PROCEEDINGS OF SCIENCE

Data and Computing Infrastructure for KM3NeT-Towards Full Deployment

Riccardo Bruno[a] and Francisco Vazquez de Sola[b] on behalf of the KM3NeT Collaboration.

a INFN Dept. Catania,
Via Santa Sofia 64, Catania, Italy

b Nikhef,
Science Park 105, 1098 XG, Amsterdam, NETHERLANDS

E-mail: riccardo.bruno@ct.infn.it, 14favd@queensu.ca

The KM3NeT Collaboration is building two Cherenkov neutrino detectors in the depths of the Mediterranean Sea to study both the intrinsic properties of neutrinos and cosmic high-energy neutrino sources. For the detection of Cherenkov light generated by charged particles from neutrino induced interactions the detectors share the same technology of large arrays of optical modules. The expected data volume grows roughly linearly with the number of optical modules, reaching hundreds of terabytes per year at full deployment. To handle this load, a Rucio file catalogue will manage distributed data storage and replication, and the DIRAC workload manager will orchestrate large-scale data processing and Monte Carlo simulations. This infrastructure allows the collaboration to scale efficiently with detector size while supporting real-time data acquisition, processing, and analysis workflows. This contribution focuses on the KM3NeT data and computing infrastructure, describing the current computing model and its evolution toward distributed Grid-based resources as the experiment approaches full deployment.

International Symposium on Grids and Clouds (ISGC2026)
15-20, March, 2026
Academia Sinica, Taipei, Taiwan

© Copyright owned by the author(s) under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0)
All rights for text and data mining, AI training, and similar technologies for commercial purposes, are reserved.
ISSN 1824-8039 . Published by SISSA Medialab.

PoS(ISGC2026)018

PoS(ISGC2026)018

1. Introduction

Neutrinos are elementary particles with almost no mass and no electric charge, interacting only via the weak force. This allows them to travel across intergalactic distances without being deflected by gravitational or electromagnetic fields, and with marginal probability of interaction. On the other hand, this makes their detection extremely challenging, as it relies on rare interactions with matter. When such an interaction occurs, neutrinos can produce charged secondary particles that release energy in the surrounding medium in the form of photons. This light emission, known as Cherenkov radiation, takes place when a charged particle moves faster than the speed of light in the medium. Neutrino detectors are therefore based on highly sensitive optical sensors deployed in large, transparent media such as deep water or ice, where a dark environment and large instrumented volumes increase the probability of detecting neutrino interactions.

KM3NeT [1] is a large European research infrastructure with two neutrino detectors being built at separate locations in the deep waters of the Mediterranean Sea. Its ARCA and ORCA detectors are optimised for neutrino astronomy and the study of the fundamental neutrino properties, respectively. The infrastructure is developed by an international collaboration involving dozens of institutes worldwide.

In section 2 the KM3NeT detectors are introduced; section 3 the data acquisition is described; section 4 the transition to a grid-based system is presented. Finally, section 5 comprises conclusions.

2. The KM3NeT detectors

The ARCA and ORCA detectors of KM3NeT in the Mediterranean Sea consist each of a three-dimensional array of optical modules, as illustrated in Figure 1.

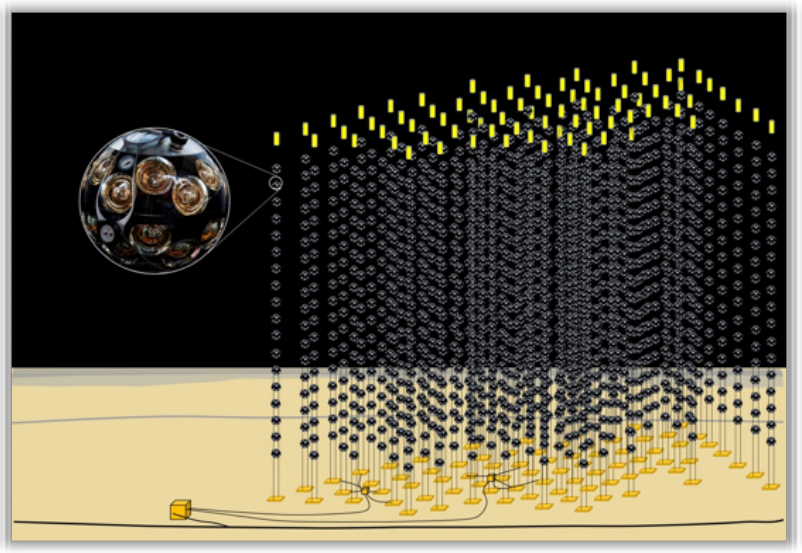

***Figure 1**: Artist's impression of a KM3NeT three-dimensional array of optical modules. Image courtesy of the KM3NeT Collaboration*

