JUNO distributed computing experience for early data-taking

Xiaomei Zhang (IHEP)

On behalf of the JUNO collaboration

MSCHEP2026, March 10-15 2026, BaiKal



Jiangmen Underground Neutrino Observatory (JUNO)

- A multipurpose Neutrino Observatory with a rich program in neutrino physics and astrophysics studying neutrinos over a large energy range
- Proposed as a reactor neutrino experiment for mass ordering in 2008.
- Construction started in 2015, and physics data taking started on August in 2025



>74 institutions, >700 collaborators

Asia: China (34), Taiwan,China (3) Thailand (3), Pakistan, Armenia

Europe: Italy (8), Germany (7), France (5), Russia (3), Belgium, Czech, Finland, Latvia, Slovakia, UK

America: Brazil (2), Chile (2), USA (2)

Real data volume and production chain (planned)

JUNO onsite | **IHEP** | **JUNO Data centers**

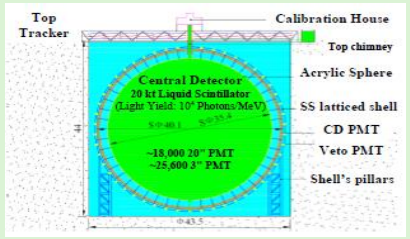

~40 GB/s

Trigger + OEC

RAW (t/q, waveform)

2 Gbps network (dedicated)

RAW data Volume: *~2 PB/year*
RTRAW data Volume: *~2 PB/year*

RAW (t/q, waveform)

Tape

100 Gbps network (LHCONE)

RTRAW

Sampling + Reconstruction

Data Quality Monitoring

Delay: ~hour

Reconstruction

ESD (KUP*)

Delay: ~half a day

one year data reconstructed within 2 months

RAW (t/q, waveform)

Tape

RTRAW

Reconstruction (new calibration)

ESD (ReProd*)

ESD (PP*)

DAQ Event Rate:
1kHz
Data Storage Rate:
~60 MB/s
Each file size:
5 GB/file
N (Events) per file:
83,333 (=5GB / 60MB/s * 1kHz)

First Reconstruction

Second Reconstruction

* KUP: Keep up production ReProd: Reproduction PP: Physics Production

Computing Model

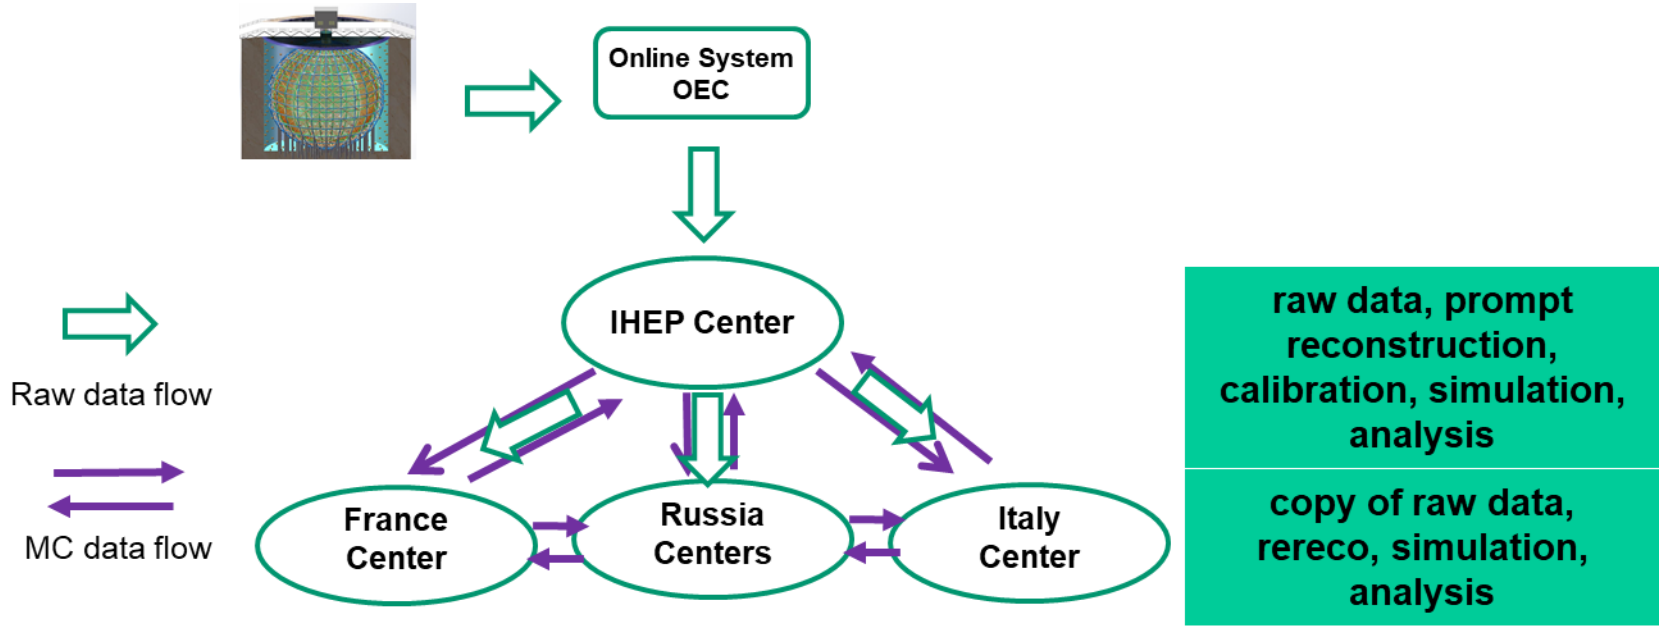

- Four data centers joined: IHEP, CC-IN2P3, INFN-CNAF, JINR
- **RAW data** flows from Online to IHEP which then distributes to other centers
- **KUP (Keep up production) and Calibration** will run in IHEP
- **MC Simulation, PP(Physics Production) and Analysis** runs in data centers
- **ReProduction** was newly added after data taking for reconstructions with new calibration parameters

Distributed Computing Infrastructure (DCI)

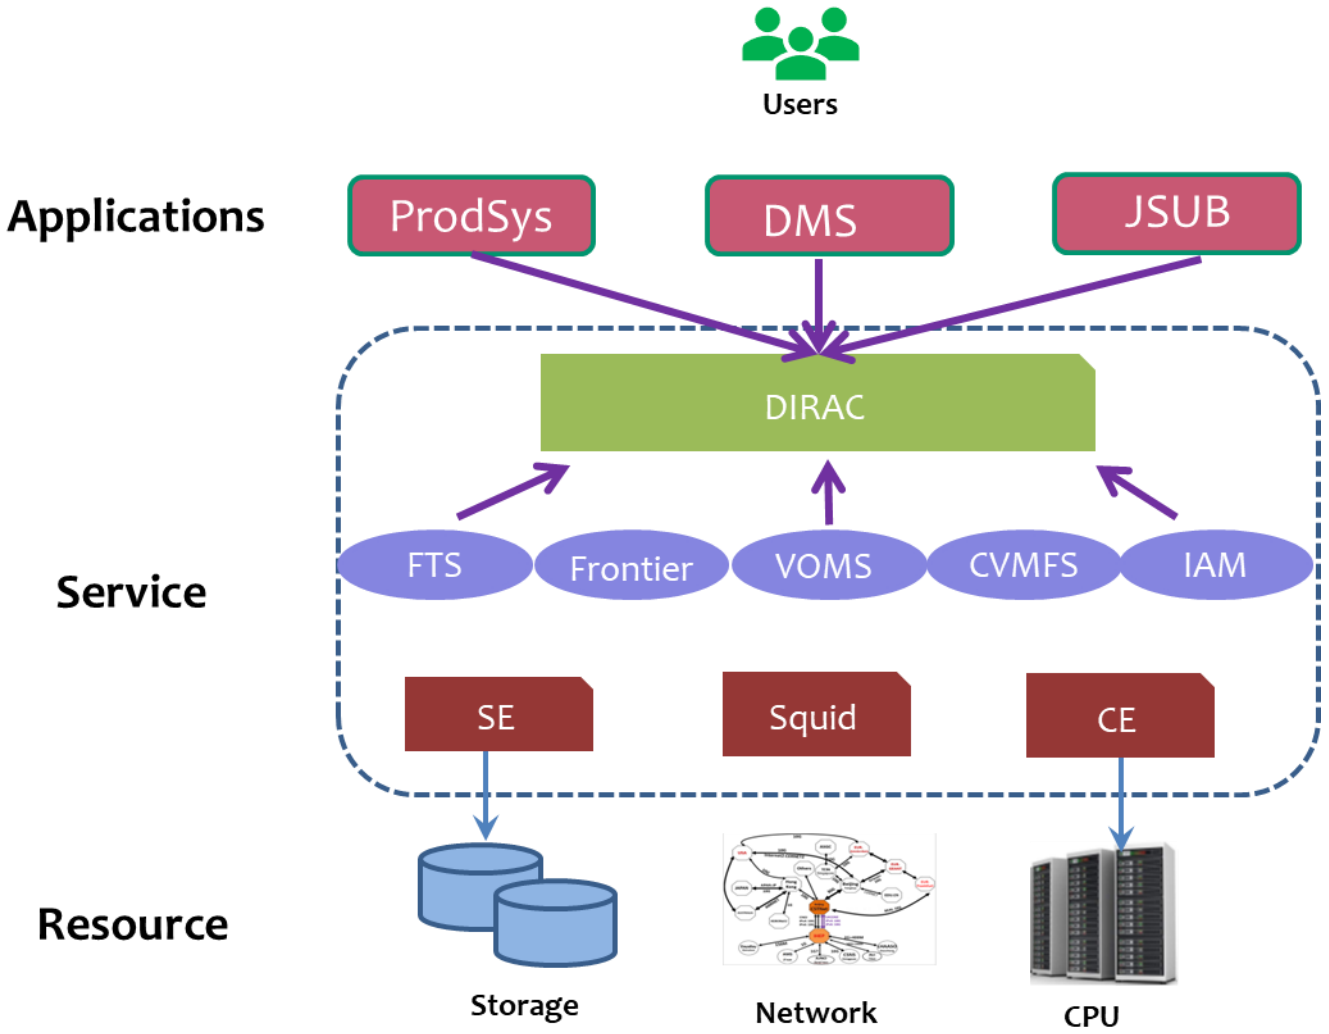

- **JUNO DCI was built based on DIRAC, use both WMS and DMS**
- **Use WLCG middleware as much as possible**
- **Develop JUNO-oriented application systems**
  - ⇨ Data production system (ProdSys)
  - ⇨ Data analysis tool (JSUB)
  - ⇨ Raw data transfer system
  - ⇨ Data operation tools

