JUNO Distributed Computing Infrastructure: Commissioning and Early Data Taking

Giuseppe Andronico

INFN Sez. CT

On behalf of JUNO DCI

Taipei March 19, 2026 - ISGC 2026

Jiangmen Underground Neutrino Observatory (JUNO)

- **Proposed as a reactor neutrino experiment for mass ordering in 2008** (PRD78:111103,2008; PRD79:073007,2009)
- **Project approved 2013, Construction started 2015, data taking started late August 2025**
- **Located Kaiping, Guangzhou, south of China (west of HongKong)**
- **JUNO baseline distance optimized to determine neutrino mass ordering**

JUNO physics program

- JUNO is a multipurpose Neutrino Observatory with a rich program in neutrino physics and astrophysics studying neutrinos over a large energy range
- Include solar, supernova, atmospheric, geo-neutrinos, proton decay, exotic searches

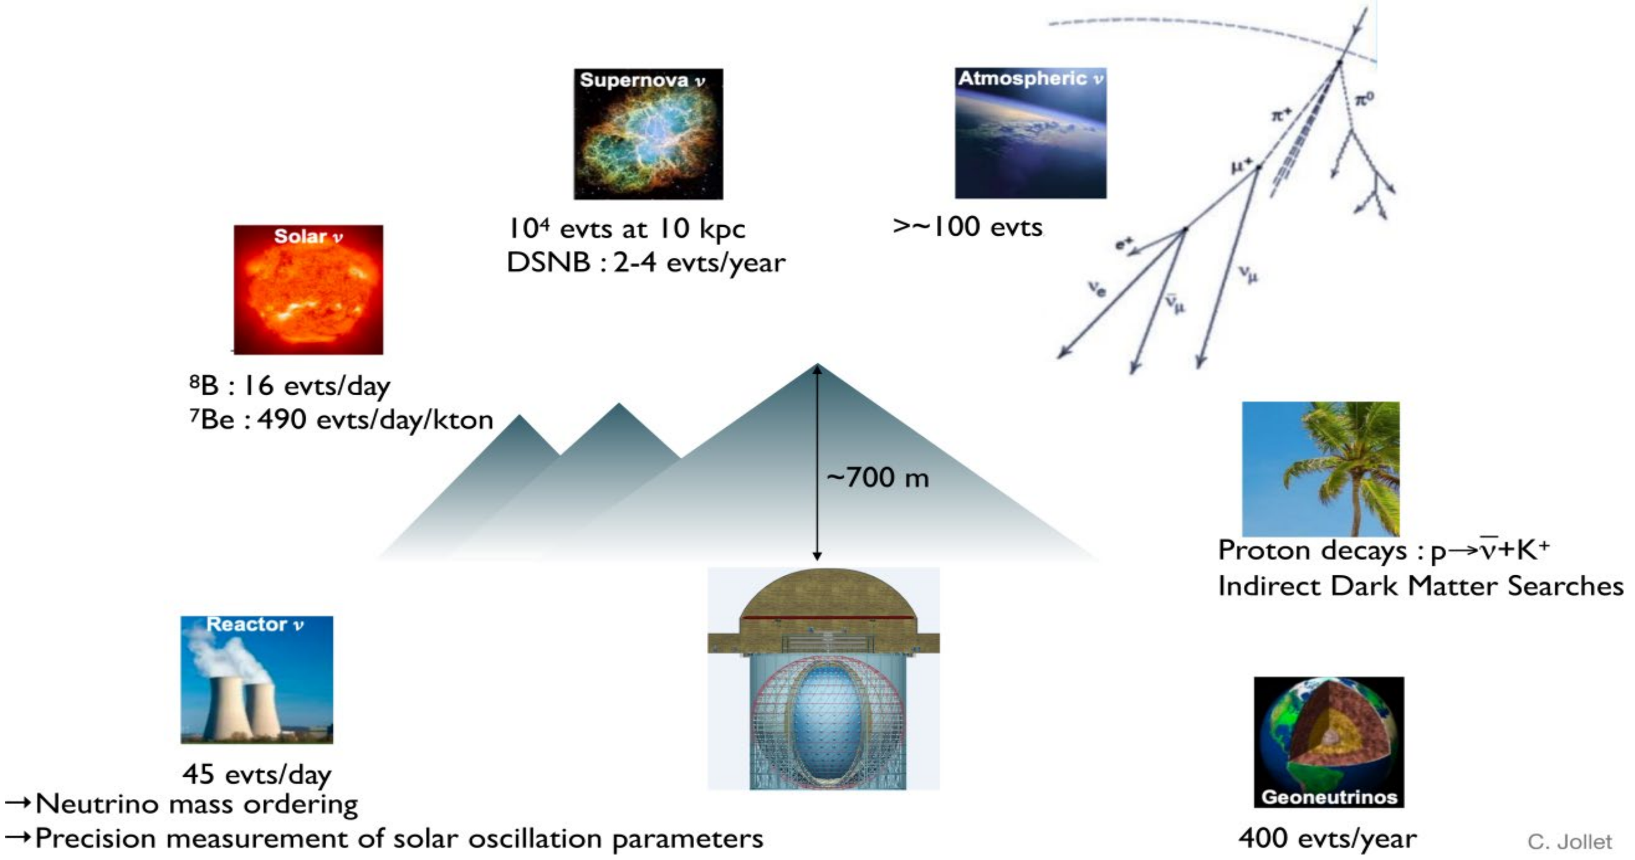

JUNO Detector

20 kton liquid scintillator, >30K PMTs, 700m underground
- **Designed to achieve 3% energy resolution at 1 MeV**