The detectors are designed to operate for at least ten years after completion and to instrument a total volume of one cubic kilometre of sea water. Each detector consists of vertical detection units of optical modules. The architecture of the system allows the detection units to be installed and operated individually, enabling data taking since the deployment of the first units, with detector efficiency improving as the number of operational units increases over time.

2.1 Detector building challenges

Neutrino detectors operating in sea water must address several critical issues. For the KM3NeT detectors, the following challenges have been taken into account.

Detected light is subject to background contributions that, for the KM3NeT detectors, come from muons generated in the collisions of cosmic rays with molecules in the atmosphere. In addition, the decay of potassium in sea water and bioluminescence also generate light. Filtering algorithms are employed to reduce this background noise. Data collected by the optical modules must be transmitted over long distances over several tens of kilometers. Bandwidth must be maximised, and transmission redundancies must be considered.

Also the engineering is challenging. The vast dimensions of the detector require dedicated working environments and devices, and pose difficulties in transmitting power and data over long distances. The optical modules, deployed in the deep sea, are subject to high hydrostatic pressure and chemical aggression from sea water. Since no deployed component can be actively maintained, high-quality components are required to guarantee 10–20 years of reliable operation. Furthermore, deployments rely on costly sea campaigns to be conducted by highly specialised companies.

Finally, the apparatus has to be designed with solutions adequate for preventing any impact on the biological environment, including a decommissioning plan at the end of the detector's lifetime.

2.2 Detector components

While the core component of the KM3NeT detectors are the Digital Optical Modules (DOMs) [2], the core component of a single optical module is its photosensor, called a photomultiplier tube (PMT). The PMT exploits the photoelectric effect to generate a cascade of electrons when a photon strikes the sensor surface. This electron cascade is measured by a dedicated electronic board and converted into a digital signal. The complete optical module comprises 31 PMTs mounted inside a pressure-resistant glass sphere. The sphere is composed of two halves (top and bottom), housing 12 and 19 PMTs, respectively. The DOM also hosts the electro-optical connectors, the Central Logic Board (CLB), an acoustic sensor, and a light emitter used for spatial positioning retrieval and time synchronisation. The CLB additionally includes a compass sensor to provide tilt information for the optical module positioning, as well as other sensors for monitoring voltage, humidity, and temperature.

Eighteen DOMs form a Detection Unit (DU). They are kept in their positions by means of two ropes attached to an anchor on the sea floor. The DOMs are connected through an electro-optical backbone cable to a collector module on the anchor called a Base Module (BM), which is connected to the seabed backbone network that finally links to the control station on shore where all detector data are processed.

PoS(ISGC2026)018

The detection units are comprised in sets of 115 units, called building blocks (BBs).

2.3 The ARCA detector

The Astroparticle Research with Cosmics in the Abyss (ARCA) detector is optimised to detect high-energy cosmic neutrinos, in the range from the TeV up to the PeV scale and beyond. It is primarily an astronomical telescope covering nearly one cubic kilometre of deep sea water, located about 100 km offshore Capo Passero, on the south-eastern coast of Sicily, Italy. The detector will consist of two building blocks (230 Detection Units in total). The DUs are anchored to the seabed at a depth of approximately 3,500 meters, interconnected both optically and electrically in a star-like topology submarine network.

The ARCA telescope is a key player in multi-messenger astronomy, in which different types of signals are studied together to provide complementary views of the same astrophysical phenomena. Beyond neutrino astronomy, ARCA also functions as a deep-sea observatory, enabling research in marine biology based on the measurements of the acoustic sensors installed in the DOMs. The detector's multidisciplinary potential also extends to Earth sciences, for instance through measurements of optical signals propagating along the long-distance submarine cable for earthquake detection.