WMS: Workload management system, DMS: Data management system

Main activities in DCI

- **Data management and transfers**
  - ⇨ Official data transfer (raw, rec, sim)
- **Reconstruction**
  - ⇨ ReProduction (monthly)
  - ⇨ Physics Production (yearly)
- **MC production**
  - ⇨ Official and user Production
- **More supports are being considered**
  - ⇨ DQM
  - ⇨ MC Tuning
  - ⇨ Group Analysis
  - ⇨ User data transfers

Real Data Management

/juno/[DETECTOR]-[TYPE]/[CAMPAIGN]/[STREAM]/[RUN_BUCKET]/[RUN_GROUP]/[RUN]

- **DETECTOR**: juno | tao
- **TYPE**: raw | rtraw | kup | reprod | pp | mcsim | mcrec
- **CAMPAIGN**: P25A | J25 | T25 etc.
  - P25A: for Physics production. P[YY][S]. [S] is suffix to represent different campaigns in this year.
  - J25.5.0[.S]: for KUP. JUNOSW version + iteration
  - T25.5.0[.S]: for KUP. TAOSW version + iteration
- **STREAM**: global_trigger, mm_trigger, etc
- **RUN_BUCKET** (for example, 1k runs)
  - such as `00001000` (run 1000-1999), `00002000` (run 2000-2999). `bucket = "%05d000"% (run//1000)`
- **RUN_GROUP** (for example, 100 runs)
  - such as `00001000` (run 1000-1099), `00001100` (run 1100-1199). `group = "%06d00"% (run//100)`
- **RUN**: run number.

◆ Three main data types: RAW, RTRAW, ESD

- ⇨ **RAW:** Byte-stream format, from online
- ⇨ **RTRAW:** ROOT format, sorted in IHEP cluster
- ⇨ **ESD:** Event Summary data, output from KUP and PP

◆ IHEP, JINR, CNAF have one copy of full data

- ⇨ RAW directly saved in TAPE (5GB/file)
- ⇨ RTRAW (5GB/file) and ESD (1-2 GB/file) in disk

◆ Data registered in DCI File Catalog

- ⇨ Folder structure and meta data arranged to meet offline needs
- ⇨ Files can be grouped into datasets by those meta data

◆ miniESD is newly introduced to reduce I/O pressure

- ⇨ <1GB/file, minimal footprint, optimized for rapid transfer and broad accessibility

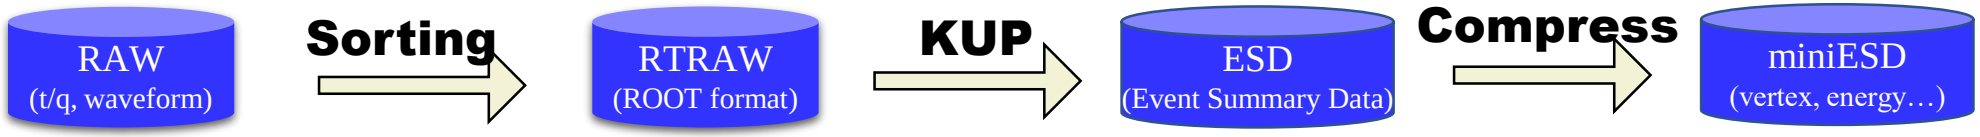

Data Growth Exceeds Projections in early data-taking

- **Data rate surpassed planning: 60 MB/s → 90 MB/s**
- **RAW data reached 1.6 PB within 7 months of data-taking (2PB one year)**
  - 37% volume increase
- **Disk and tape space tighter than expected**
  - **CNAF & IN2P3** usage exceeded **90%, JINR 74.2%**
  - **IHEP** expanded disk and tape capacity post data-taking

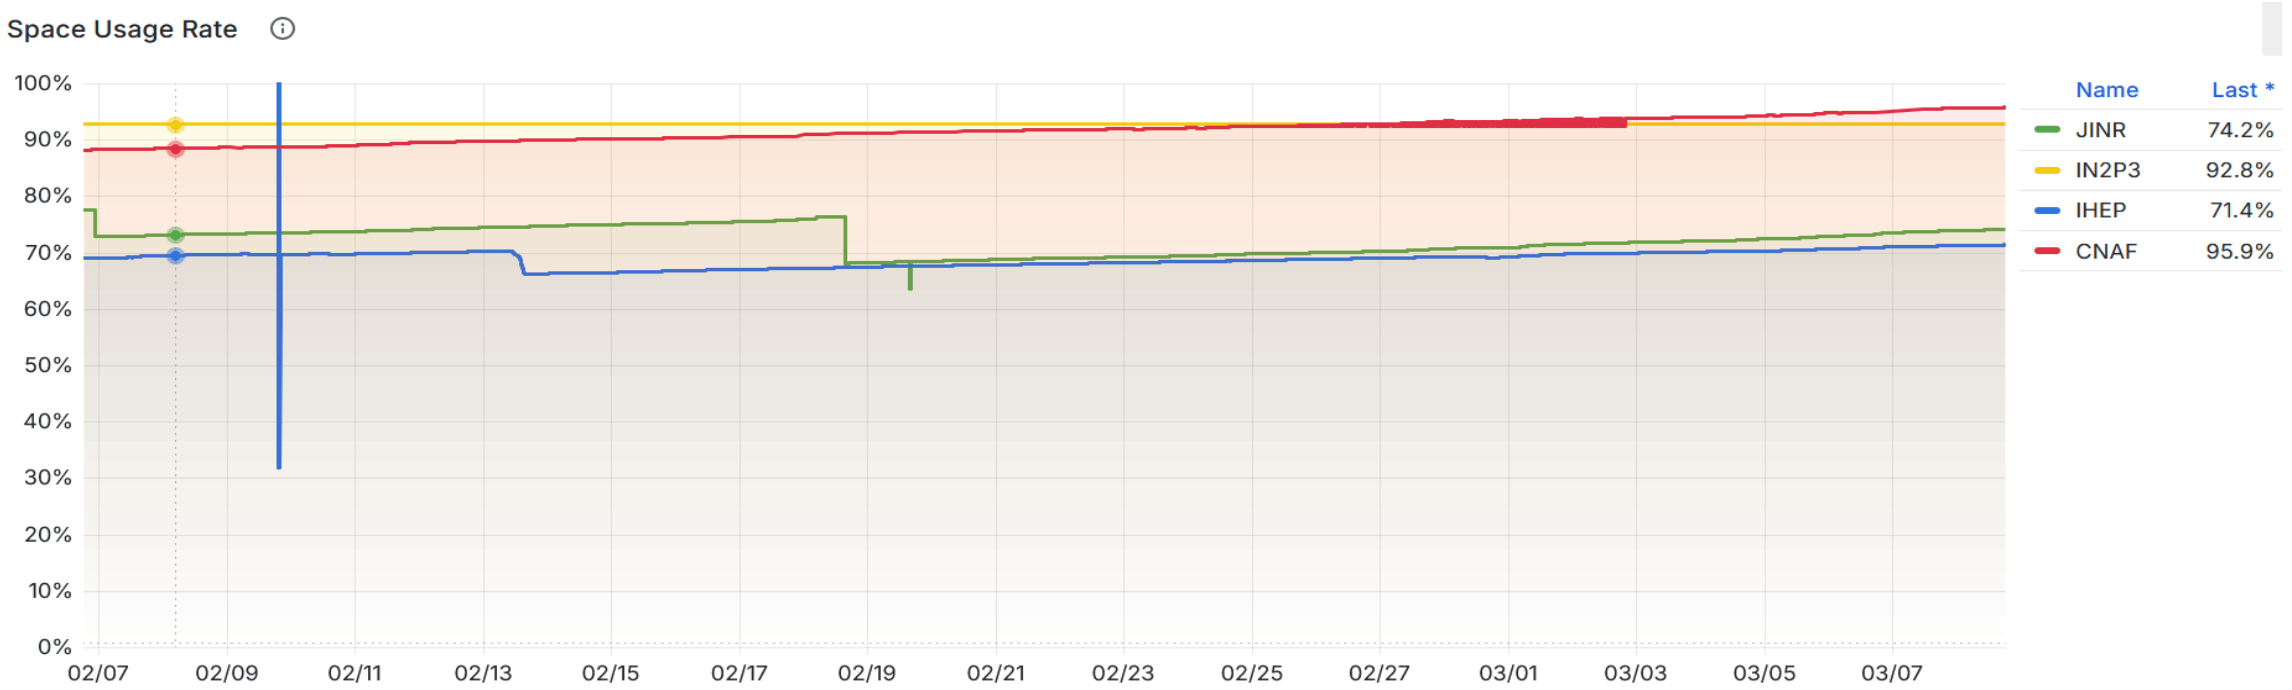

Challenges of DCI data management

◆ Data Scale in DCI File Catalogue

- Over 8 PB total registered data
- More than 31 million files

◆ Data Growth Challenges

⇨ Performance of File Catalogue

- As data volume grows, file catalogue queries and updates become slower
- More frequent data operations can become bottlenecks for large-scale workflows