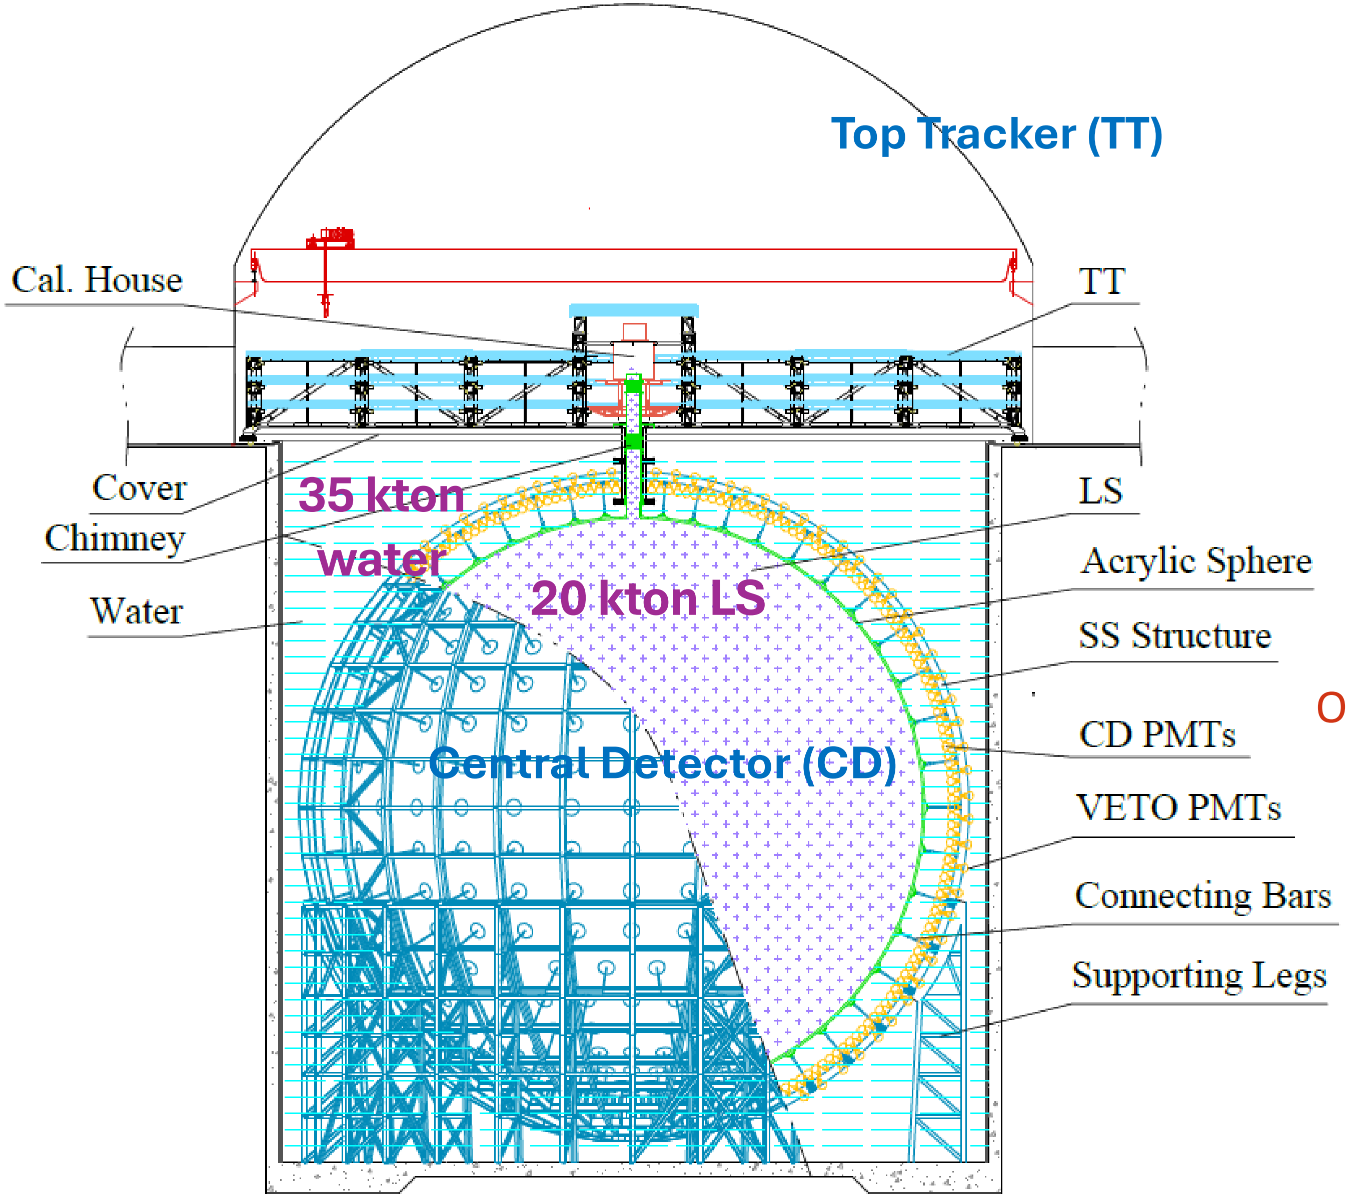

Acrylic Sphere:

Inner Diameter (ID): 35.4 m
Thickness:12 cm

Stainless Steel (SS) Structure:

ID: 40.1 m, Outer Diameter (OD): 41.1 m
17612 20-inch PMTs, **25600** 3-inch PMTs

Water pool:

ID: 43.5 m, Height: 44 m, Depth: 43.5 m
2400 20-inch PMTs

OSIRIS

Online Scintillator Internal Radioactivity Investigation System

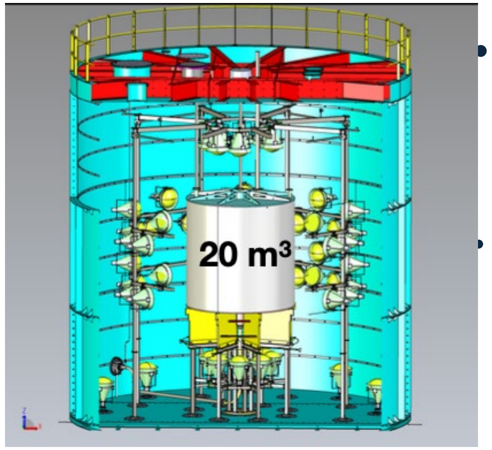

- **A dedicated pre-detector to measure radioactivity levels of the LS**
- 20 tons of LS in 3m-by-3m acrylic vessel, 76 MCP-PMTs, 3m of water shielding

Computing Model

- Raw data volume: ~2–3 PB/year.
- Computing power needed: ~300 kHS25
- Hierarchical Model: Tier 0 (IHEP) for primary storage/reconstruction; Tier 1s (CNAF, IN2P3, JINR) for replicas, re-reconstruction, and simulation.