2.4 The ORCA detector

The Oscillation Research with Cosmics in the Abyss (ORCA) detector is located approximately 40 km offshore from Toulon in France, at a depth of about 2,500 metres in the Mediterranean Sea. ORCA instruments a smaller volume of sea water than ARCA, featuring a much denser arrangement of optical modules. The detector will consist of one building block, corresponding to 115 Detection Units (DUs) in total. The DUs are anchored to the seabed and interconnected both optically and electrically in a tree-like topology network. The compact and dense geometry is optimised for the detection of atmospheric neutrinos in the GeV range. The reduced spacing between optical modules enhances the reconstruction capabilities, improves the sensitivity to neutrino oscillations, and enables the determination of the neutrino mass ordering.

3. Data acquisition and flow

As described above, the modular design of the detectors enabled KM3NeT to start data taking with ORCA in 2019, when four DUs were operational, and with ARCA in 2021, when six DUs were operational. This section describes the data flow from the raw detector output to higher-level analysis files.

3.1 All data to shore

The light which reaches the PMTs of the DOMs is detected through a configurable discrimination stage. The original analogue signal from each photomultiplier is converted into a digital signal called a hit. No hit selection is performed offshore. An FPGA on board the Central Logic Board (CLB) of each DOM sends the hits via UDP datagrams, which are transmitted through submarine optical fibres to the control station on the shore, where the datagrams are collected by the Data Queue, which organises the incoming information into time slices of 100 ms each. A Data Queue process then distributes each time slice, via TCP packets, to a configurable number of Data Filters, each responsible for processing a single time slice. Each Data Filter applies the triggering algorithms and forwards the selected events to the Data Writer module,

PoS(ISGC2026)018

where the data are stored as raw files. This step reduces the original data volume by a factor of $O(10^3)$.

The raw files are grouped into run data files, each covering approximately three hours of data taking, and are collected and stored securely on a computer centre outside the shore station via a daily file copy process.

This data acquisition and storage process is a distinctive feature of the KM3NeT. It provides a dynamic and configurable architecture for applying triggering, and offers the opportunity to discriminate against different sources of noise [3].

3.2 The computing model

The computing infrastructure of the KM3NeT project is organised into three distinct tiered levels, following the same hierarchy established for the WLCG infrastructure built for the LHC experiments [4]. This structure has been formally described in the Review of the KM3NeT Data Management Plan (Deliverable D4.1) [5] of the European INFRADEV2 project [6], which focuses on the data management plan. The overall architecture is illustrated in Figure 2.

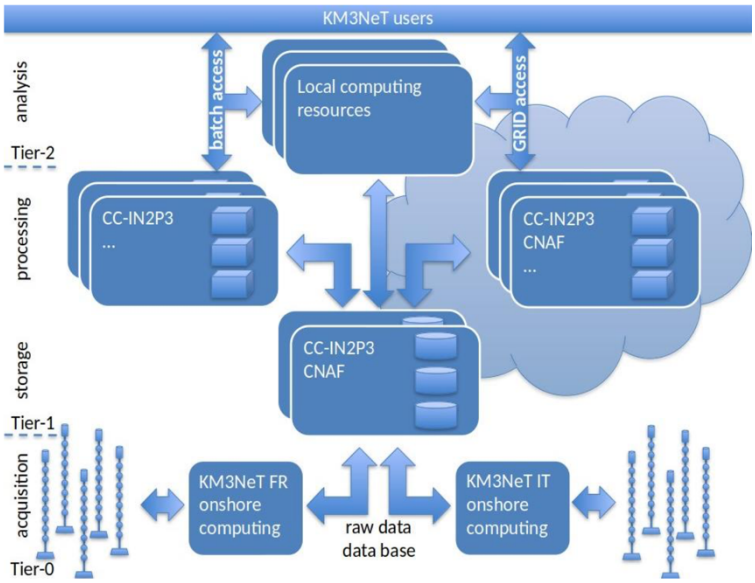

Figure 2: *The KM3NeT Computing model, with the three tiers, reproduced from Ref. [5]*

From the bottom to the top of the figure, the Tier 0 corresponds to the detector level. At this stage, hits originating from the instrumentation are collected at the shore stations using the "All data to shore" strategy. The hits, organised into time slices as explained in the previous section, are stored as filtered data (raw files) covering three hours of detector data acquisition. These raw files are then transferred to the Tier 1 level via a daily file transfer operation. Each raw file is replicated across two different sites for security purposes.