⇨ Consistency and Cleanup in distributed computing env

- Large scale Consistency check and Cleanup is not just a simple "push-button" operation with storage systems
- Need Cross-group cooperation and essential for defining retention policies with tools
- Consistency checks and Cleanup on huge datasets are time-consuming and never to be easy

◆ Hybrid of Grid and Local usage in JUNO introduces complexity

- File metadata stored in both local databases and grid catalogues
- Synchronization between local and grid systems is critical
- WLCG DMS was designed for pure grid environments, offering limited guidance for hybrid setups

Real data distribution with MessageQueue 1/2

- **Multiple real data activities**
  - SPADE->RAW
  - RAW-> IHEP Tape
  - RAW->RAW2RTRAW-> RTRAW
  - RTRAW -> KUP -> ESD
  - RAW, RTRAW, ESD -> DCI
- **A message-driven system based on Kafka is introduced to allow smooth communications among those activities**

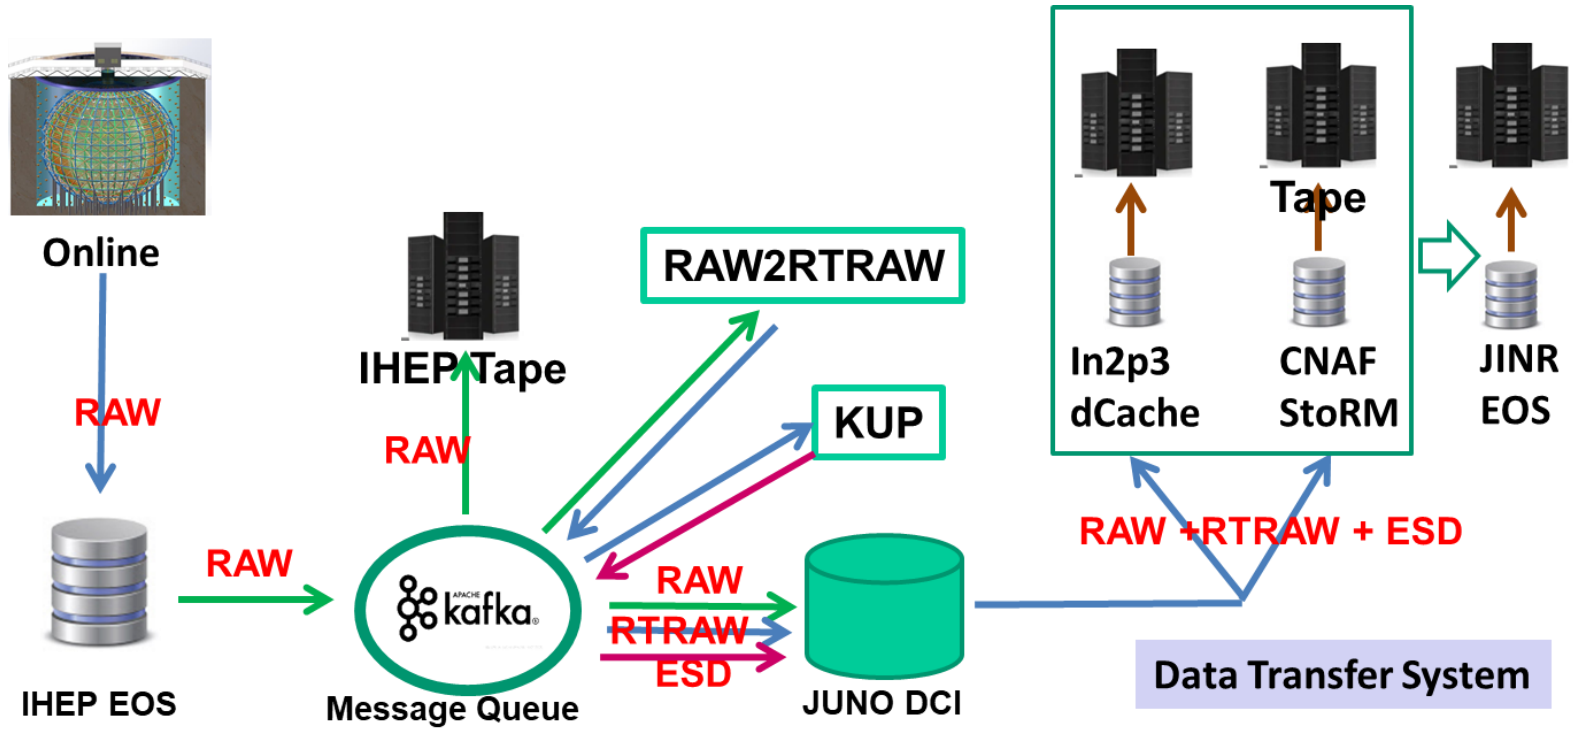

Real data distribution with Message Queue 2/2

Advantages

- Highly flexible and extensible — easy to integrate new activities
- Enables smooth automation of data workflows

Challenges and Solutions

- Issues with producers can propagate to all consumers
  - Some problems harmless to producers may severely impact consumers
- Certain failures require manual intervention, which can be disruptive
  - These failures caused by peak load pressure on our critical systems, such as the computing cluster, storage, and network
  - Need robust monitoring and alerting systems to detect issues early
- Solutions: smarter tools for automatic detection and recovery

DCI automatic system for global data transfers

- **Automatic system** developed on DCI DMS
  - ⇨ **DIRAC DMS** manages central scheduling for datasets
  - ⇨ **FTS** takes care of file-level replications
- **Monitoring stack**
  - **ElasticSearch** — Central log aggregation
  - **Grafana** — Real-time dashboards for FTS
- **Accounting using DIRAC**

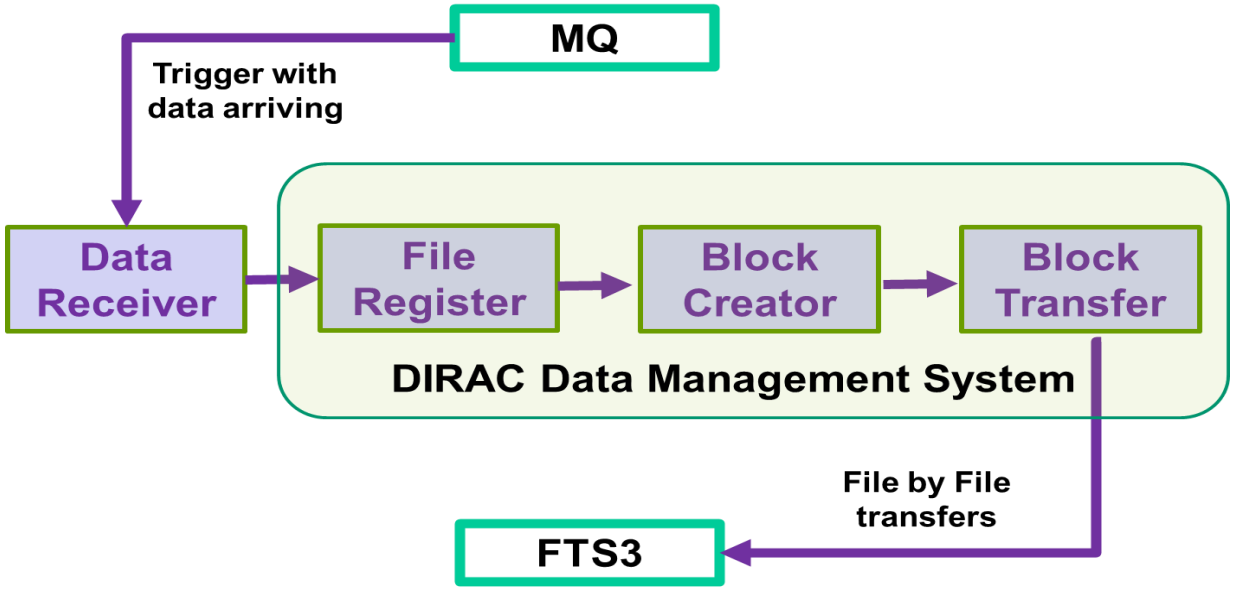

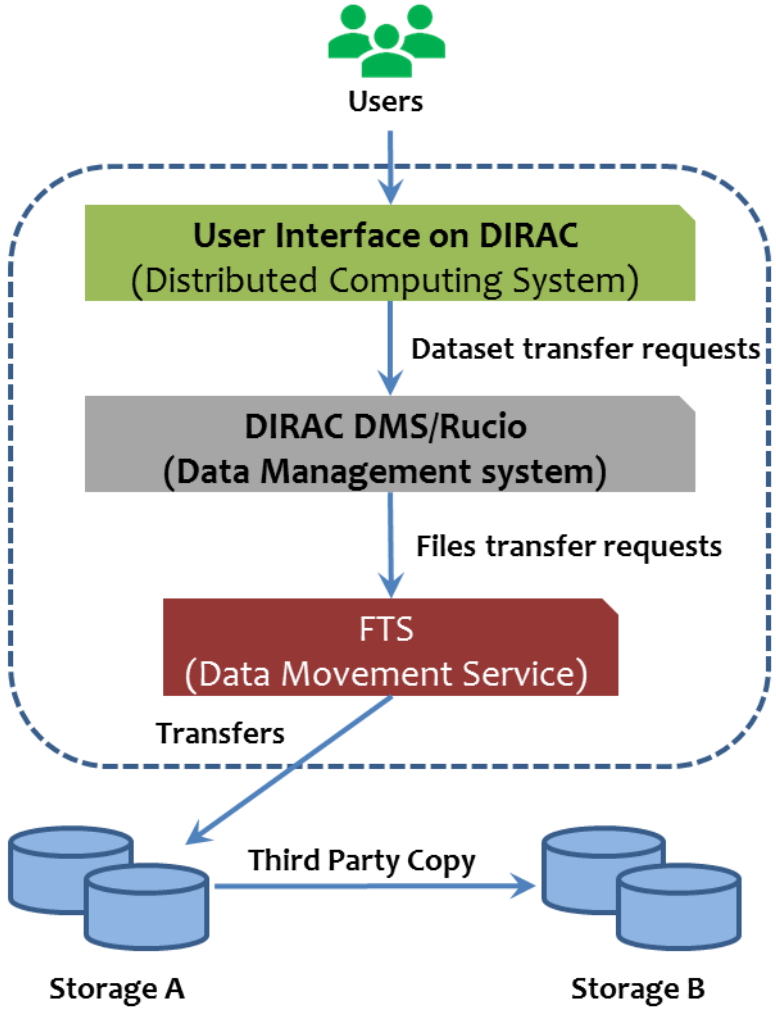

Status of data transfers

- **Global data transfers operate smoothly throughout the year**
  - ⇨ Peak transfer speeds exceed 30 Gb/s
- **~7.3 PB of real data transferred to CNAF, JINR and IN2P3**
  - ⇨ ~3.8 PB to tape storage
  - ⇨ ~3.4 PB to disk storage
  - ⇨ ~0.1 PB to IN2P3 (TT data)