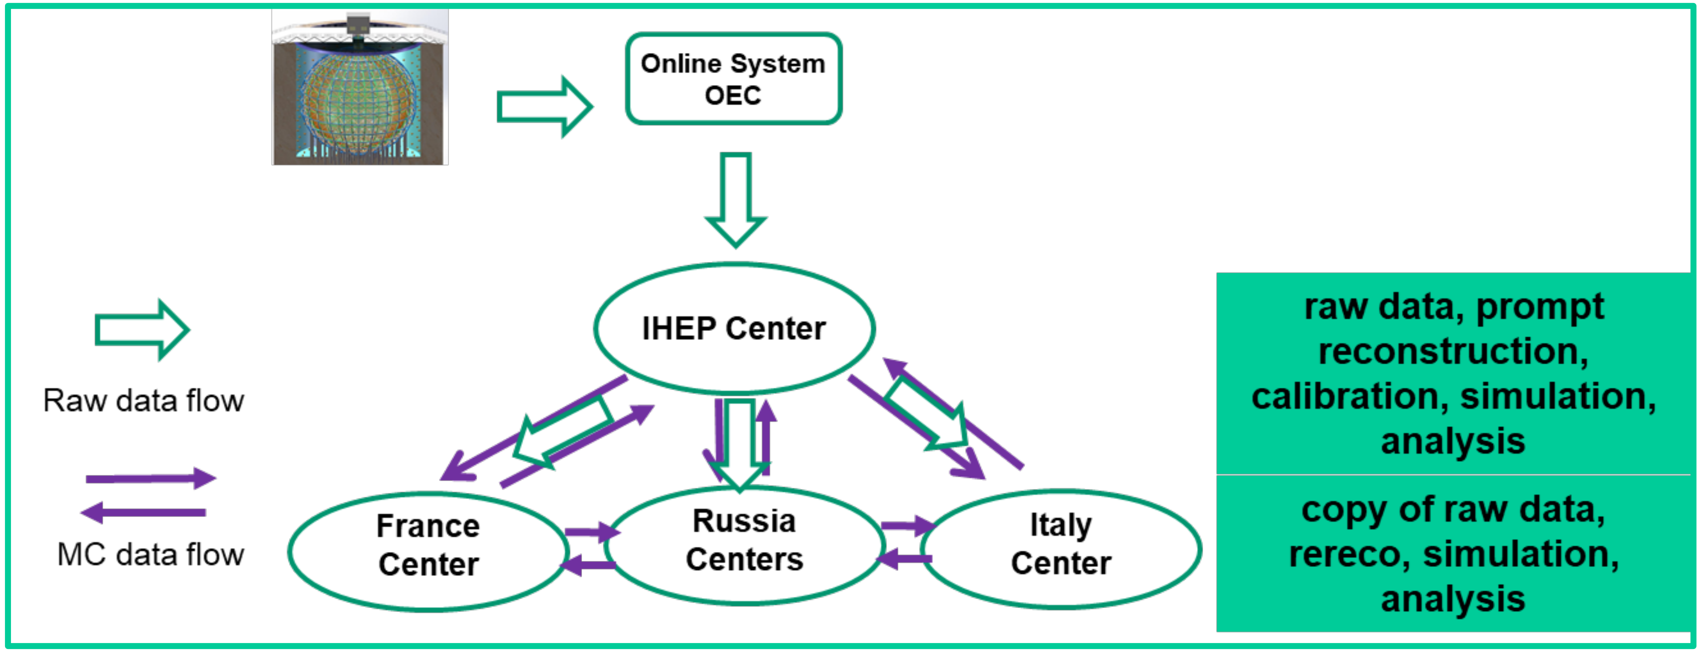

- **Raw data** flows from Online to IHEP which then distributes to other centers
- **KUP (Keep Up Production, 1st reconstruction)* and Calibration** will run in IHEP
- **MC Simulation, PP(Physics Production, 2nd reconstruction) and Analysis** runs in data centers

Distributed Computing Infrastructure (DCI)