At the Tier 1 level, raw files are processed as:

- **Calibration**, to retrieve true detector geometries;
- **Reconstruction**, to calculate event trajectories by applying different algorithms;
- **Simulation**, to model the detector behaviour and enable the verification of the detector efficiency.

At the top of the hierarchy, the Tier 2 level is dedicated to physics analysis. Here, stored reconstructed data, organised into summary files, are used to obtain meaningful scientific results.

PoS(ISGC2026)018

The core of the computing infrastructure comprises two main Computing Centres (CCs): CNAF in Italy and CC-IN2P3 in France. Both centres provide disk and tape storage, computing resources, and user interface access. Beyond these primary centres, other CCs are also involved, primarily for mass data processing and analysis, though they are generally not open to all collaboration members. CC-IN2P3 hosts a database that not only contains generic experiment data but also detector sensor data, such as acoustic data and internal compass data, which are essential for calibration [7]. This information is required at a processing stage, and the results are stored inside a Git repository at ECAP (Germany). Other CCs provide their resources mainly for internal institutional use and for performing analyses. The database also provides an Identity Provider service named KM3SAML and APIs to dynamically retrieve detector information. Table 1 lists the computing centres currently involved in the KM3NeT computing infrastructure.

Table 1: *KM3NeT computing centres and corresponding tier levels*

Tier	Name	Country
1	IN2P3 (CCLyon)	France
1	CNAF (INFN Tier 1)	Italy
2	Nikhef (Stoomboot, ELPROD)	Netherlands
2	ECAP (Woody)	Germany
2	AGH (Cyfronet)	Poland
2	IFIC (Gluon, Artemis)	Spain
2	HULL (Viper)	United Kingdom
2	Demokritos	Greece

3.3 Data processing

The data processing in KM3NeT consists of several steps, as represented by the simplified structure in Figure 3.

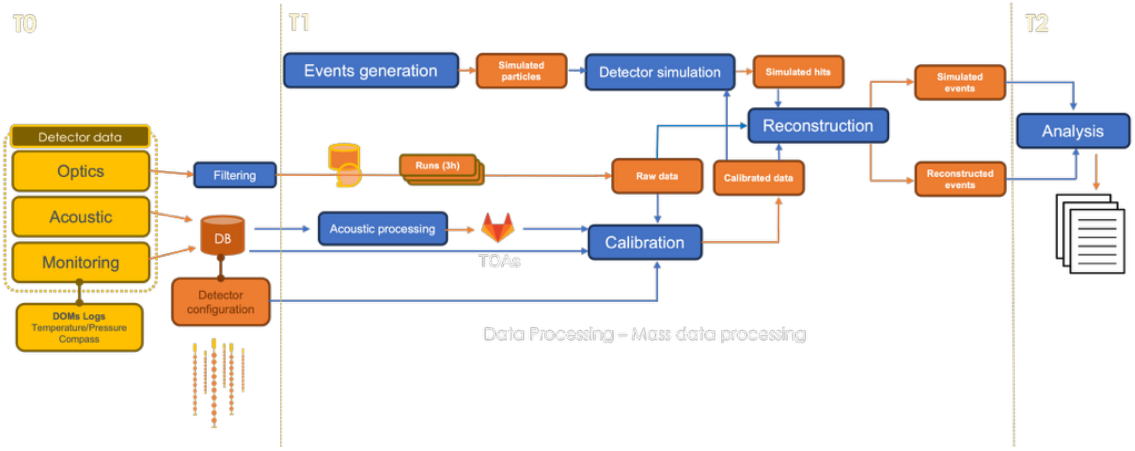

Figure 3: *Simplified representation of the KM3NeT data-processing workflow. Processing steps are shown in blue and data flows in orange*

First, the detector configuration is loaded, including the geometry (DU and DOM positions, etc.). Acoustic and tilt sensor logs are analysed to determine the precise position and orientation of each optical module at a given time. This information is used to calibrate the detector.

In the second step, physics events are reconstructed from the filtered raw data. Accurate calibration is essential at this stage to ensure reliable event reconstruction.

PoS(ISGC2026)018

The third phase consists of Monte Carlo simulations. Using the same calibration data, the interactions of atmospheric muons and neutrinos are simulated, together with the detector response. This serves a dual purpose: monitoring detector performance and comparing observed events with theoretical expectations.