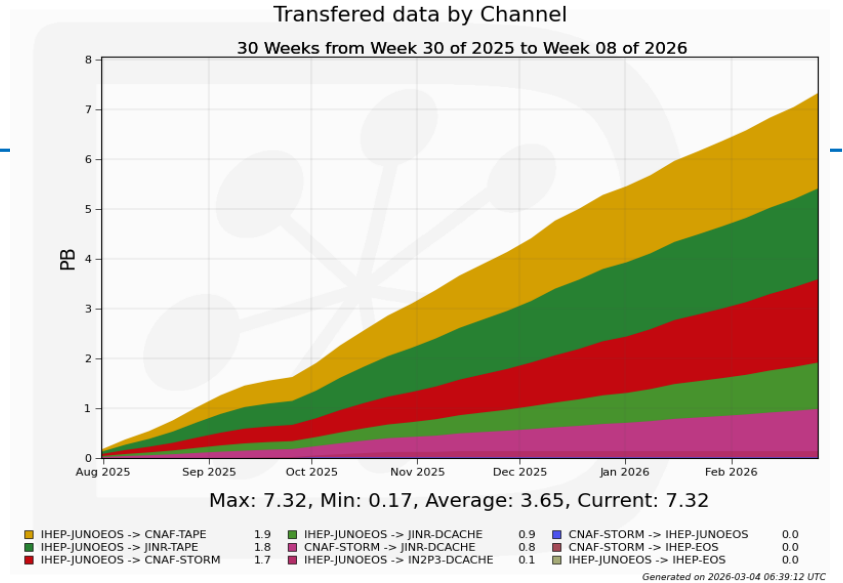

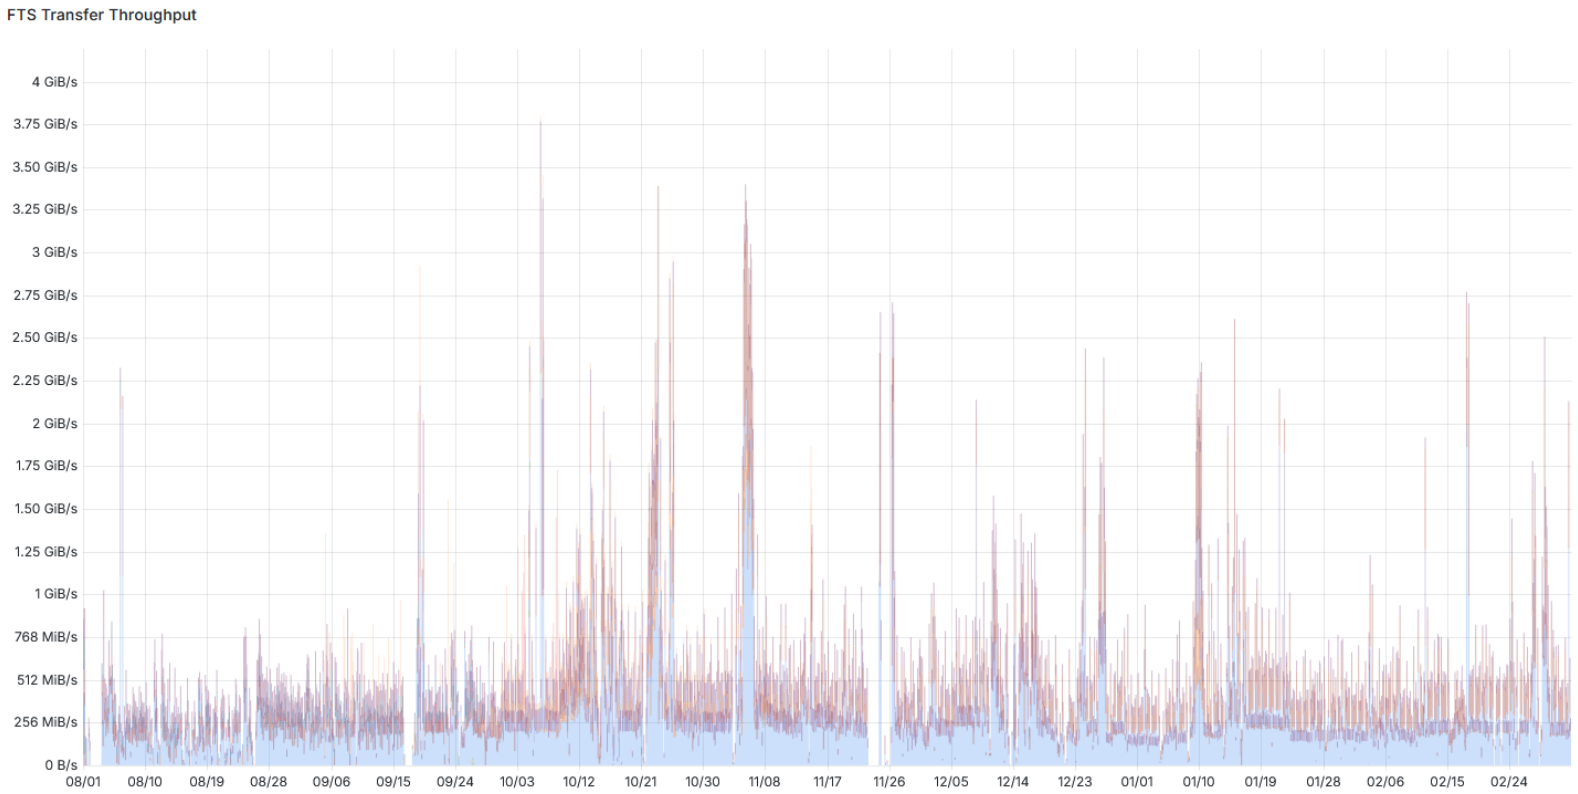

Name	Max	Mean	Last *
CNAF to CNAF	36.6 MiB/s	8.04 MiB/s	68.8 KiB/s
IHEP-STORM to CNAF	37.5 MiB/s	5.00 MiB/s	176 KiB/s
IHEP to CNAF	2.31 GiB/s	270 MiB/s	207 MiB/s
JINR to CNAF	673 MiB/s	69.8 MiB/s	2.38 MiB/s
JINR-EOS to CNAF	39.8 MiB/s	10.4 MiB/s	2.33 KiB/s
IN2P3 to CNAF	189 MiB/s	15.8 MiB/s	543 KiB/s
CERN to IHEP-STORM	30.8 MiB/s	1.14 MiB/s	1.42 MiB/s
CNAF to IHEP	1.63 GiB/s	119 MiB/s	505 KiB/s
IHEP-STORM to IHEP	51.2 MiB/s	14.3 MiB/s	1.68 MiB/s
JINR to IHEP	5.31 MiB/s	2.50 MiB/s	1.59 MiB/s
JINR-EOS to IHEP	65.7 MiB/s	14.0 MiB/s	1.98 MiB/s
IN2P3 to IHEP	396 MiB/s	45.5 MiB/s	24.1 MiB/s
CNAF to JINR	1.60 GiB/s	133 MiB/s	15.9 MiB/s
IHEP to JINR	204 MiB/s	55.6 MiB/s	937 KiB/s
IHEP to JINR-Tape	1.59 GiB/s	106 MiB/s	58.7 MiB/s
JINR-EOS to JINR-Tape	548 KiB/s	547 KiB/s	548 KiB/s
IN2P3 to JINR-Tape	66.0 MiB/s	2.56 MiB/s	271 KiB/s
IN2P3-xrd to IN2P3-xrd-egee	80.5 MiB/s	80.5 MiB/s	80.5 MiB/s
CNAF to IN2P3	2.11 MiB/s	2.11 MiB/s	2.11 MiB/s
IHEP to IN2P3	1.28 GiB/s	71.3 MiB/s	1.33 MiB/s
IN2P3 to IN2P3-Tape	1.24 GiB/s	546 MiB/s	10.5 MiB/s

Status of IHEP->JINR raw data transfers

- **~3.5PB data transferred to JINR disk and tape after data-taking**
- **Channel of IHEP → JINR dCache is unstable since data-taking**
  - Use both **IHEP → JINR** and **CNAF → JINR** to reduce pressure on this channel
  - Performance degraded further since October, 2025
  - Workaround: Route transfers via **IHEP → CNAF → JINR**
- **IHEP → JINR Tape uses LHCOPN and not affected**

