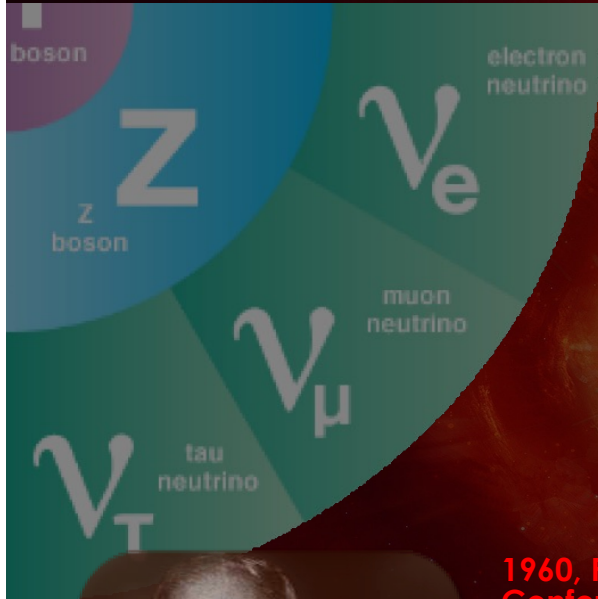


# Data and Computing Infrastructure for KM3NeT: Towards Full Deployment

**R. Bruno**<sup>\*</sup>, F. Vasquez<sup>+</sup> de Sola  
on behalf of the KM3NeT Collaboration

18/03/2026 ISGC 2026 - <sup>\*</sup>INFN - KM3NeT Computing&Software co-coordinator  
<sup>+</sup>Nikhef

- Neutrino detectors, the KM3NeT collaboration
- ARCA and ORCA detectors
- All data to shore and the tiered computing model
- Computing centres and their services
- The data flow and processing
- Transition to the Grid
- EGI, Rucio, Dirac and other Grid services
- Conclusions



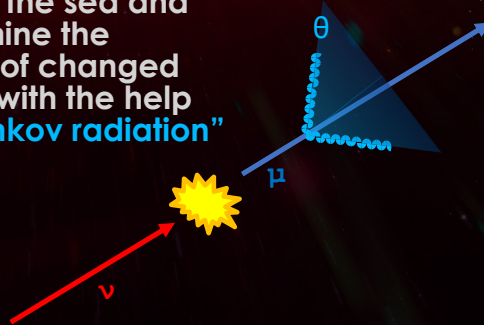
## Neutrino particles

- Almost null mass and no electric charge
- Difficult to detect
- Travel straight along intergalactic distances
- Astrophysical messengers
- Flavour change (oscillation)
- Rarely interaction with matter occurs



1960, Rochester Conference  
M. Markov

“We propose to install detectors deep in a lake or in the sea and to determine the direction of changed particles with the help of Cherenkov radiation”



## Detector needs

- Capture light → **Optical Modules**
- Small quantity of photons → **Dark**
- Transparent material → **Water/ICE**
- Improve interaction probability →