Recently, the full workflow has been integrated into a single Snakemake-based pipeline, replacing the previous script-by-script approach. This improves automation, reproducibility, and consistency of the processing chain. However, this design is naturally optimised for execution on local clusters. When moving to distributed Grid infrastructures, several limitations emerge. The current workflow assumes direct access to local storage via direct storage streaming or file transfers, and it is not yet fully optimised for fine-grained job distribution across heterogeneous Grid resources [8]. A Grid-compatible prototype has been developed in parallel, starting from the original Snakemake project, and it was successfully tested, but it is not yet in production. The main challenges include:

- Two different software projects on the project software repository;
- Version mismatch between the production Snakemake workflow and the Grid prototype components;
- The need for multi-core worker nodes, which reduces scheduling flexibility;
- Limited parallelisation due to suboptimal data distribution.

Improved scalability is expected through tighter integration with Rucio [9] for data management, and potentially through the development of a dedicated workflow plugin. Additional challenges are related to the limited Grid experience within the collaboration and the ongoing transitions in Grid infrastructures, such as the migration from X.509 certificates to token-based authentication and the recent decommissioning of the Sectigo CA [10].

4. The Grid transition

To support the growth of the detectors and the corresponding increase in the amount of data over the ten-year lifetime once the detectors are completed, a huge volume of data must be managed. As depicted in Figure 4, the two detectors will need to handle O(PB) of data per year, a quantity that cannot be easily supported by a few data centres alone and requires distributed data solutions.

PoS(ISGC2026)018

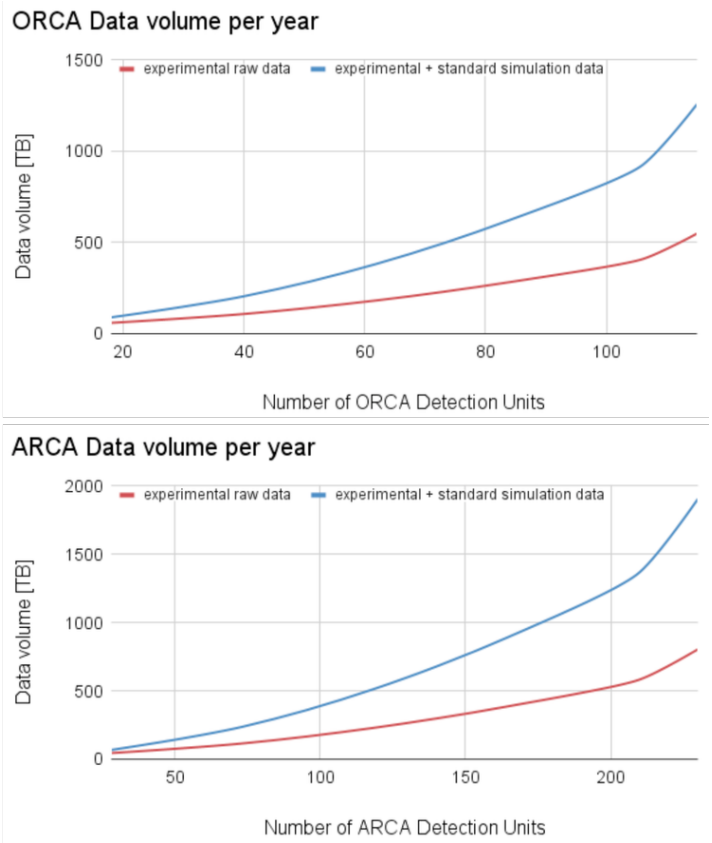

Figure 4: *Estimated KM3NeT data-volume requirements as a function of the number of deployed DUs, reproduced from Ref. [5]*

This conclusion is consistent with the outcomes of the INFRADEV2 project [6, 11] and with the technical evaluation performed by the independent Scientific and Technical Advisory Committee (STAC).

For this reason, the first step towards a distributed computing infrastructure is the adoption of Rucio, together with the identification of additional Tier-1 storage elements.

The Rucio service is already in place and will soon be integrated into the run-based data processing workflow. Rucio will then become the central service for accessing detector files within the collaboration. This implies that each collaboration member must access Rucio using either tokens or legacy X.509 certificates. Rucio will also be served via Science Gateways and Open Science data repositories.

The second step consists of merging the two separate Snakemake [12] projects into a single, harmonised solution, with the goal of making Grid-based data processing the main production system. This step also involves optimising the current approach by distributing jobs across less demanding resources and improving the number of parallel jobs. This can be further achieved by integrating Rucio directly into the Snakemake process.

