-Fir	50 K
PATh_facility-CE1	42 K
LIGO-CIT-CE	20 K
Georgia_Tech_PACE_CE_LIGO	18 K
AEI-Hypatia-CE1	11 K
LSU-SuperMIC-CE1	2 K
LIGO-CIT-CE1	513
OSG_IN_ICTS_ALICE_HUMPTY	162
SU-ITS-CE2	11
SU-ITS-CE3	1
N/A	0

Network traffic of the Pelican Cache at PIC



Personnel contribution to IGWN Computing

PIC
port d'informació
científica

IFAE contributes 0.5 FTE to IGWN Computing Operations (Jam Sadiq):

- **HTCondor support for IGWN** (main focus): helps users diagnose and fix job issues that run gravitational-wave analyses on shared resources.
- **System upkeep and upgrades:** Maintains the monitoring dashboard and documentation, and coordinates upgrades with developers.
- **PyCBC liaison:** supports European users with HTCondor issues in PyCBC workflows and raises them with the developers.
- **IGWN computing helpdesk:** handles tickets and supports the wider IGWN computing team.
- **Open Science Grid:** joins biweekly meetings with other organizations that use HTCondor to discuss future developments.

Computing for the Einstein Telescope

The Challenge: The big difference with respect to 2nd Gen interferometers is the increased event rate, thanks to the much improved sensitivity.

- ~1000× more detections than 2G detectors: >100,000 BBH plus >10,000 of BNS/NSBH per year.
- Data volume (~10 PB/year per interferometer) is large for GW but manageable by HEP standards; the real challenge is **compute**.
- Focus on software **sustainability**: being able to reliably run for decades with small, stable effort.

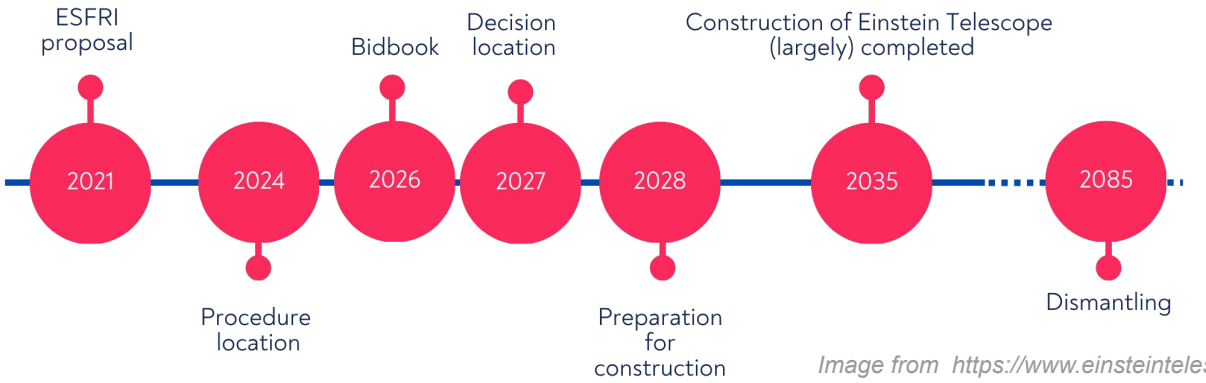


Image from https://www.einsteintelescope-emr.eu

Computing for the Einstein Telescope

ET E-Infrastructure Board (EIB) mandate

"...to design, create and operate an evolving, efficient and functional e-infrastructure environment at a reasonable cost for the collaboration. Initially the focus will be the development of a Computing Model for the ET".

Organization in four divisions plus a Technology Tracking Group

- Division 1: Software, frameworks, and data challenge support
- Division 2: Services and Collaboration Support
- Division 3: Computing and data model, Resource Estimation
- Division 4: Multimessenger alerts infrastructure
- TTG: Technology Tracking Working Group

Coordination of EIB Div3 and ET-PP WP8

EIB Div3 scope (led by PIC)

- Develop the Einstein Telescope Computing Model
  - Overall architecture of the e-Infrastructure, either as a single integrated system or as a few separate systems (online, low-latency, and offline)
  - Collaborate with OSB and define a process to evaluate computing needs from the evolving scientific program of the collaboration.
- Provide a running estimate of the computing resources needed for all computing domains.

ET-PP WP 8 Computing and Data Access (led by BSC)

- Goal: deliver ET's computing and data access model: workflows, resource estimates, online Tier-0 design, and data storage and access policy.