**Large Spaces**

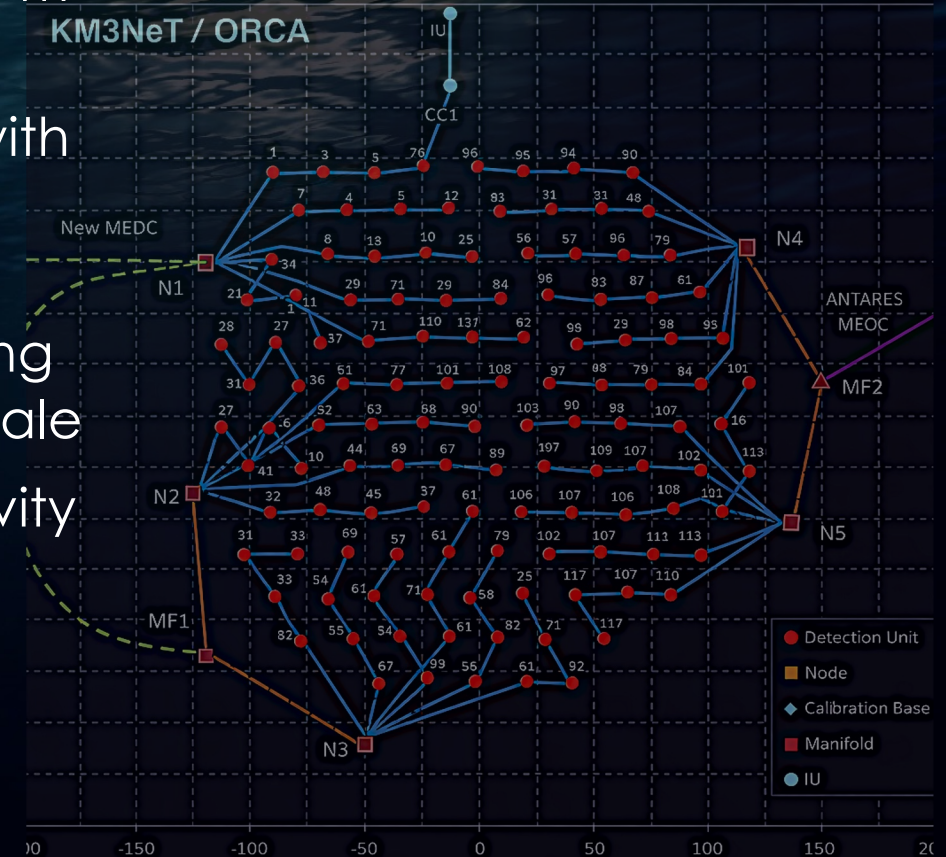
**Deeps & Earth's Mass**



# Oscillation Research with Cosmics in the Abyss (ORCA)



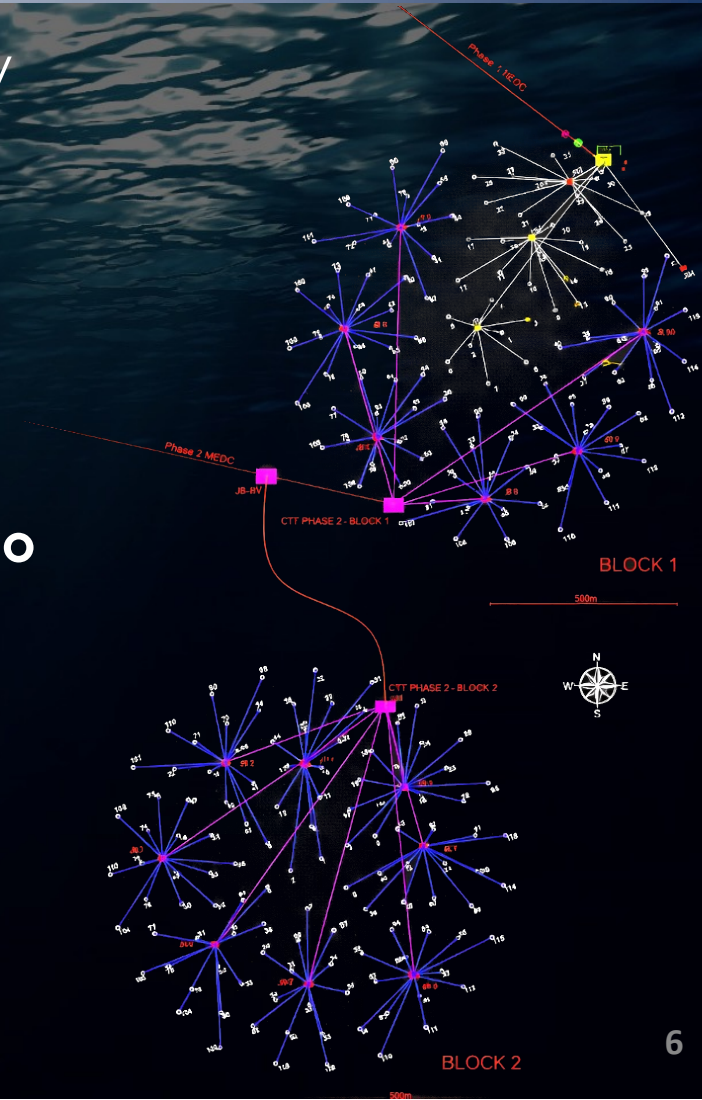
- Located 40Km off Toulon, France at ~2500 m depth
- Volume is about 7 Mtonnes ( $\sim 0.007 \text{ km}^3$ ) with **one building block** of 115 DUs
- **Neutrino oscillation physics**, including determination of the neutrino mass ordering using atmospheric neutrinos in the GeV scale
- Denser photosensor array improves sensitivity to neutrino oscillation effects
- Deployed 33 DUs so far



# Astroparticle Research with Cosmics in the Abyss (ARCA)



- Located 100Km off Capo Passero (Sicily), Italy at ~3500 m depth,
- Focus: **Astro particle physics and neutrino astronomy** - finding astrophysical sources of neutrinos
- Designed for **high-energy** cosmic neutrinos (TeV–PeV)
- Volume is about 2\*500 Mtonnes (~1 km<sup>3</sup>) ~ **two building blocks** of 115 DUs each
- Deployed 51 DUs so far



- Noise sources
  - Bioluminescence
  - K40 radiation
- Calibration
  - Position and orientation
    - Acoustic & Compass position system
  - Time synchronization
    - Light emitters
- Data
  - Long-distance data and power transmission
  - Maximum bandwidth exploitation
  - Redundancy
- Building
  - Dimensions
  - High pressure
  - Sea water chemical aggression
  - No maintenance (10-20 years)
  - High-quality components
    - Same as aerospace
  - Costly marine campaigns
    - Specialized Equipment/Crews
    - Weather dependency
    - 1,2 Marine Campaigns a year at most
- Environment
  - No impacts on the biodiversity
  - Decommissioning plan

# KM3NeT – Main detector components




White Rabbit Switches  
ns – 1GBs Optical Network



## 3" Photomultiplier (PMT)

High sensible photo-sensor, detects even a single photon  
A PCB controls voltage and converts the analog input to a digital signal



## Digital Optical Module (DOM)

A glass sphere containing **31 photomultipliers**  
A PCB with FPGA, sensors, and optical SFP/Glen-air transceiver  
A hydrophone and a light emitter for positioning and time synchronisation