4.1 EGI.eu

Grid resources in KM3NeT have been used since 2013, although it was limited to isolated initiatives by individual researchers. A more structured adoption of Grid computing has been encouraged since 2023 within the European INFRADEV2 project, recognising the continuous growth in data volume over time. An official Memorandum of Understanding (MoU) with EGI.eu was recently signed, and the configuration of operational services such as the Grid Operation

PoS(ISGC2026)018

Center (GOCDB), the user support unit, and related components is still ongoing. The current Grid sites serving the Virtual Organisation (VO) km3net.org are:

- SURF-SARA / SARA MATRIX (Netherlands),
- ELPROD (Netherlands),
- CNAF Tier-1 (Italy),
- IN2P3 (France).

EGI.eu provides its DIRAC instance for Grid data processing and CVMFS hosted at RAL (UK). CERN currently hosts the FTS service used by Rucio, which itself is hosted at Nikhef. In addition, an IAM server hosted at CNAF provides authentication and authorisation, supporting both legacy VOMS proxies and most modern token-based access.

Several issues remain: some have already been addressed, while others are still under resolution. The user community has consolidated procedures that are tightly coupled to local computing clusters, especially at IN2P3. Moreover, the current level of expertise is not yet sufficient to widely promote Grid usage, nor to establish a core community team that could support the collaboration in user-oriented activities such as training, troubleshooting, and assisting members in writing and adapting their application code for Grid environments.

Over the past year, the decommissioning of Sectigo certificates had to be managed and the migration to the HARICA Certification Authority (CA) was performed, which was not immediately available in all countries at the same time.

Since Rucio will become the main data management system of the project, token-based authentication is essential. However, not all Rucio Storage Elements (RSEs) currently support tokens.

Finally, not all KM3NeT members are yet linked to the IAM service. This will become possible once the KM3SAML identity provider is fully integrated.

4.2 Rucio, DIRAC, and other services

Rucio will be one of the most important services for the collaboration. Its installation started in 2023 within the European project INFRADEV2, initially operating as a prototype. Since late 2025, concrete actions are ongoing to deploy Rucio as the primary data management system serving the entire community.

A first key step is enabling user access through IAM. This requires a dedicated MoU and the integration of the KM3SAML Identity Provider.

The Rucio data management model is not yet well known within the scientific community. Dedicated training initiatives are therefore needed to help users properly organise Scopes, Containers, Datasets, and Replication Rules.

One of the initial technical steps is the migration of the IN2P3 'semi-permanent storage' directory named **`sps`**`/...` to a Rucio RSE. This directory currently contains a huge amount of data, ranging from personal analysis output files to processing and mass production intermediate and output files. The directory requires frequent manual intervention and its capacity is constantly under pressure. Thus, migrating to Rucio provides a more scalable and automated data management solution for the long term.

Once Rucio is fully configured, the Snakemake-based data processing workflows will be migrated to the Grid. In this context, a development effort has started under the name "KM3NeT File System". At the same time, the available RSEs must be properly configured to improve data organisation, distribution, and replication efficiency.

A new dedicated DIRAC instance hosted at IN2P3 will soon become available, replacing the current general-purpose EGI instance. The transition will not have impacts on the existing configurations. However, several issues remain unsolved, such as linking IAM membership with DIRAC, the fact that token support is not yet available, and the need to wait for a stable DIRAC

PoS(ISGC2026)018

X instance that supports tokens. Finally, there is not yet an official solution to integrate Rucio with DIRAC; nevertheless, the current Grid data processing uses alternative ways to access Rucio data.

Another important service is CVMFS for software distribution, currently provided by EGI at RAL (UK). The KM3NeT collaboration has encountered difficulties in integrating the GitLab-based CI/CD pipeline to automatically update CVMFS content, as programmatic access was not initially guaranteed. As a temporary workaround, GSI access has been reactivated for the collaboration. A more permanent solution, based on a dedicated publisher service, is currently being installed. Meanwhile, CNAF offers a CVMFS instance that already provides a publisher service. The collaboration is planning to test this instance as a potential long-term alternative.