Deliverables

- D8.1: Computing and Data Requirements
- D8.2: Computing and Data Model
- D8.3: Data Access Implementation Guidelines

Computing Model

D8.1 and D8.2 documents provide the baseline for the resource requirements: hardware and personnel.

- On-site: 2,150–6,450 cores and 10–20 PB per interferometer
  - 2 × 100 Gb/s network per site
- Rapid alerts: ~4 MHS06
- Offline: ~0.6 MHS06
- 2 × 10 PB/year custodial storage per interferometer
- ~50 FTE for S&C operations

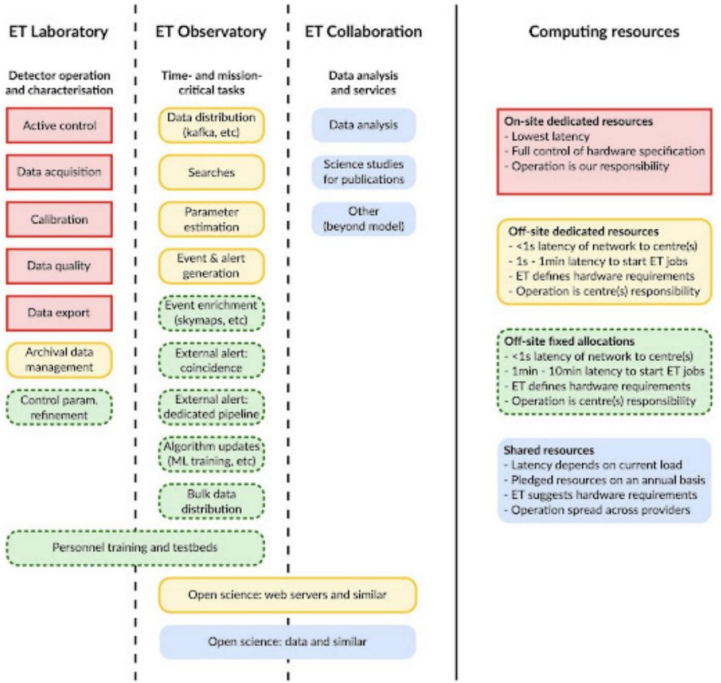


Figure 8.1: ET computing tasks and their computing resource categories.

Work Breakdown Structure

The Computing Model enabled us to define a work breakdown structure for ET computing.

It was useful for this year's costing exercise that was part of the GCR, and will also be useful to guide the organization of the work from now on.

An interesting outcome of the GCR costing exercise: Computing represents

- a negligible contribution to the construction cost (~1%)
- a sizable contribution to the operations FTE cost (>30%)

Einstein Telescope Computing – Product Breakdown Structure

PBS Code	System / Sub-system / Component
1	**On-Site Computing**
1.1	**DAQ, Detector Control, Calibration and Data Quality**
1.1.1	On-site compute nodes
1.1.2	On-site storage
1.2	**On-Site Computing Operations**
1.2.1	Database DevOps
1.2.2	On-site computing personnel (per site)
2	**Rapid GW Alert System**
2.1	**Rapid Alert compute cluster**
2.1.1	CPU/GPU compute nodes
2.2	**Rapid Alert System Operations**
2.2.1	GW event & alert services DevOps team
3	**Data Management & Storage**
3.1	**Raw data archival distributed copy**
3.1.1	Tape storage pledges at T1 computing centres
3.2	**Data Management Operations**
3.2.1	Data Archival DevOps team
3.2.2	Data Streaming service DevOps team
3.2.3	Distributed data service DevOps team
4	**Distributed Computing Services & Middleware**
4.1	**Services Cloud System**
4.1.1	Cluster nodes for Services Cloud
4.2	**Middleware & Infrastructure Operations**
4.2.1	Core infrastructure team
4.2.2	WMS; accounting & monitoring DevOps team
4.2.3	Coordination team
5	**Software Infrastructure & Engineering**
5.1	**Software development & CI/CD Infrastructure**
5.1.1	Software development and version control platform (GitLab/GitHub)
5.1.2	CI/CD pipeline & software packaging infrastructure

Estimation of ET Computing Needs

The EIB launched a process in 2025 to gather requirements from OSB, ISB, SCB. First estimates were collected, along with info on emerging SW/HW R&D technologies.

Numbers show large variation and uncertainty. Follow-up will be needed to sharpen the picture.