## Detection Unit (DU)

**18 DOMs** linked together by two ropes, and an electro-optical cable

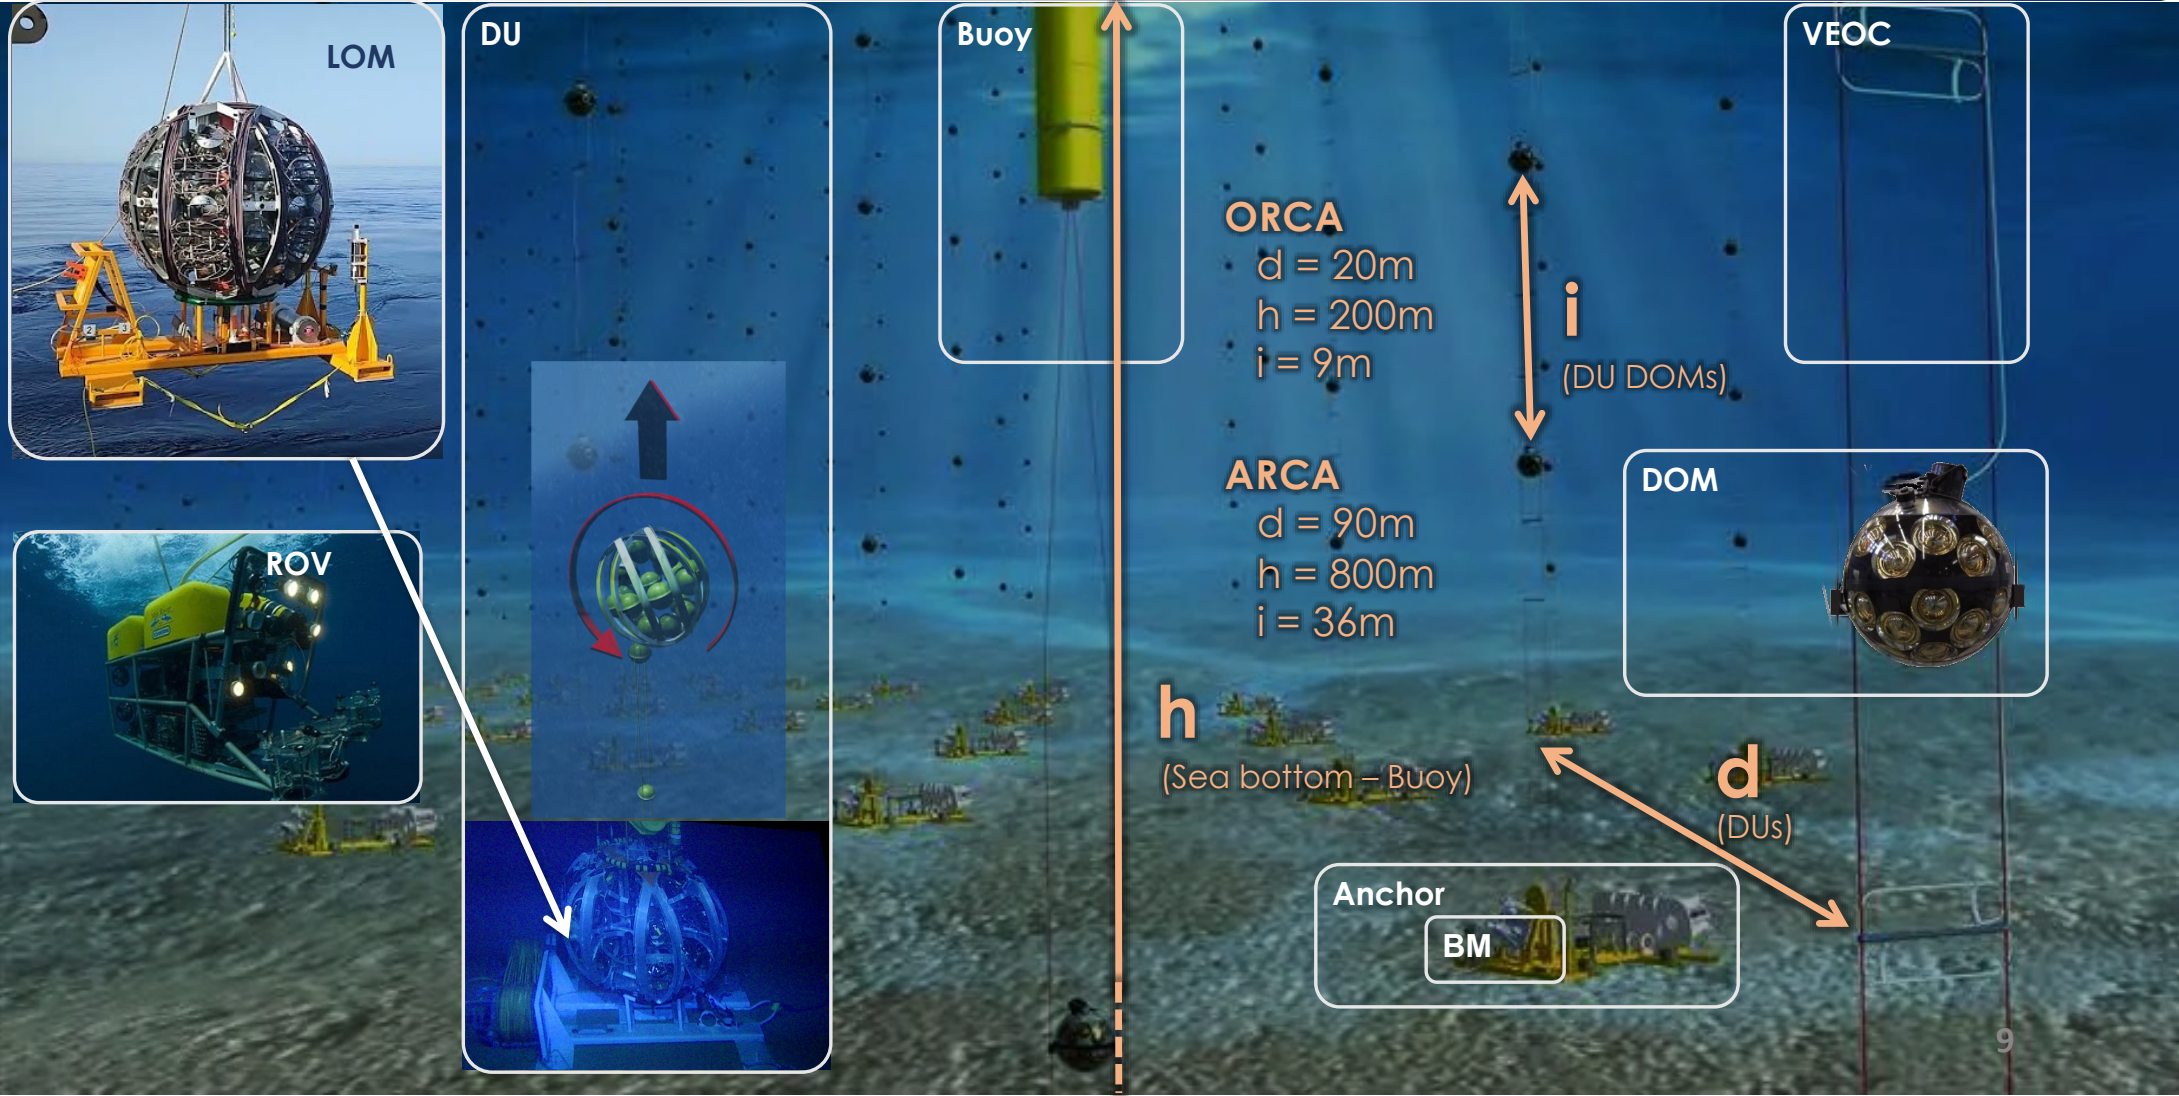
## Base Module (BM)