Transfer Efficiency

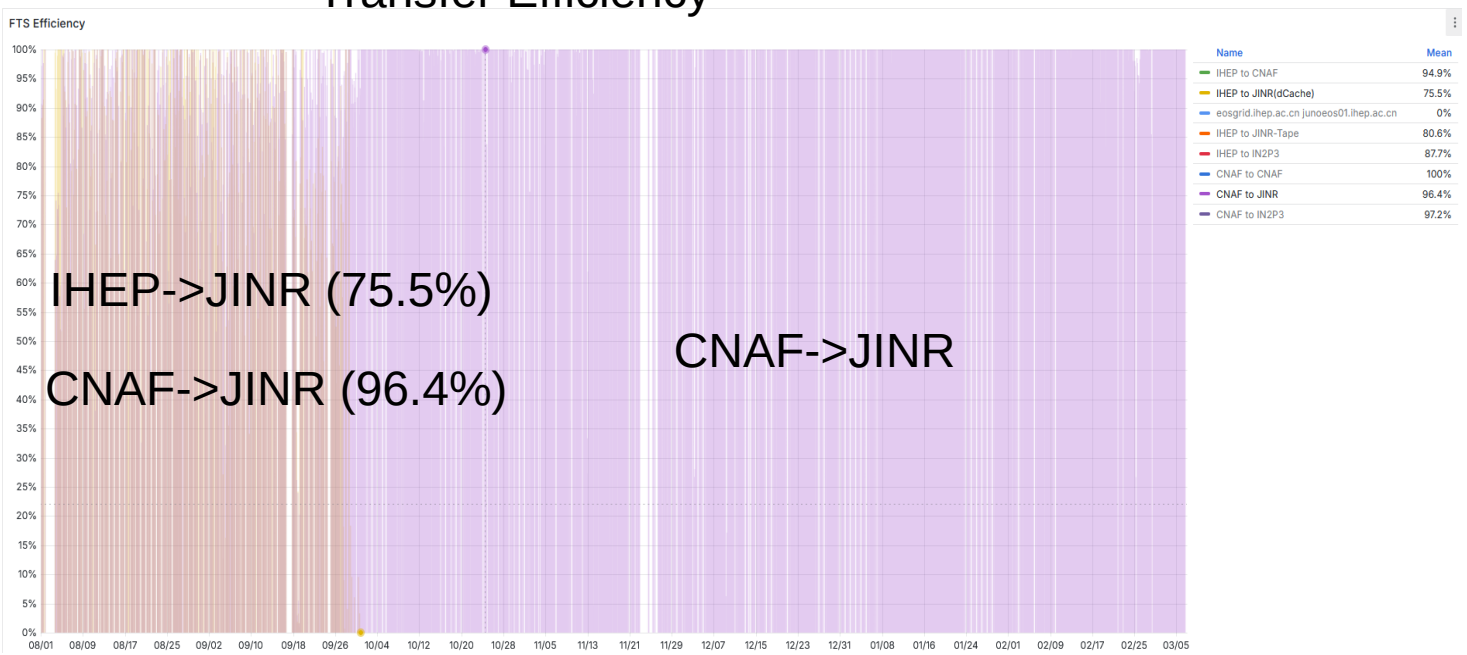

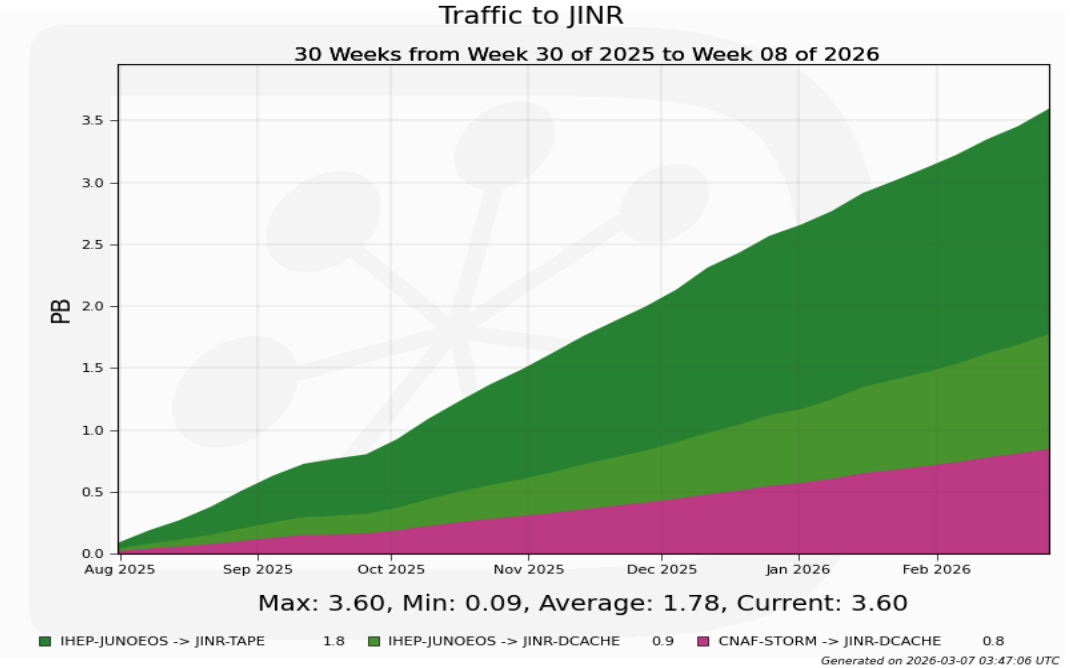

Three rounds of ReProd in 2025 after data taking

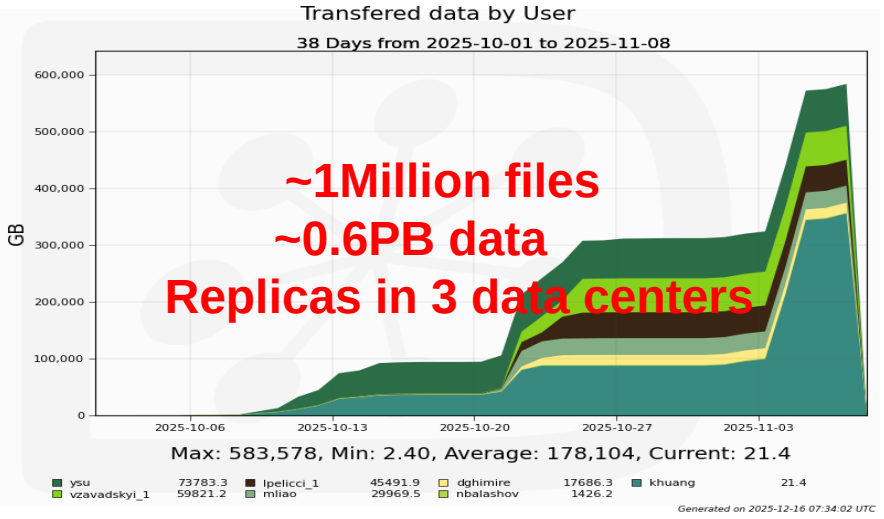

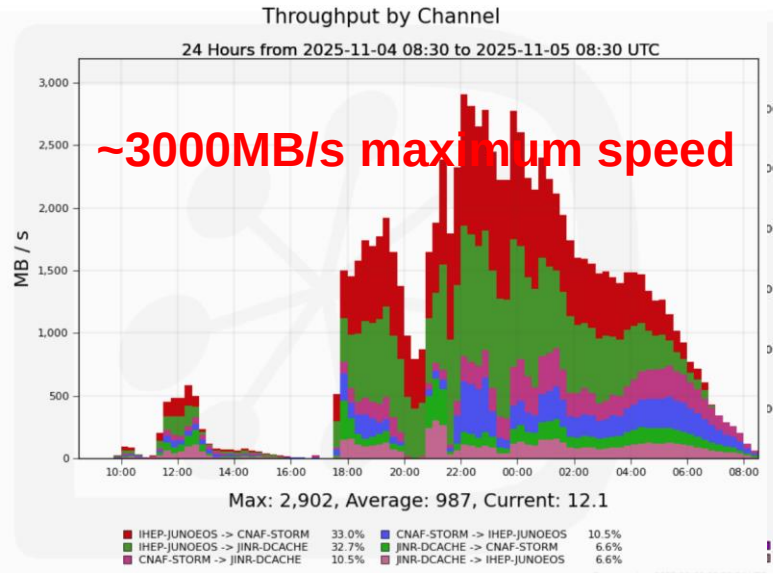

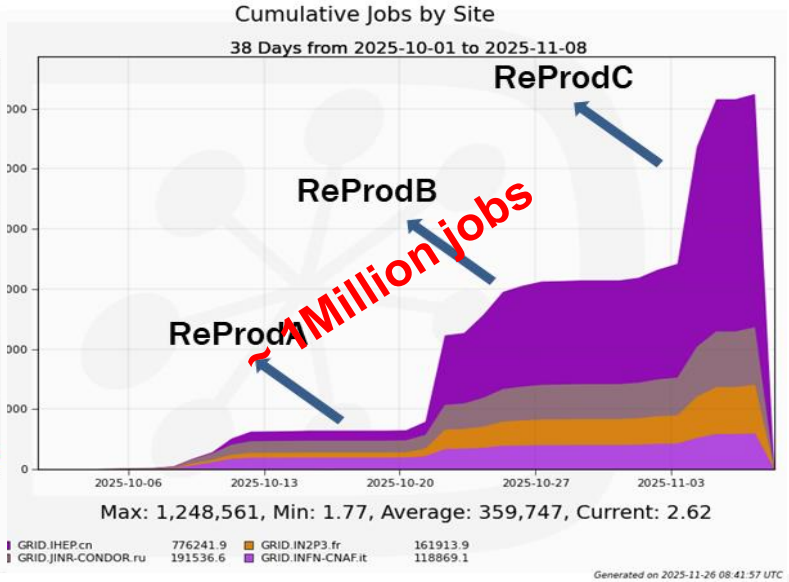

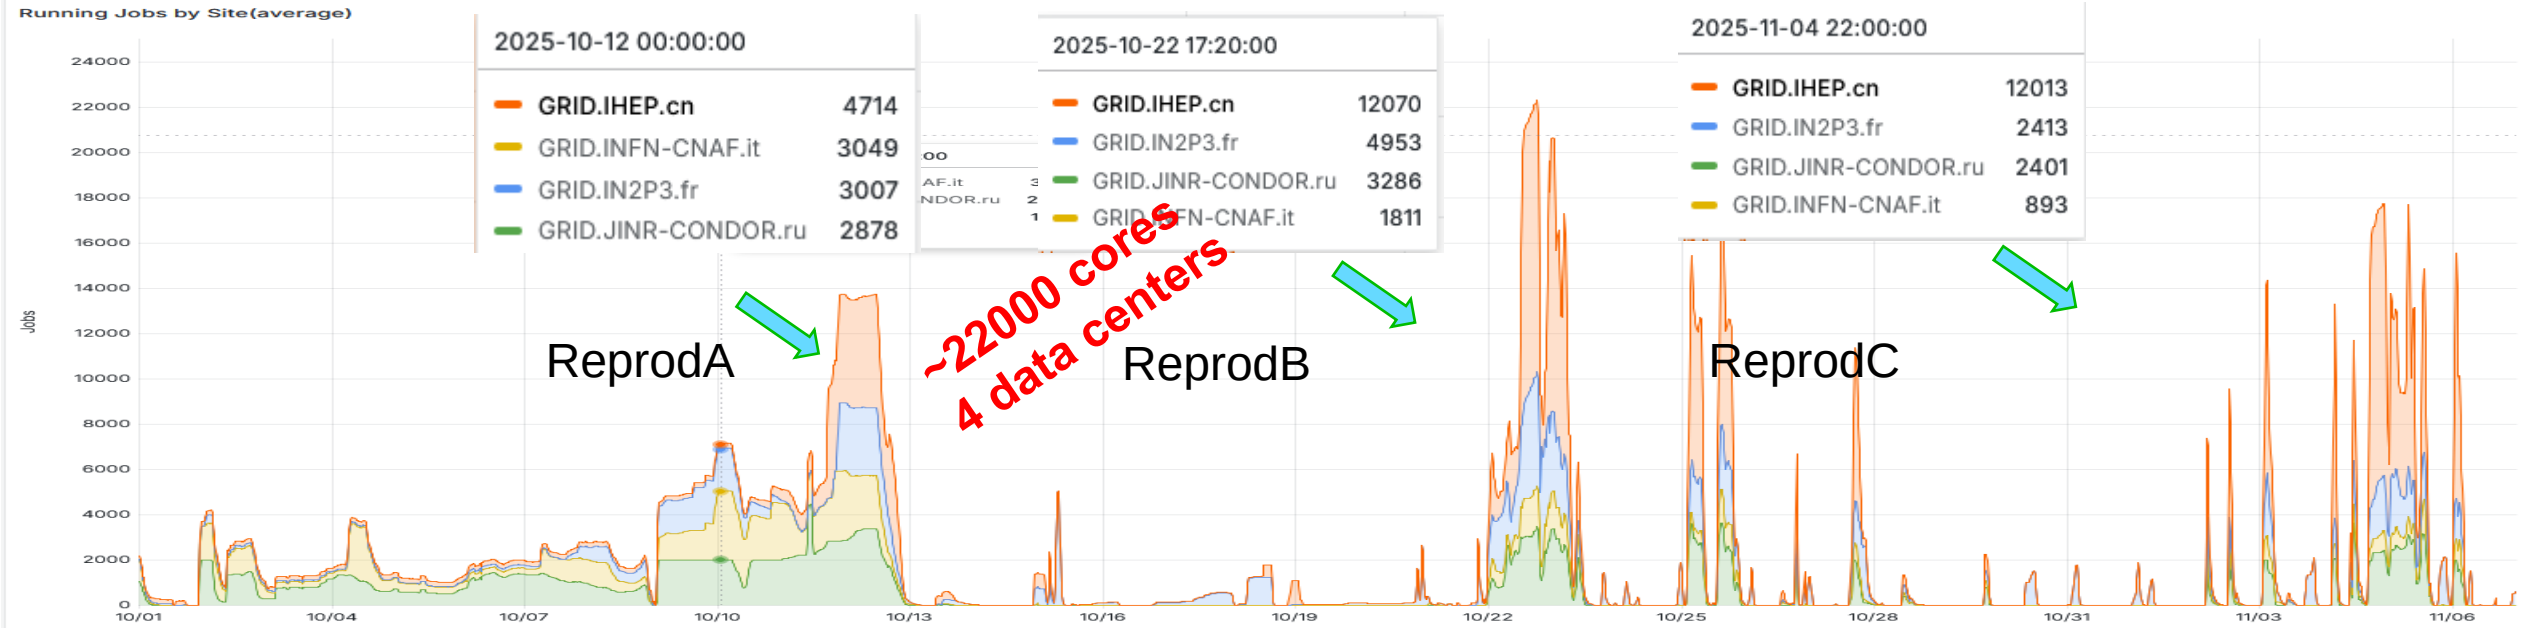

CPU consumption and Site contribution for ReProd

- **CPU consumption**
  - ReProdA + ReProdB + ReProdC = 8kHS06 year
  - ReProdC (two month's data) used ~1.8kHS06 year
- **Normalized CPU contributions by Site**