5. Summary and outlook

The modular design of the KM3NeT detectors has enabled data taking since the early stages of deployment. These early operations have already produced major scientific results, including the observation of an ultra-high-energy cosmic neutrino [13]. The increasing number of operational detection units, however, produces a corresponding growth in data volume and computing requirements. The existing computing model, historically centred on a limited number of Computing Centres and workflows optimised for local execution, must therefore evolve towards a more distributed infrastructure capable of supporting data replication, repeated processing campaigns, Monte Carlo production, and physics analysis throughout the experiment's operational lifetime.

This contribution has described the KM3NeT tiered computing model, the data flow from the detector to the analysis level, and the current Snakemake-based processing workflow. Although this workflow provides reproducible and automated processing on local clusters, its deployment on heterogeneous Grid resources still presents limitations related to data distribution, software harmonisation, job parallelisation, and the availability of suitable worker nodes.

To address these challenges, KM3NeT has signed a Memorandum of Understanding with EGI.eu and is adopting Rucio as its central distributed data-management system. Work is also progressing on the deployment of a dedicated DIRAC instance, the integration of Rucio into the processing workflows, the configuration of additional storage elements, and the use of CVMFS for software distribution. Several technical and organisational issues remain open, including the transition from X.509 certificates to token-based authentication, the integration of IAM with Rucio and DIRAC, the automation of CVMFS publishing, and the development of sufficient Grid expertise within the collaboration.

The transition to distributed Grid computing is therefore both a technical requirement and an organisational process. Future work will focus on completing the integration of Rucio with the production workflows, stabilising the dedicated DIRAC and CVMFS services, improving data and job distribution, and expanding training and support activities for KM3NeT users. The experience gained may also be relevant to other medium-sized scientific collaborations approaching multi-petabyte data volumes.

References

[1] Adrián-Martínez S et al., "Letter of intent for KM3NeT 2.0", J. Phys. G: Nucl. Part. Phys. 43 (2016) 084001, DOI: 10.1088/0954-3899/43/8/084001

[2] Aiello S et al., "The KM3NeT multi-PMT optical module", JINST 17 (2022) P07038, DOI 10.1088/1748-0221/17/07/P07038

[3] de Jong M et al., "The online data filter for the KM3NeT neutrino telescopes", NIM A 1083 (2026) 171097, DOI: 10.1016/j.nima.2025.171097

PoS(ISGC2026)018

[4] Bonacorsi D and Ferrari T, "WLCG Service Challenges and Tiered architecture in the LHC era", in Montagna G, Nicrosini O and Vercesi V (eds), IFAE 2006 (Springer, Milano, 2007), DOI: 10.1007/978-88-470-0530-3_68

[5] Bouwhuis M et al., "KM3NeT Data Management Plan", Deliverable D4.1 of the KM3NeT-INFRADEV2 project (Grant Agreement No. 101079679), European Commission, 2023, KM3NeT-INFRADEV2-WP4-D4.1-v.approved

[6] KM3NeT-INFRADEV2 project, European Union Horizon Europe Research and Innovation Programme, Grant agreement No. 101079679, DOI: 10.3030/101079679

[7] Albert A and Bozza C, "The relational database system of KM3NeT", EPJ Web Conf. 116 (2016) 07004, DOI: 10.1051/epjconf/201611607004

[8] Sinopoulou A et al., "Management of the data processing & Run-by-Run simulations in KM3NeT", EPJ Web Conf. 337 (2025) 01046, DOI: 10.1051/epjconf/202533701046

[9] Vázquez de Sola Fernández F, "Rucio for KM3NeT", presented at the Eighth Rucio Community Workshop, CERN, 2025

[10] WLCG AuthZ WG, "The WLCG Virtual Organization VOMS configuration", WLCG AuthZ Documentation, wlcg-authz-docs

[11] Schnabel J, "WP4: Data Management and Open Science" [Oral presentation], KM3NeT-INFRADEV2 Final Review Meeting, CERN, Geneva, Switzerland, 2026, KM3NeT-INFRADEV2 Final Review

[12] Sinopoulou A (for the KM3NeT Collaboration), "Advancing KM3NeT Data Management: Harnessing Snakemake and Grid Computing", PoS(EPS-HEP2025)571 (2026)

[13] The KM3NeT Collaboration, "Observation of an ultra-high-energy cosmic neutrino with KM3NeT", Nature 638 (2025) 376-382, DOI: 10.1038/s41586-024-08543-1

PoS(ISGC2026)018