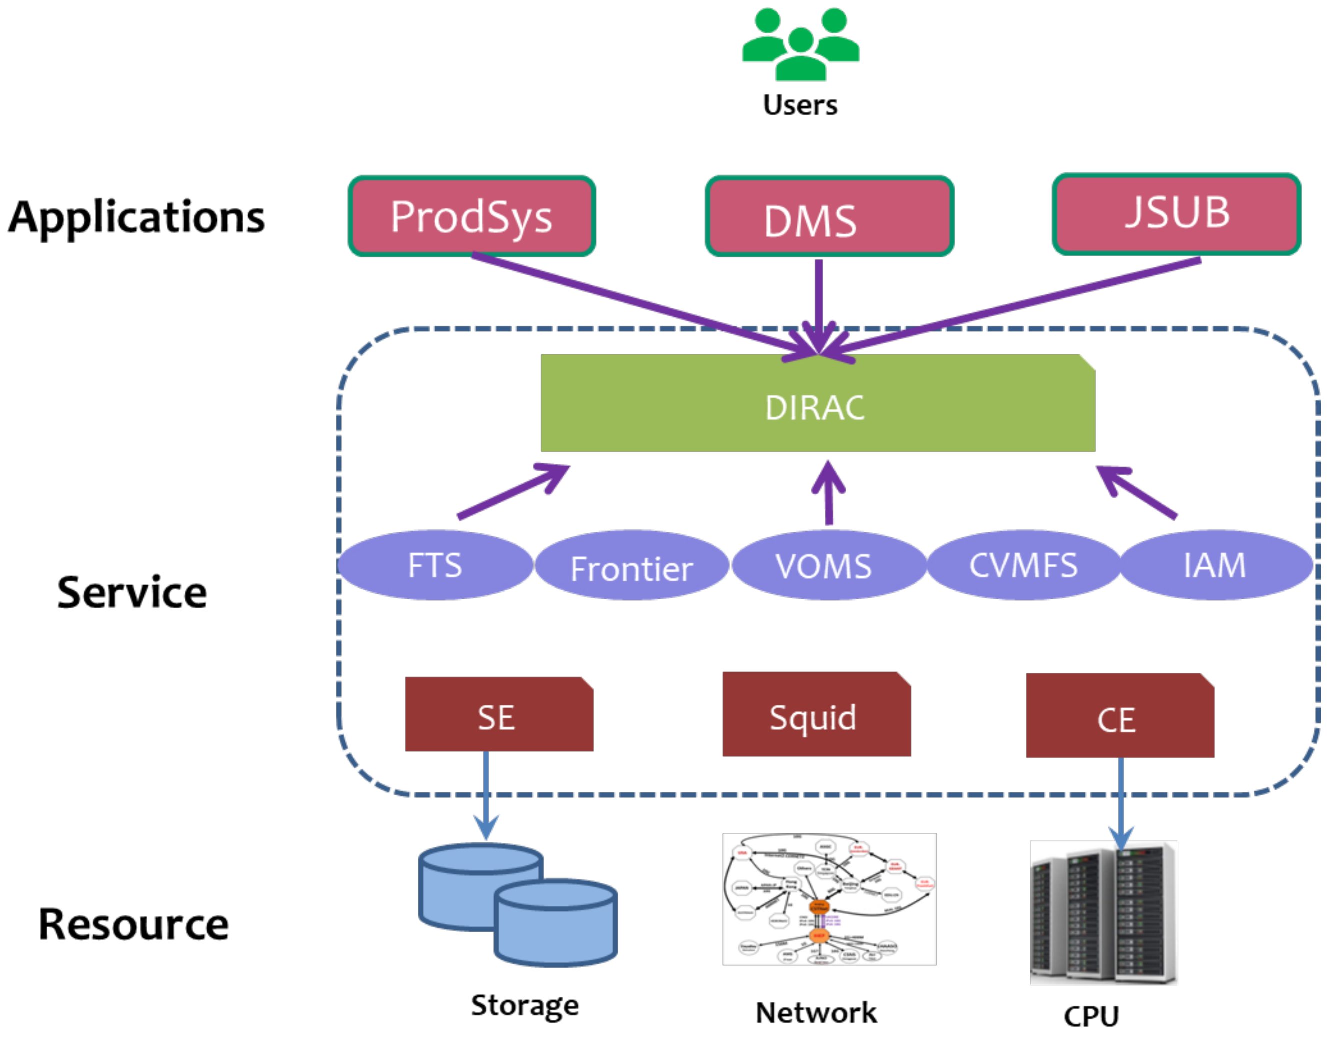

- **JUNO DCI was built based on DIRAC, using both WMS and DMS**
- **Use WLCG middleware as much as possible**
- **Software Distribution: CVMFS**
- **Condition Data: FroNTier/Squid; reduces central DB load by 99%.**
- **Transitioning from VOMS to token-based IAM.**
- **Develop JUNO-oriented application systems**
  - Data production system (ProdSys)
  - Data analysis tool (JSUB)
  - Raw data transfer system
  - Data operation tools

WMS: Workload management system, DMS: Data management system

Data rate and offline processing flow

- **KUP (Keep Up Production): Automated pipeline using Kafka for converting RAW to RTRAW (ROOT format) and producing ESD.**
- **SPADE: Automated transfer from the detector site to IHEP.**
- **Replication: Automated registration in DFC and parallel replication to Tier 1 sites.**