  - IHEP : JINR : CNAF : IN2P3 ~= 13 : 7 : 5 : 4.5
  - JINR contributes ~24% CPU resource
  - IHEP using opportunistic resources during ReProd

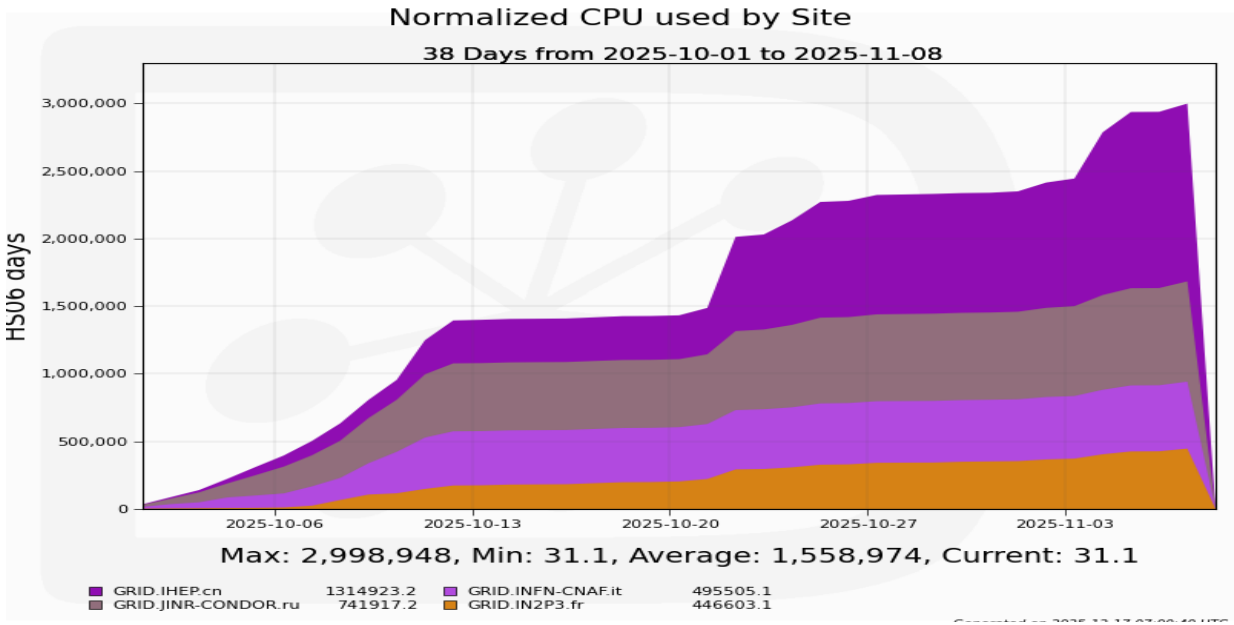

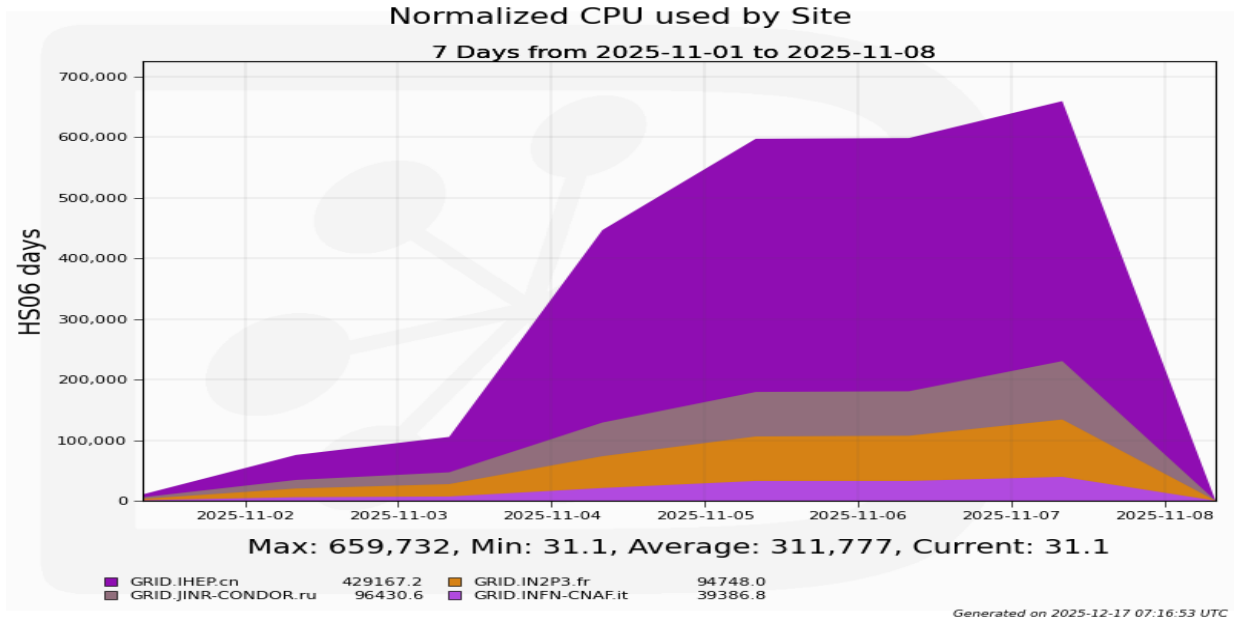

Current procedures of reproduction in DCI

- **Task Definition: Shifters create ini files, describing production parameters**
- **Task Submission: submit through ProdSys API**
- **Automated Processing: Job splitting, Chain creation, Jobs and transfers submission**