Link DOMs to the shore

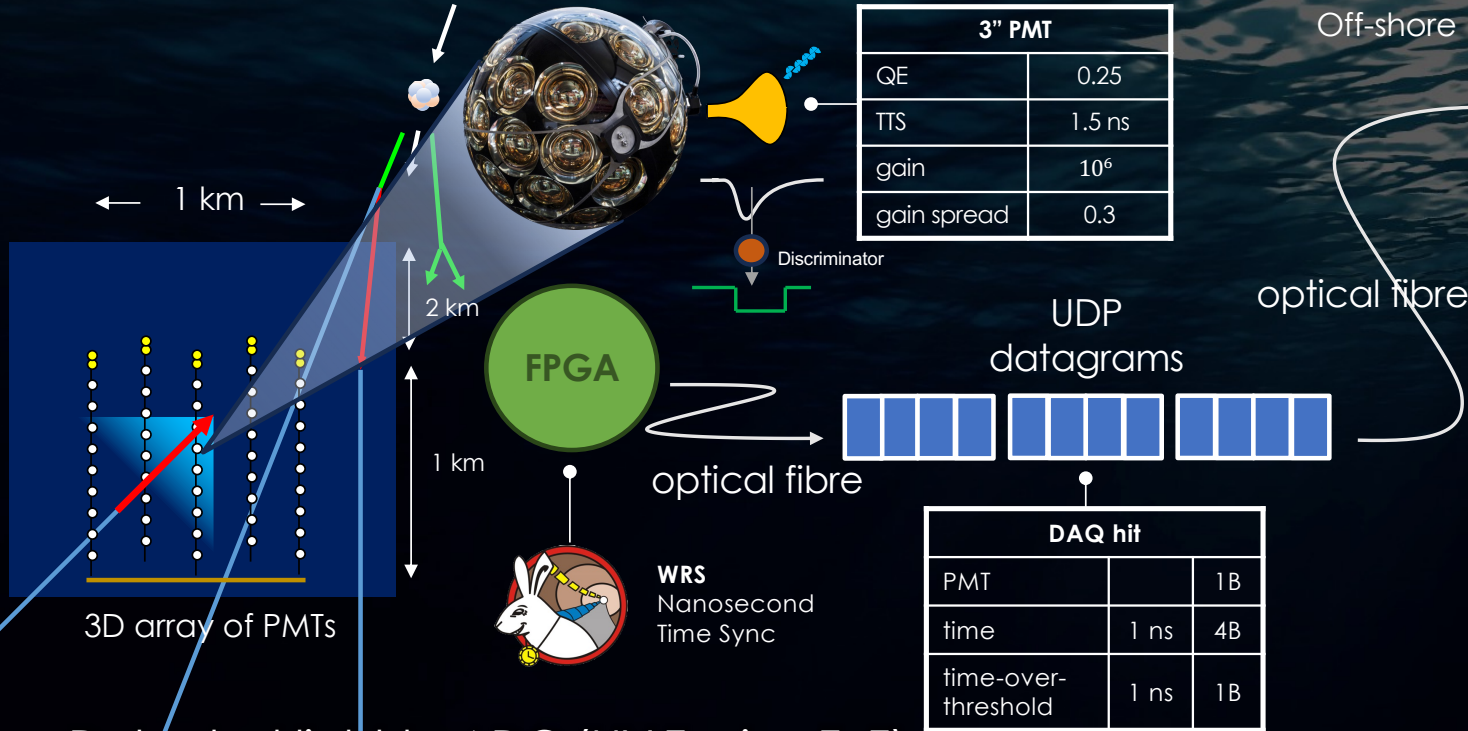


Dry (ARCA/ORCA)  
Wet (ARCA)

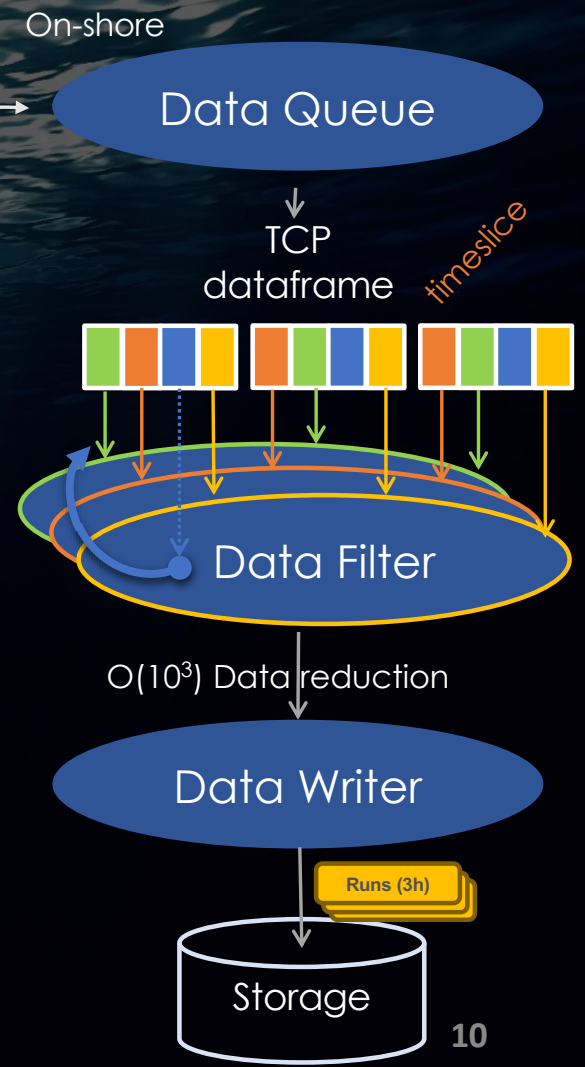
# Detector deployment and shape



# All data to shore



- Detected light to ADC (HV Tuning ToT)
- The DOM's FPGA sends information via UDP
- Time unsorted datagrams ordered by DQ
- DQ sorts and dispatches to DFs (TimeSlices 100ms)
- DFs Process a timeslice and applying triggering algorithms
- DW stores the the raw data to the Disk



# Computing Model – Data management Plan



## Tier 2 – Higher-level analyses

Local clusters, Personal computers, Clouds, ...

## Tier 1 – Processing

Calibration sets (Git) ← KM3Net DB, raw data

Reconstruction ← raw data + Calibration

Modeling fit ← find energy and direction of neutrino candidates

Simulated Data ← Monte Carlo

## Tier 0 – Detector

DAQ filtered 'raw data' (onshore)