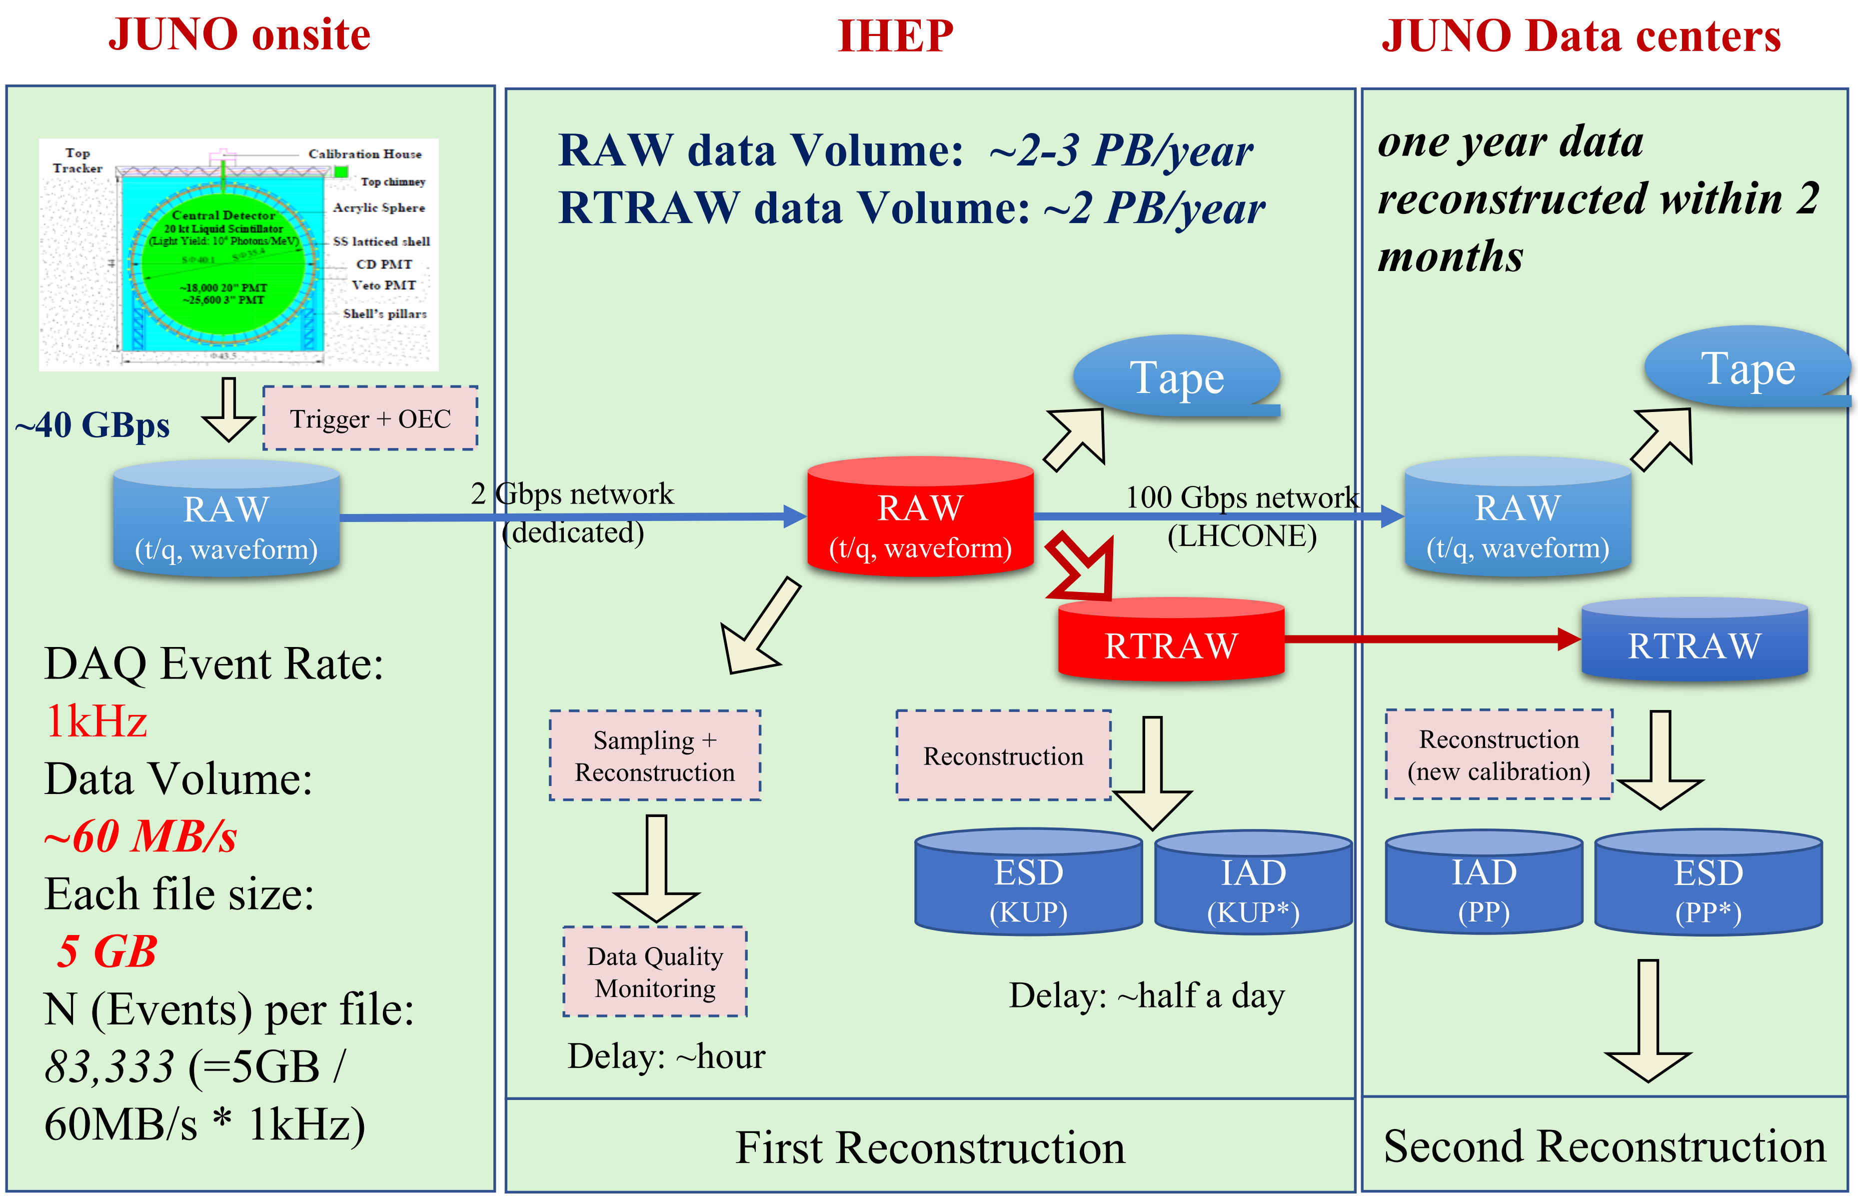

* KUP: Keep up production PP: Physics Production

International network

JUNO joined LHCONE in April, 2021

- **IHEP ⇔ Europe upgraded to 100Gbps in 2023**
  - ⇨ LHCONE 80Gbps (JUNO up to 35Gbps), LHCOPN @20Gbps
- **IHEP ⇔ JINR is 10Gbps, CNAF ⇔ JINR is 100Gbps**