✓**The automatic system handles the majority of operations**
✓**Human intervention (shifters + experts) focuses on exception handling**
✓**Next Step:** Enhance exception handling through smarter automation

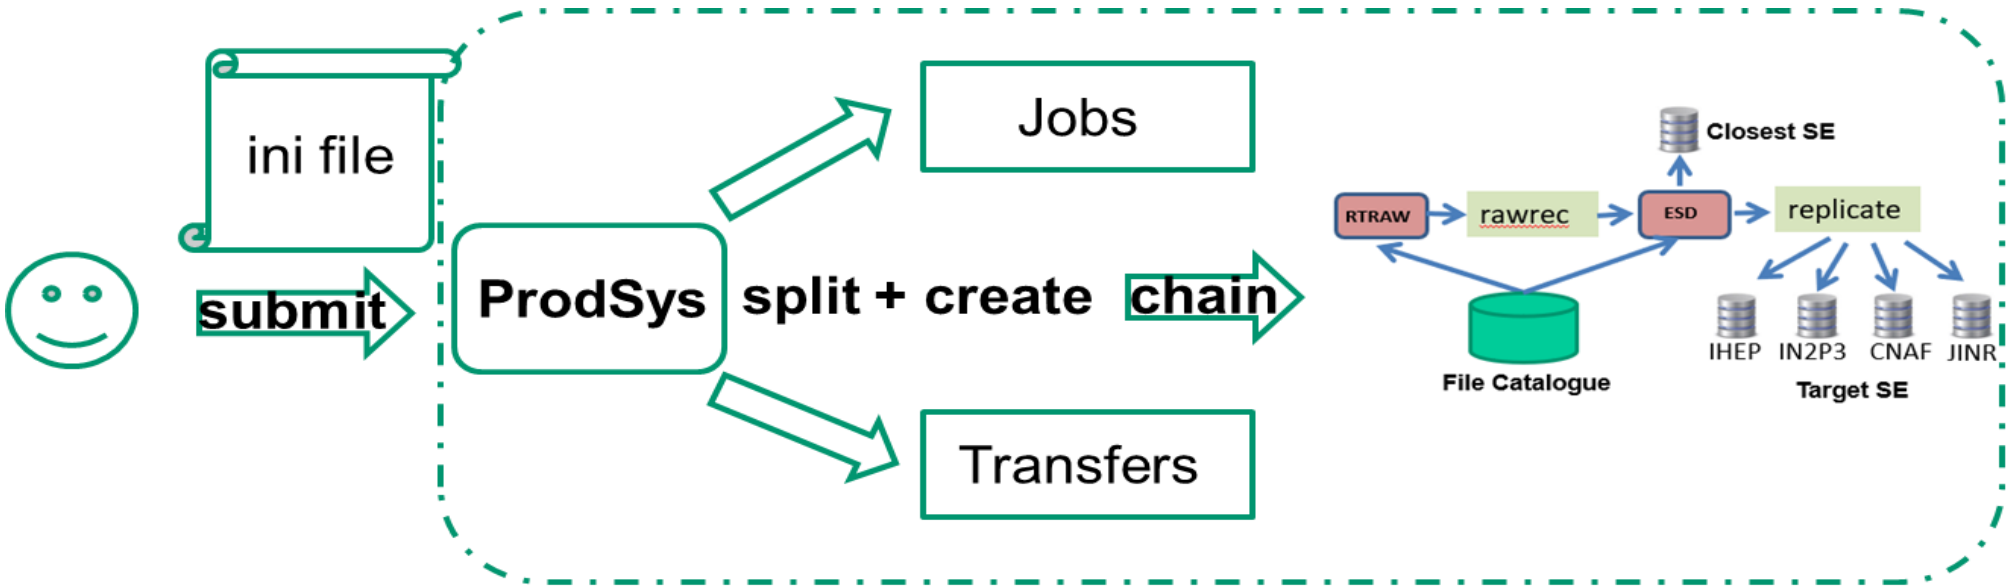

Success rate and error analysis

- **Success rate**
  - ReProdA: 98%, ReProdB: 95.3%, ReProdC: 97%
  - Errors are evenly distributed across all sites among reproductions
- **Main reasons for failed jobs**
  - Pressure of the system (central services, SE, CE)
  - Network thrashing
- **All these errors recoverable after rescheduling**

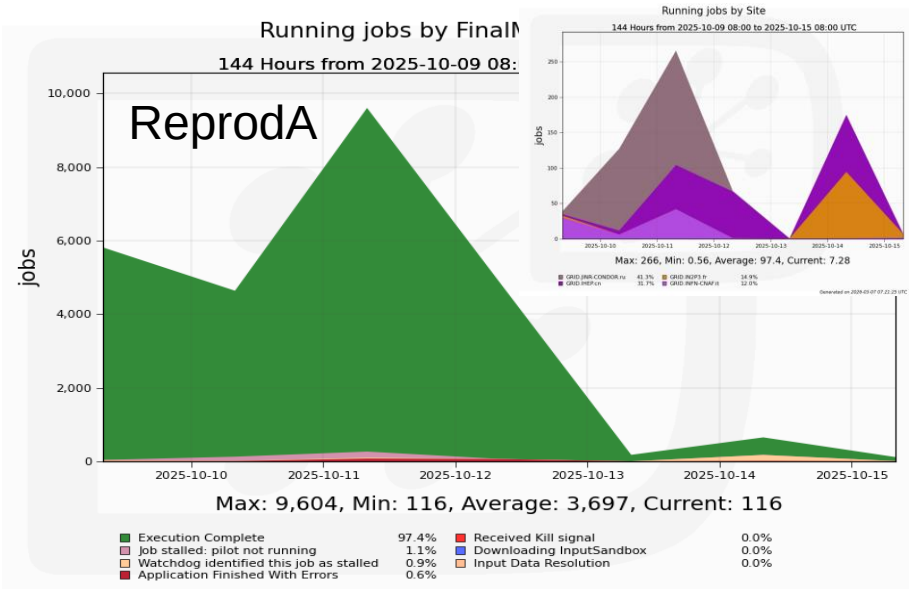

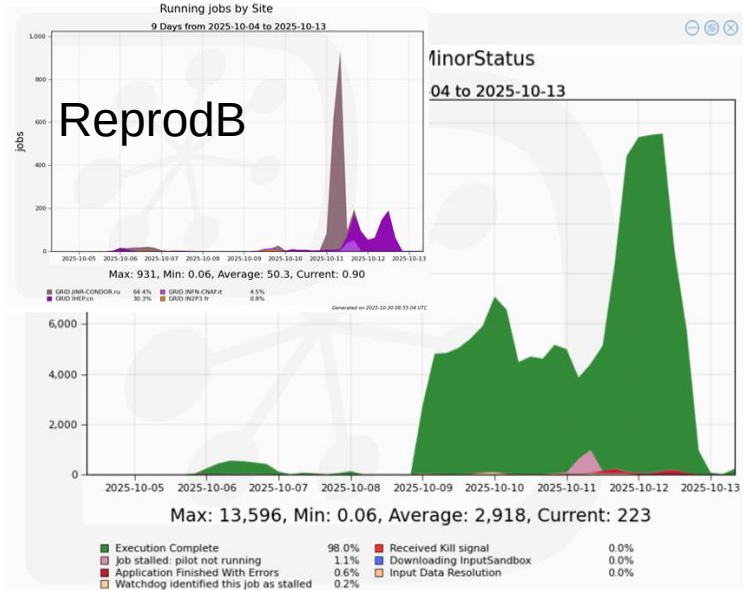

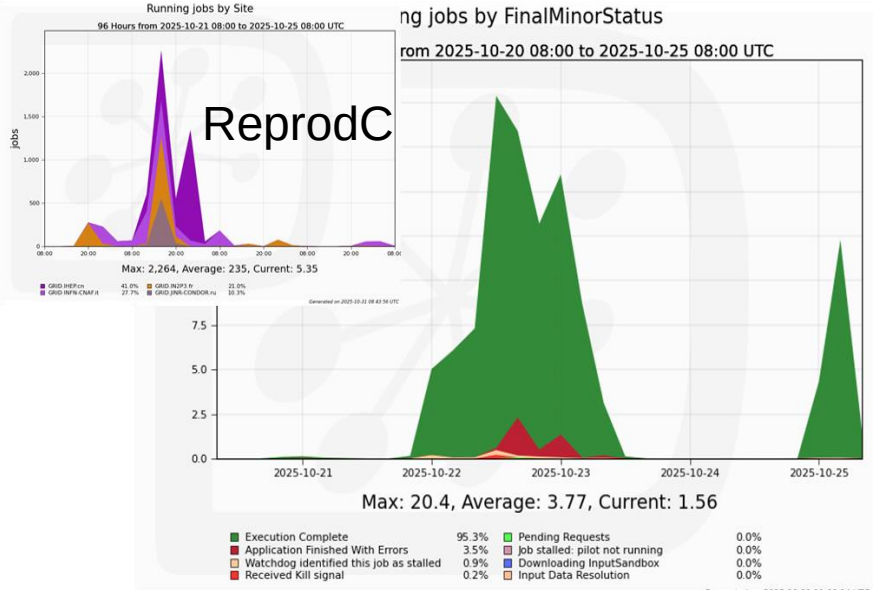

System Performance Observations

- **RepProdA: No significant system pressure detected**
- **RepProdB & RepProdC: Experienced system crashes under heavy load**
- **Main Cause: Short job durations combined with increased resource demands placed considerable strain on DCI services**