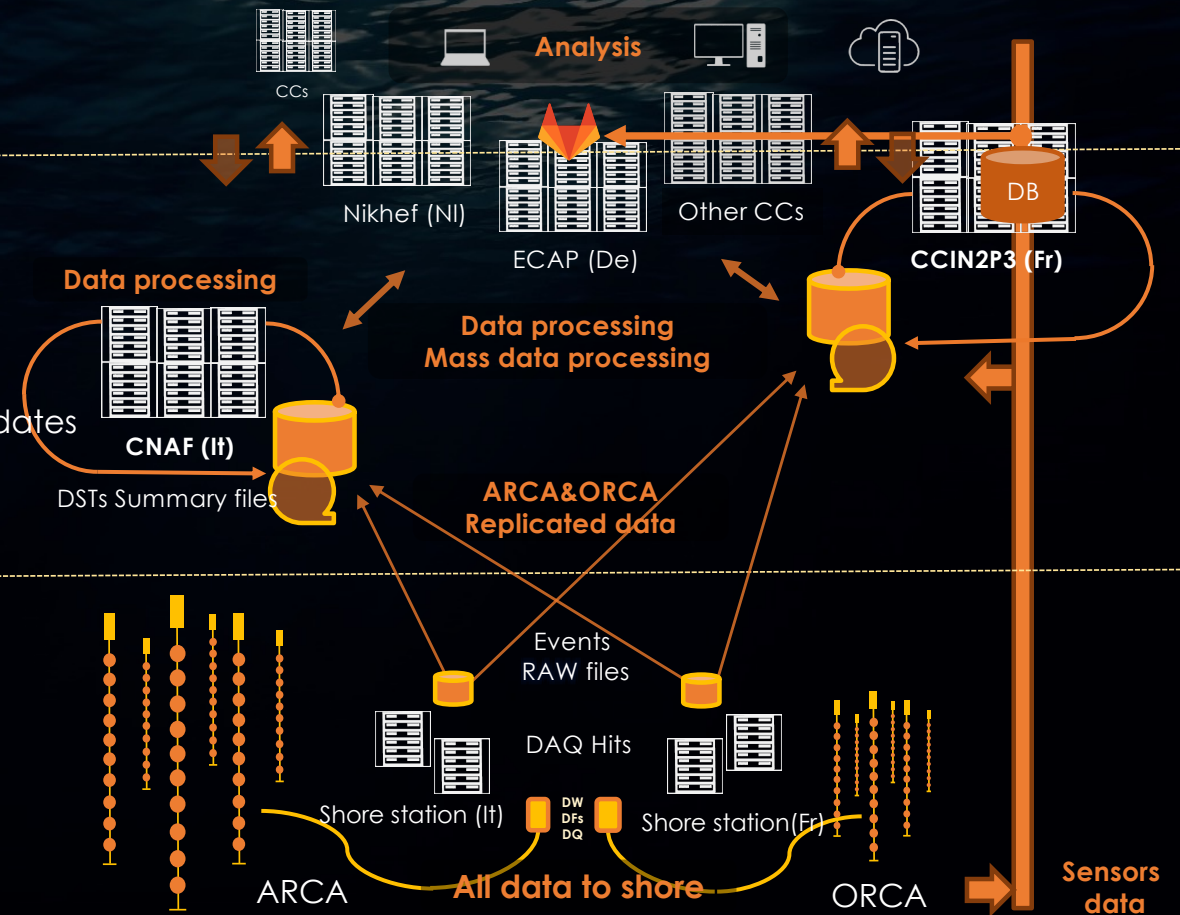
~300 Mb/s x DU (~30% of the final detectors)

Data daily sent to Tier 1

10Gb/s - 520 TB/year

Files made of three hours each (**run**)

Tiered hierarchical structure (LCG computing model)





# Other CCs and ICT Services



## Services

- T2 Nikhef (NI)
- T1 Stoomboot (Local cluster) **MP**
- ELPROD (Grid)**
- T1 SURF (EGI) (NI)
- SARA, SARA-MATRIX (Grid)**
- T2 ECAP (De) **MP**
- Woody
- T2 AGH (PI)
- T1 Storage: O(PB)
- CPU: O(100MHS06h)
- Cyfronet
- T2 IFIC (Es)
- T2 Storage: O(100TB)
- CPU: O(10MHS06h)
- Gluon Artemis
- T2 HULL (UK)
- Viper
- T2 Demokritos (Gr)
- T2 UM6P (Ma)

## IN2P3

- ORACLE DB + KM3SAML
- Production services
- CORAL/CORMORANT
- Sardine (Shift)
- DIRAC (EGI → KM3Net)

## CNAF

- IAM
- CVMFS

## NIKHEF

- RUCIO

## ECAP

- RocketChat
- Wiki
- GitLab
- eLog
- sFTP

## CERN

- Indico
- FTS

## RAL (EGI)

- CVMFS

## CCPM

- ORCA Monitoring

## INFN

- ARCA Monitoring
- DIA (Dom Integration)
- VPN

## Third parties

- GOOGLE (km3net.de)
- Drive
- Doc
- Sheet
- Calendar
- Groups
- ZOOM
- MS Teams (Marine Campaigns)