Findings

- Members with local/LVK access already covered. Priority should go to those with none.
- Need an ET-specific environment similar to IGWN (HTCondor-based, proven robust).
- Growing GPU demand needs dedicated planning.
- Final needs will depend on whether an MDC with Continuous Wave signals is done and how large a parameter space is explored (adjustable to available resources).
- A dedicated GitLab CI/CD for sw development and testing + container registry to build images is an identified need.
- Tutorials/documentation is needed for new users.

Next steps:

- Developing formal mechanisms for resource pledging, institutional participation, and establishing a recognition scheme for contributions to the distributed infrastructure.

Mock Data Challenges

The current estimate of computing power needs relies heavily on assumptions about future algorithmic speed-ups, such as leveraging ML techniques and GPUs.

- MDCs will be necessary to quantify the impact of these technological improvements and refine the estimates, as well as guide the evolution of the software framework.

To ensure this feedback is effectively implemented, a governance of the MDC with participation of the OSB and the EIB has been proposed (see ET-0235A-26).

Another example of effective use of common infrastructure at PIC:

- Sep'26: Urgent need for MDC code "gwmock" review and dataset production.
- Swift reaction from PIC with ready-to-go functionality and resources.
  - Interactive notebooks and batch access to CPU & GPU + 20 TB disk
  - Controlled access via ET Indigo IAM
    - 6 people (reviewers/reviewees) working at PIC to generate and review the MDC data set

ET Data Lake development

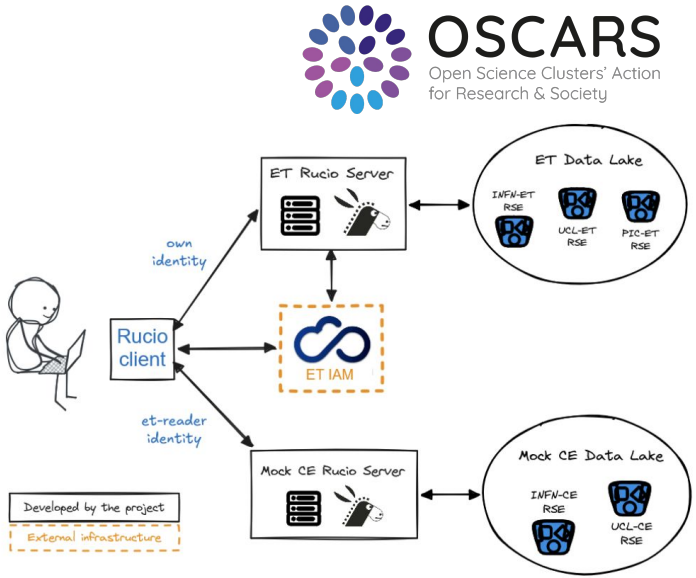
MADDEN project funded by the OSCARS HE project. Coordinated by INFN and UCLouvain.

Goals:

- Showcase Rucio for ET data management
- Build a multi-RI Rucio Data Lake
- Provide POSIX view of Rucio catalog
- Rucio “mini data lake” for Virgo as technology demonstrator for LVK.

PIC provides a storage endpoint for the Data Lake demonstrator.

- Will also use this for MDC data distribution.



Summary

Spain contributes actively to GW computing, both in IGWN today and in ET for the future

- Our approach builds on distributed computing expertise and on neighbouring communities: HEP and gamma-ray astronomy.
- The activity is largely unfunded. The only dedicated effort is 1 FTE through ET-PP WP8 at BSC and 0.5 FTE for IGWN at IFAE.

Grid: from IGWN to ET

- Through PIC, Spain already provides a sizable share of IGWN computing: >4 M CPU hours/yr, and #1 GPU provider last year.
- ET is now setting up a similar distributed infrastructure, and Spain is well placed to contribute cycles to it.
- Opportunity to steer the infrastructure development and policies, such as the pledging process and the incentive scheme.
- Turning this into reality requires minimal dedicated personnel funding (1-2 FTE).

HPC: RES as an additional resource

- The Spanish Supercomputing Network (RES) includes MN5 at BSC and several mid-sized HPCs with substantial CPU & GPU.
- Access today is only through competitive calls with 6-month allocations, which makes it hard to plan a long-term programme.
- Proposal: present the case to the Ministry to recognise ET as a "Strategic RES project". This would guarantee a CPU-hour quota (not so clear for GPU).