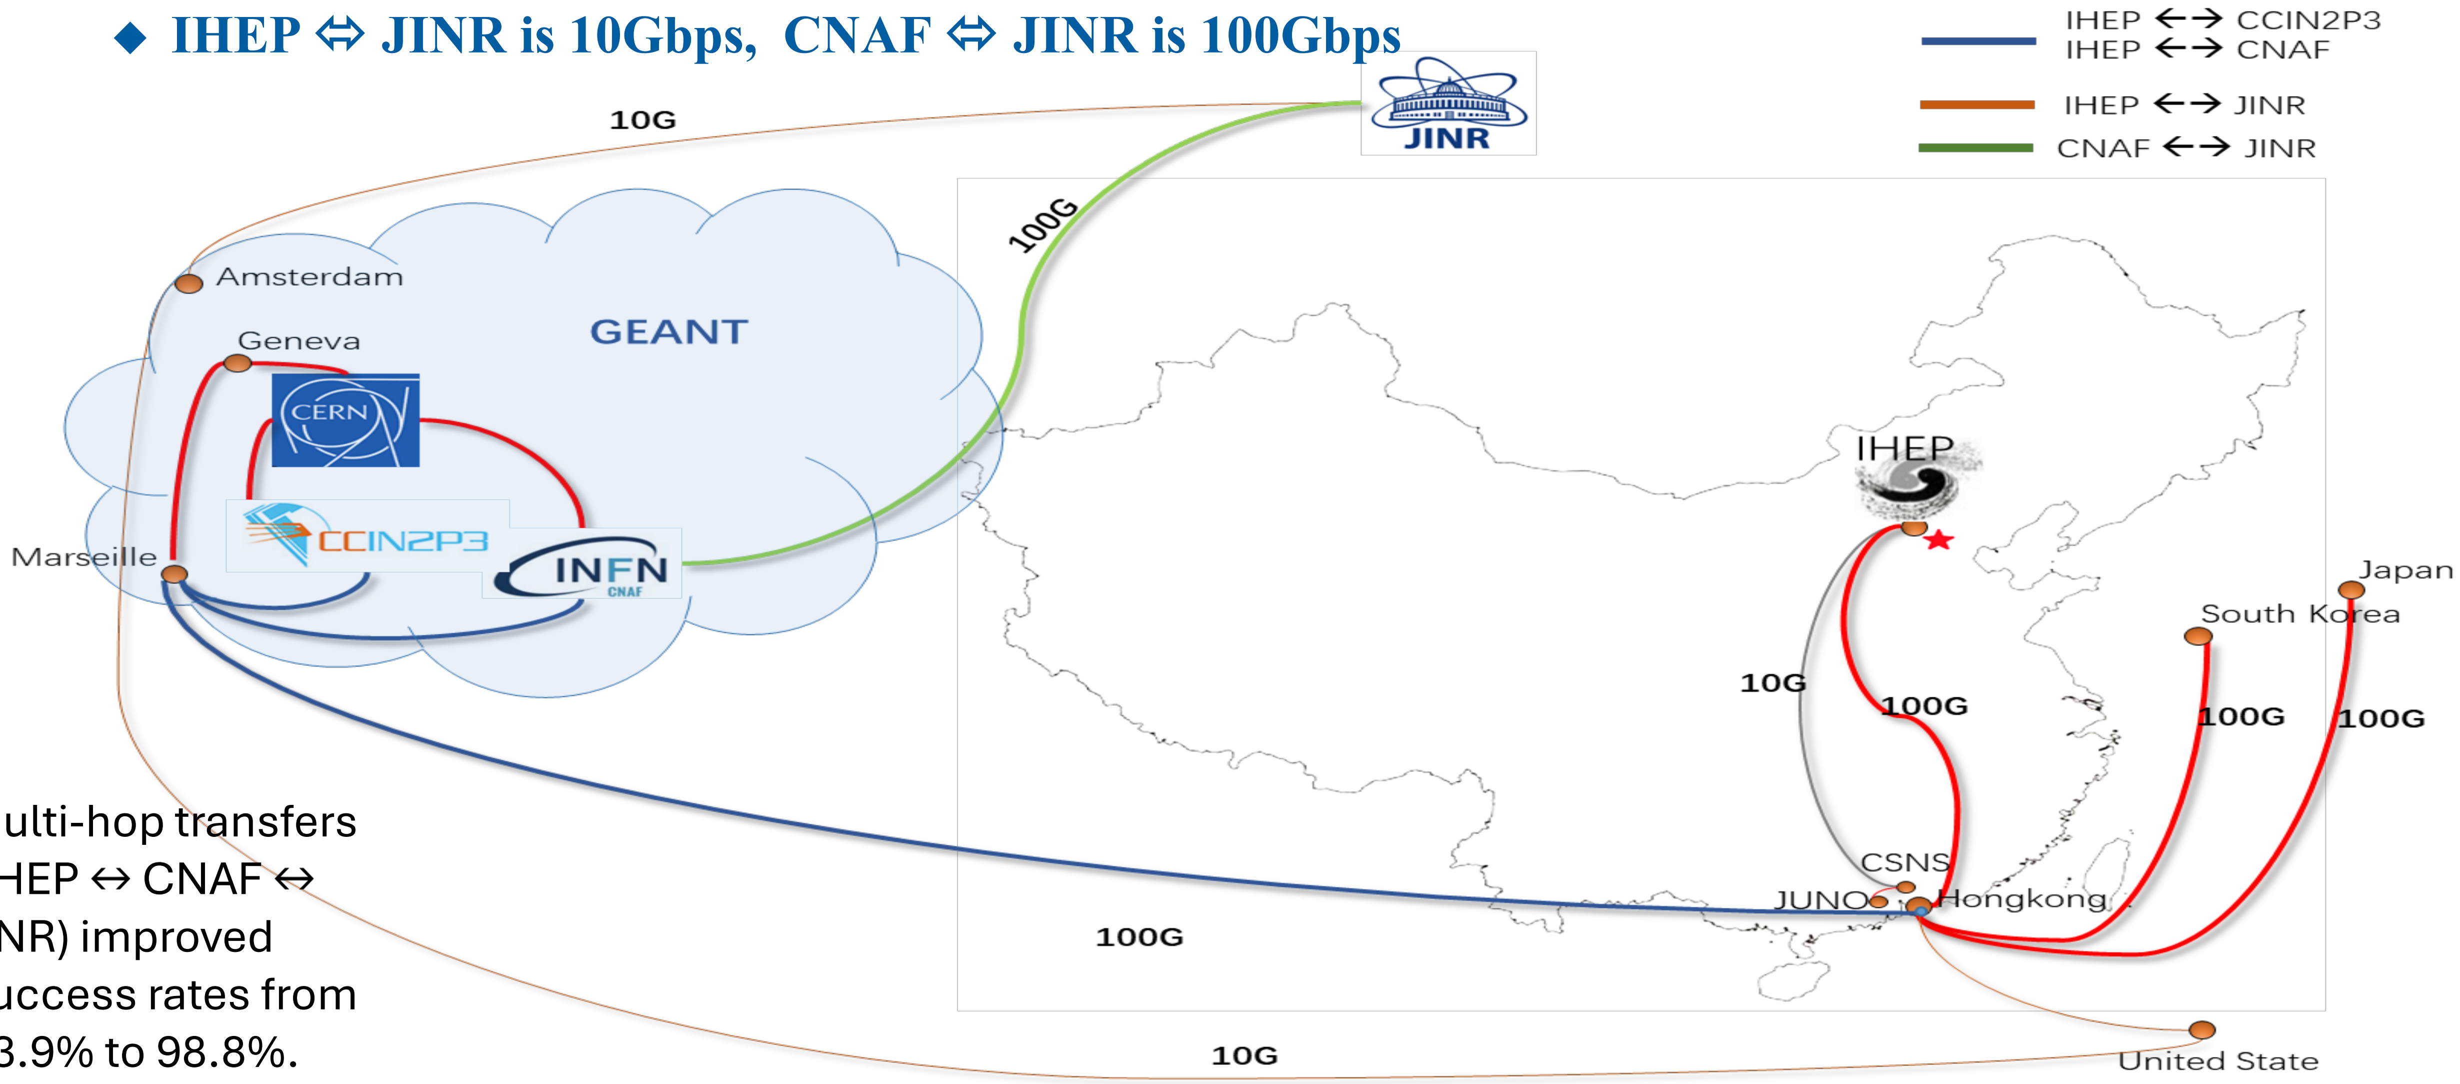

Multi-hop transfers (IHEP ↔ CNAF ↔ JINR) improved success rates from 73.9% to 98.8%.

Re-Production (ReProd 25) Campaigns Results

- Analysis of campaigns ReProd 25A, 25B, 25C, 25D
- Success rates: 95.3% to 98%.
- Scale: ~1 Million files, 0.6 PB replicas, and peak of >20k concurrent jobs.