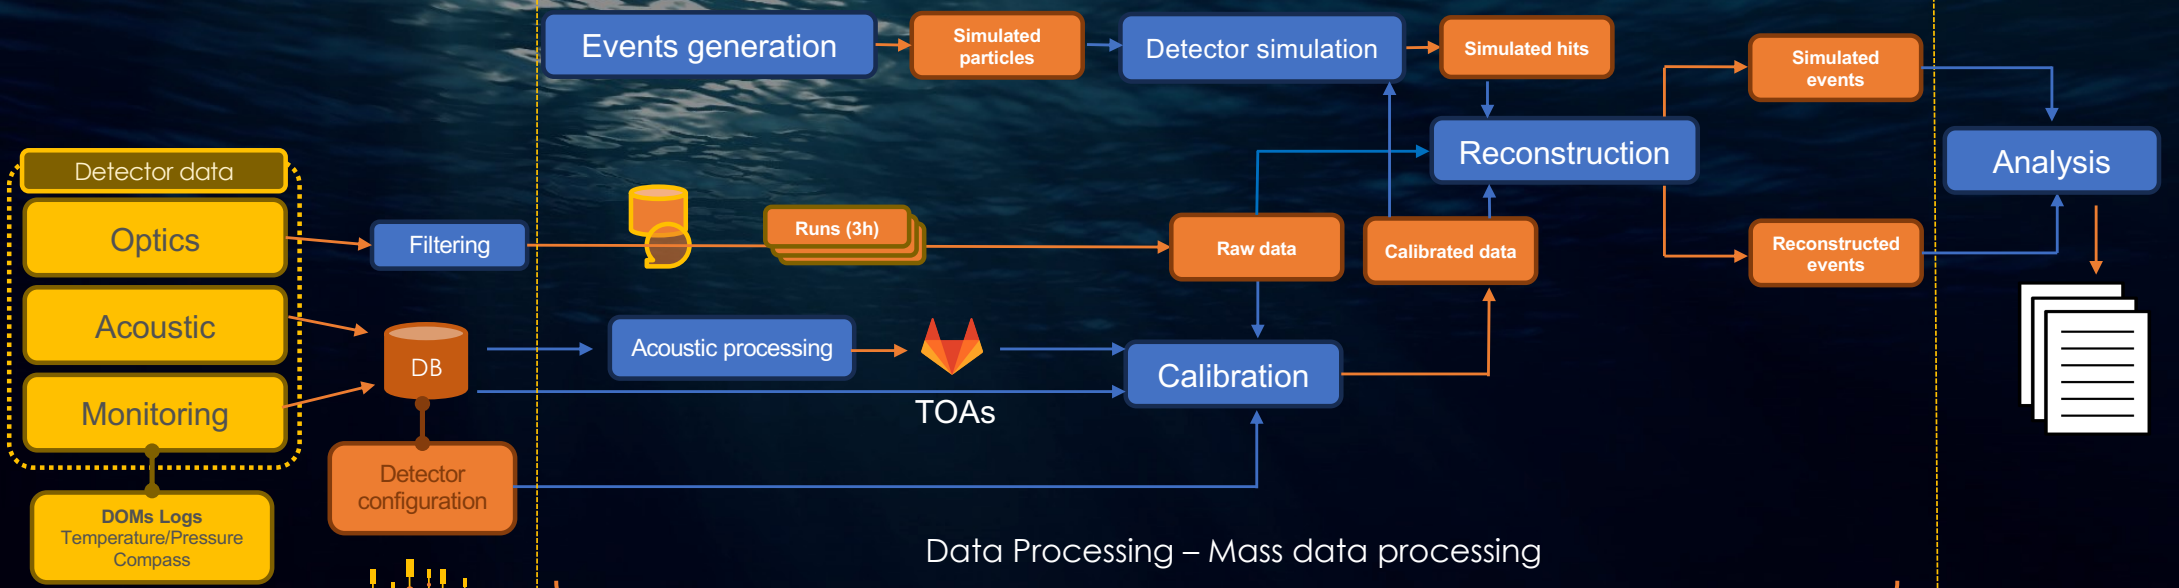
# Run based data processing



T0

T1

T2



Data Processing – Mass data processing

Run Based Data Processing
 ➔
 Run Based Grid Data Processing

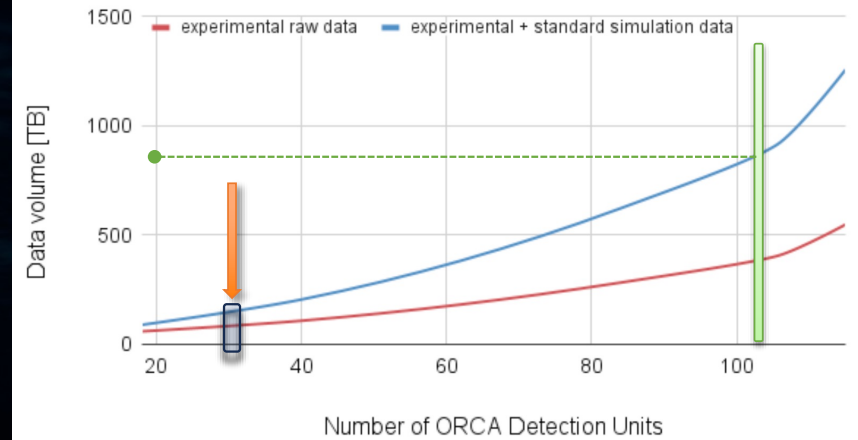
- Developed for Clusters (IN2P3 - Slurm)
- Full working prototype
- Maintainable
- Need multicore WNs
- Configurable
- Direct storage access (xrood)
- Adapted for CNAF (HTCondor)
- Not yet in production
- 450 HS06h per run