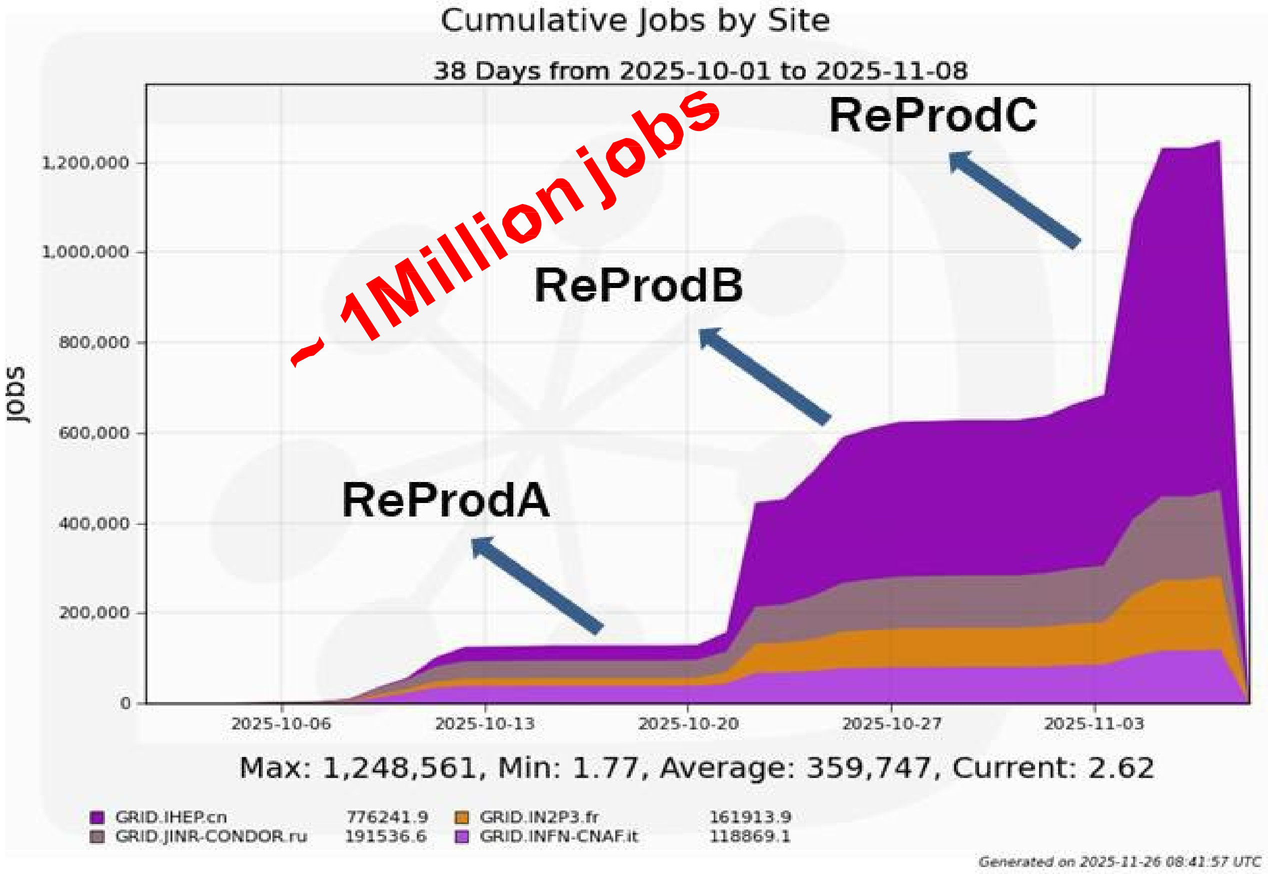

Resource Consumption & Site Contributions

- Total CPU used: ~8k HS06-year for the major campaigns.
- Normalized contributions: IHEP:JINR:CNAF:IN2P3 ≈ 13:7:5:4.5.
- Support for multi-core jobs to optimize HPC usage.

ReProdB

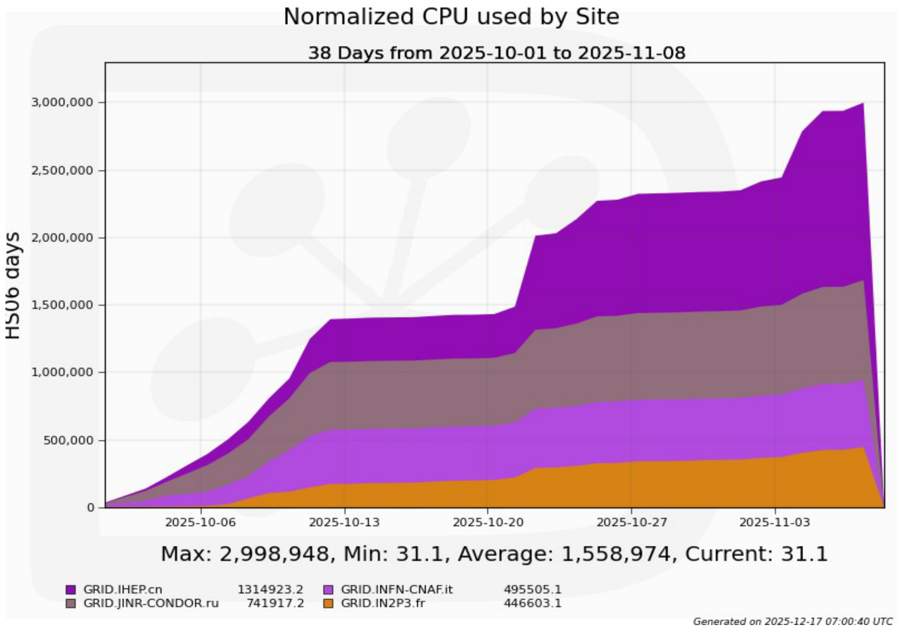

ReProdC

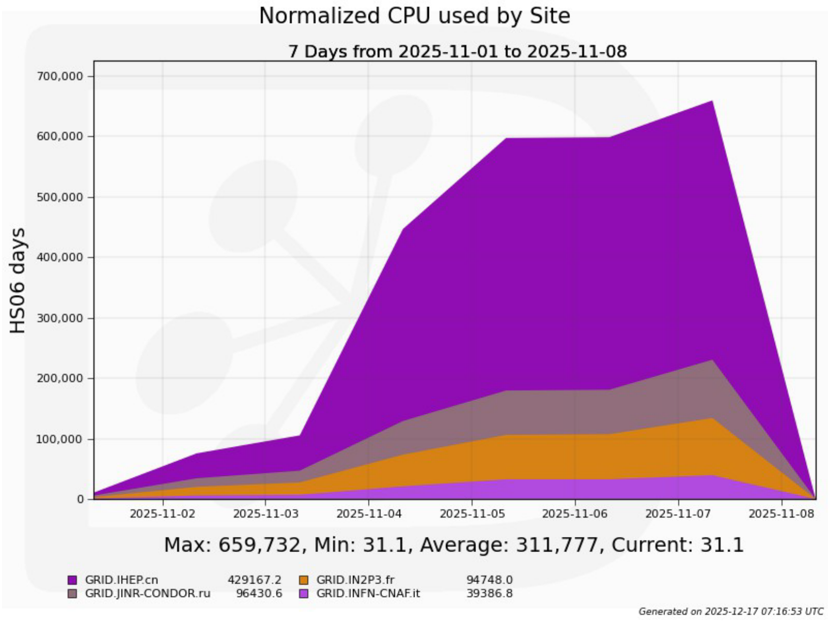

Identified Bottlenecks and Solutions

- Strains on the Configuration Service (CS) and JobStateUpdate when exceeding 10,000 concurrent jobs.
- Memory and Disk I/O spikes in the SandboxStore during simultaneous job completions due to SandboxStore contention.
- Solutions: Scaling out services to multiple machines, extending CS to Tier 1 sites like IN2P3.

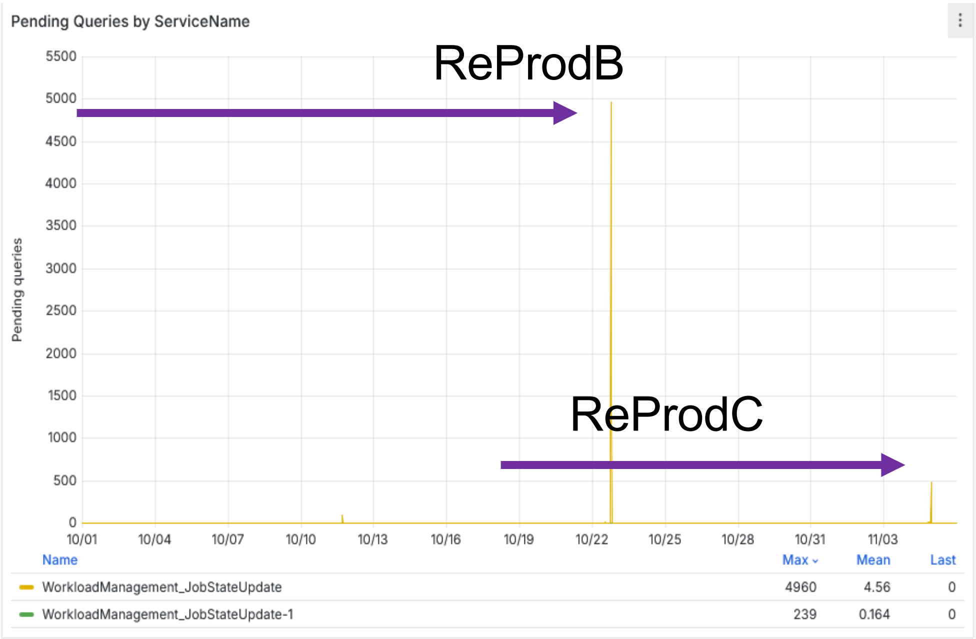

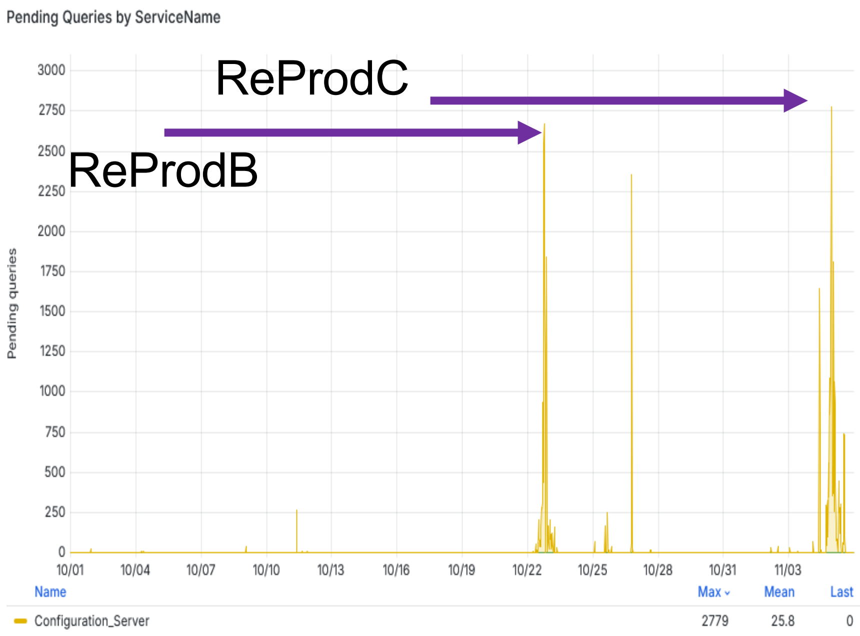

Summary and Future Plans

- DCI is stable and ready for full detector operation.
- Ongoing transition to full token-based AAI.
- Future: Supporting large-scale user analysis and metadata-driven queries.

Thank you

Any question?

Backup

JUNO physics goal

Neutrino Oscillation

- **Mixing matrix(PMNS):**

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} V_{e1} & V_{e2} & V_{e3} \\ V_{\mu 1} & V_{\mu 2} & V_{\mu 3} \\ V_{\tau 1} & V_{\tau 2} & V_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}$$

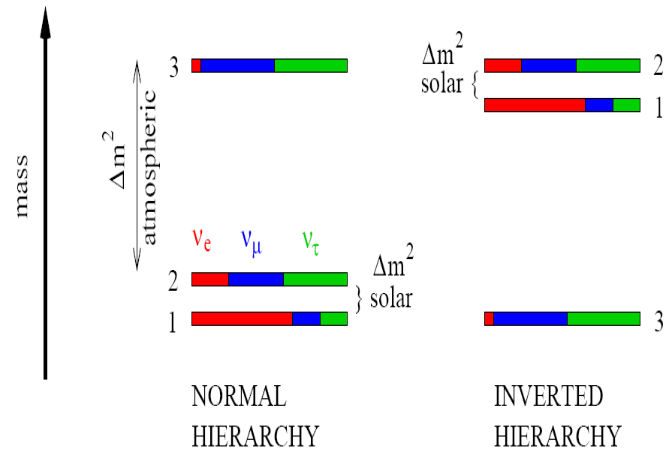

Known parameters: θ_{23}, θ_{12}, θ_{13}, $|\Delta m^2_{32}|$, Δm^2_{21}
Unknown parameters: sign of Δm^2_{32}, CP phase δ

1998 Atm. ν osc. θ_{23}, $|\Delta m^2_{32}|$
2002 Solar ν osc. θ_{12}, Δm^2_{21}
2012 reactor ν osc. θ_{13}

- **Determine Neutrino mass ordering***
- **Precise measure solar oscillation parameters**

* We already know ν1 lighter than ν2, which is the lightest neutrino, ν1 or ν3?

- **JUNO expects to achieve ~3% energy resolution @ 1MeV**
- **Neutrino Mass Ordering (NMO) @ 3σ with 6.5 years of DAQ time (reactors only)**
- **Combined reactor and atmospheric neutrino analysis would further improve the NMO sensitivity**
- **Best precision for the measurement of 3 oscillation parameters ($\sin^2\theta_{12}$, Δm^2_{21}, $|\Delta m^2_{31}|$) within 59 days, precision <0.5% in 6 years**

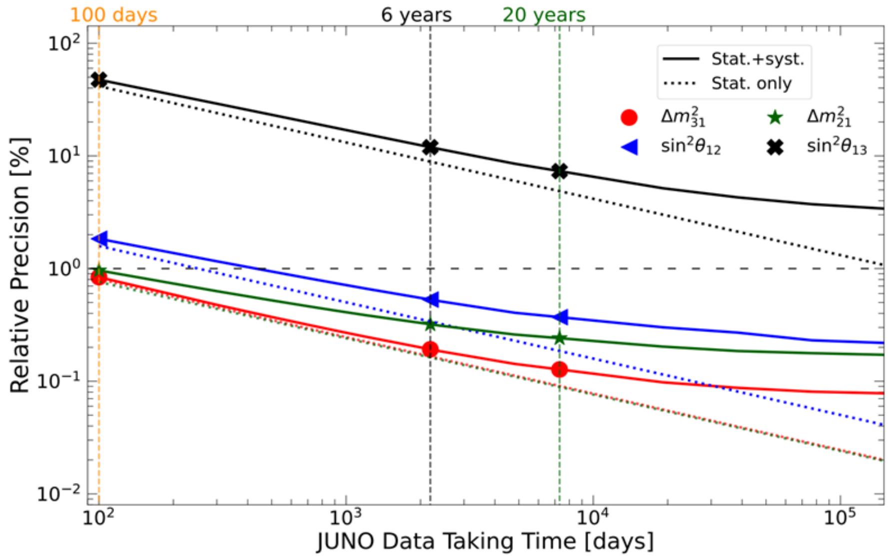

Chin.Phys.C 46 (2022) 12, 123001