# Grid transition

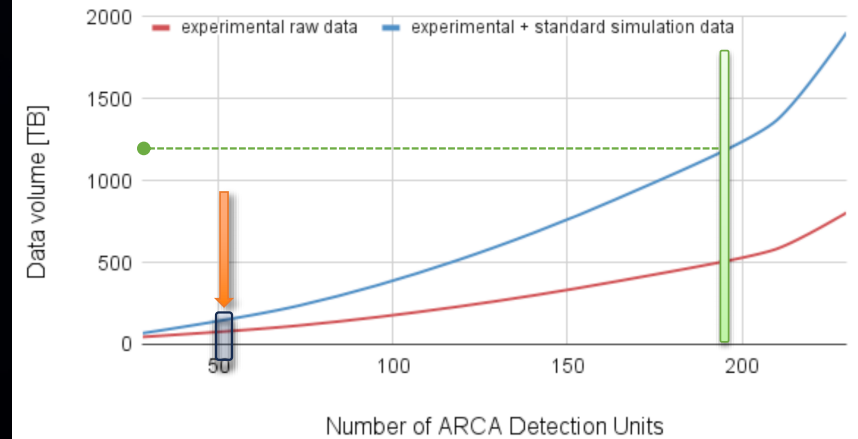


- Scaling KM3Net Computing for 10 years after deployment
- Both detectors will generate  $O(\text{PB})$  year of data
- Current Storage and computing resources will be not enough
- A distributed solution is required
  - INFRADEV2 outcomes
  - Confirmed by STAC (Technical evaluation)
- Two phases
  1. Data management transition
    1. Adoption of RUCIO as central DMS
    2. Integration of RUCIO in data processing
    3. Identification of Tier1 storage elements (RSEs)
  2. Toward a Grid-Native Production model
    1. Merge the two Snakemake projects
    2. Integrate RUCIO and Snakemake
    3. Switch to the Grid data processing as the main production system

ORCA Data volume per year



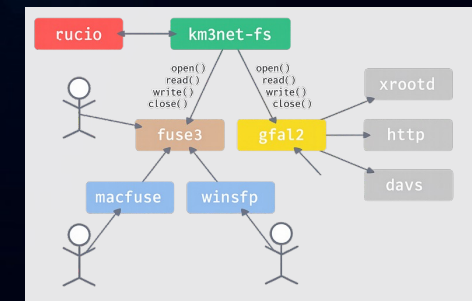
ARCA Data volume per year



- EGI MoU – November 2025 (Initial relationships since 2013)
- EGI Operations services (GOCDDB, GGUS, O.Portal, ...)
- Sites
  - SURF-SARA, SARA-MATRIX, ELPROD, CNAF T1, IN2P3, ... (VO: [km3net.org](http://km3net.org))
- Services
  - CVMFS (RAL), **RUCIO** (Nikhef), **DIRAC** (IN2P3), FTS (CERN), **IAM** (CNAF)
- Grid Issues
  - Consolidated procedures and low know how
  - Legacy certificate transition (X509 → Tokens)
  - Sectigo decommissioning → HARICA/CERN/Other IGTF CAs
  - Rucio RSEs support for Tokens
  - Link of IAM users to KM3Net users (KM3SAML)



- Hosted at NIKHEF (Kubernetes) – FTS (CERN)
- Transiting from a tester prototype to a production service
- IAM users integration (MoU) + KM3SAML IdP
- Spread the usage know-how
- Data Efficiency: Scopes/Containers/Datasets, R.Rules, M.data
- User storage procedures (IN2P3' /sps → Rucio)
- Snakemake integration (KM3NeT filesystem)
- RSEs - identification, organization, data-distribution, replication





- DIRAC EGI (dirac.egi.eu) instance hosted at IN2P3
- New dedicated instance (dirac.in2p3.fr)
  - Seamless migration: Users, Resources, ...
- Membership not yet integrated with IAM
- Tokens support not yet ready (DIRAC X)
- RUCIO integration (rucio-plugin)
  
- EGI CVMFS @RAL
- Issues integrating our software CI/CD (GSI)
- CNAF with publisher, ready for CI/CD (Evaluating)
  
- Integrating KM3SAML in EduGAIN (IAM)

5 Sites

GRID.CNAF.it  
GRID.CPPM.fr  
GRID.IN2P3-CC.fr  
GRID.NIKHEF.nl  
GRID.SARA.nl

18 CEs

57 Queues

# Conclusions



Modular shape allowed to use the detector since the beginning



Currently at ~30% of the full detector size, completion ~2031



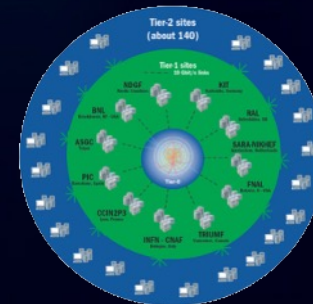
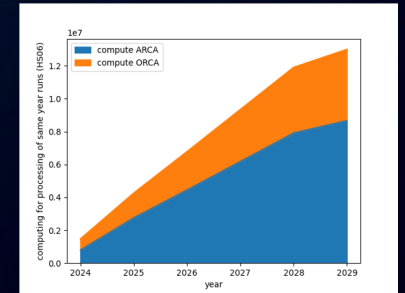
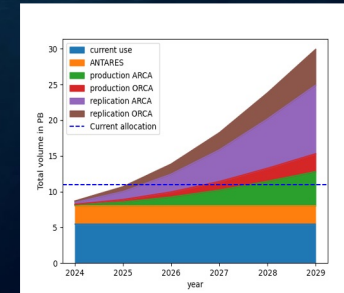
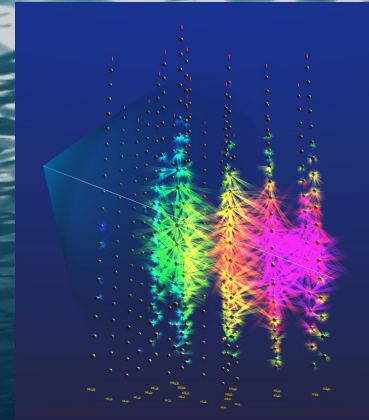
10 Years of of data taking after deployment



Huge amount of data to process and to analyse



Requiring to join dedicated infrastructures (WLCG, LHCONE, ...)





Thank you