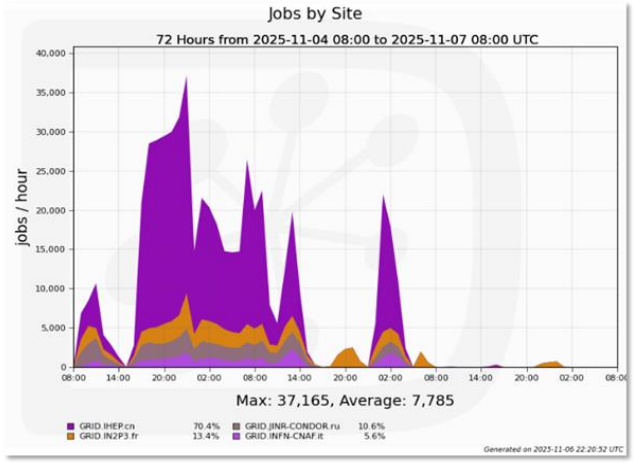

- DIRAC service at IHEP for JUNO, BES III, CEPC, ...
  - Example: JUNO production data processing in November 2025:
    - Up to 10Hz rate of executed jobs
    - **Two times more than the LHCb's average value !**

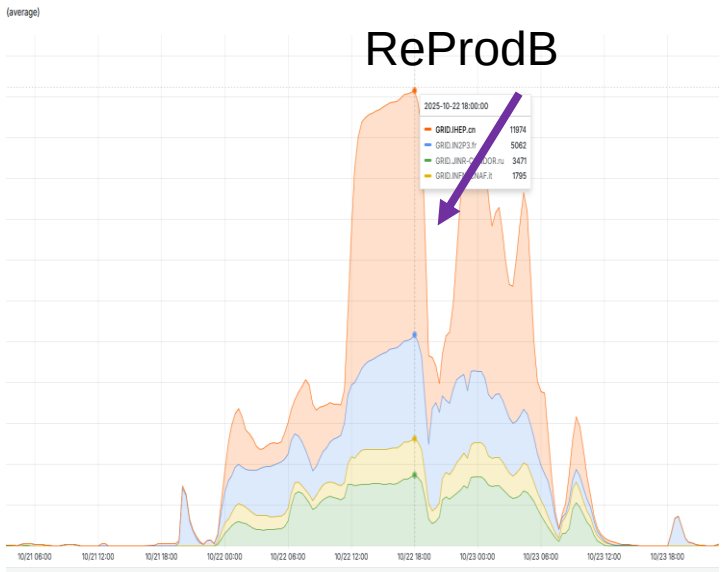

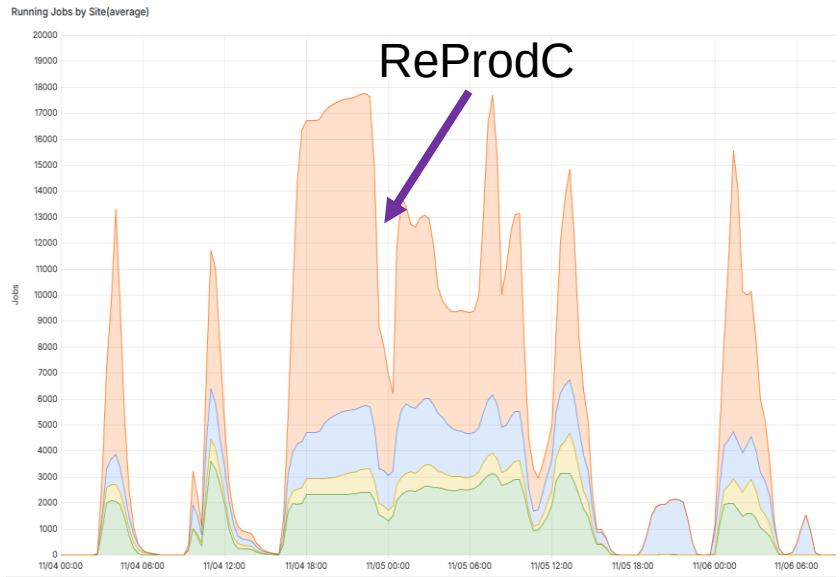

System pressures and bottlenecks

- **Configuration Service bottleneck: Performance degrades**
  - Currently 2 CS instances at IHEP, Planned extension: Deploy additional instance at IN2P3
- **SandboxStore instability: High memory peaks lead to crashes in ReProdD**
  - Requires investigation and optimization
- **JobStateUpdate Congestion in ReProdB**
  - Scaling out successfully resolved congestion in ReProdC
  - Similar approach may benefit other services under high load

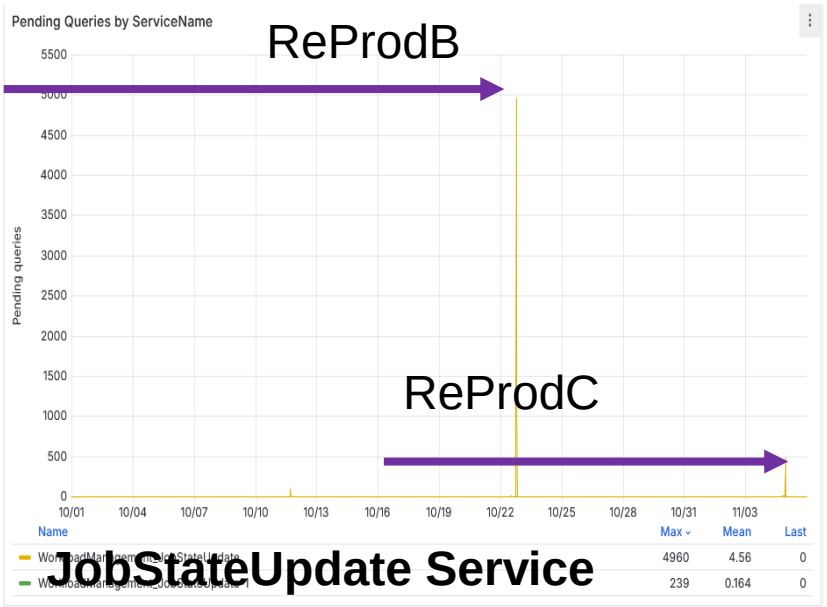

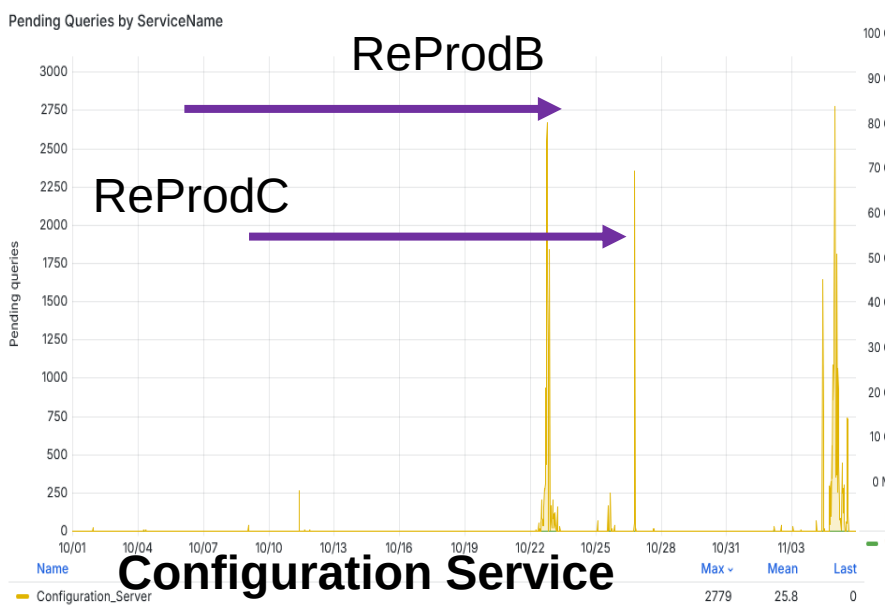

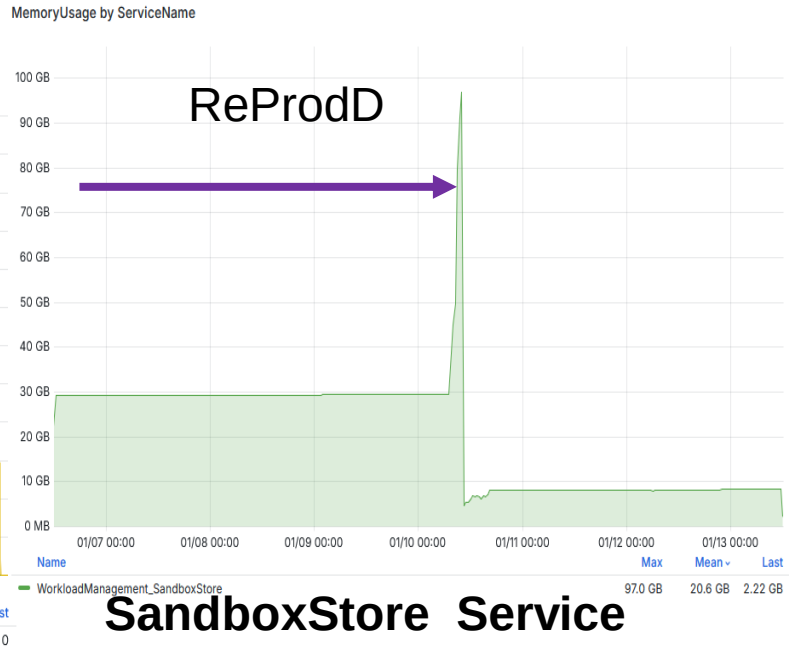

Data transfers during ReProd

- **Total 583TB data are replicated to three data centers**
- **25A have some delays, caused by IHEP->JINR network issues**
- **25B and 25C got improvements with multi-hops using CNAF as “Hub”**
  - Significant improvements made by Multi-hop transfers 73.9%->89.7%->98.8%

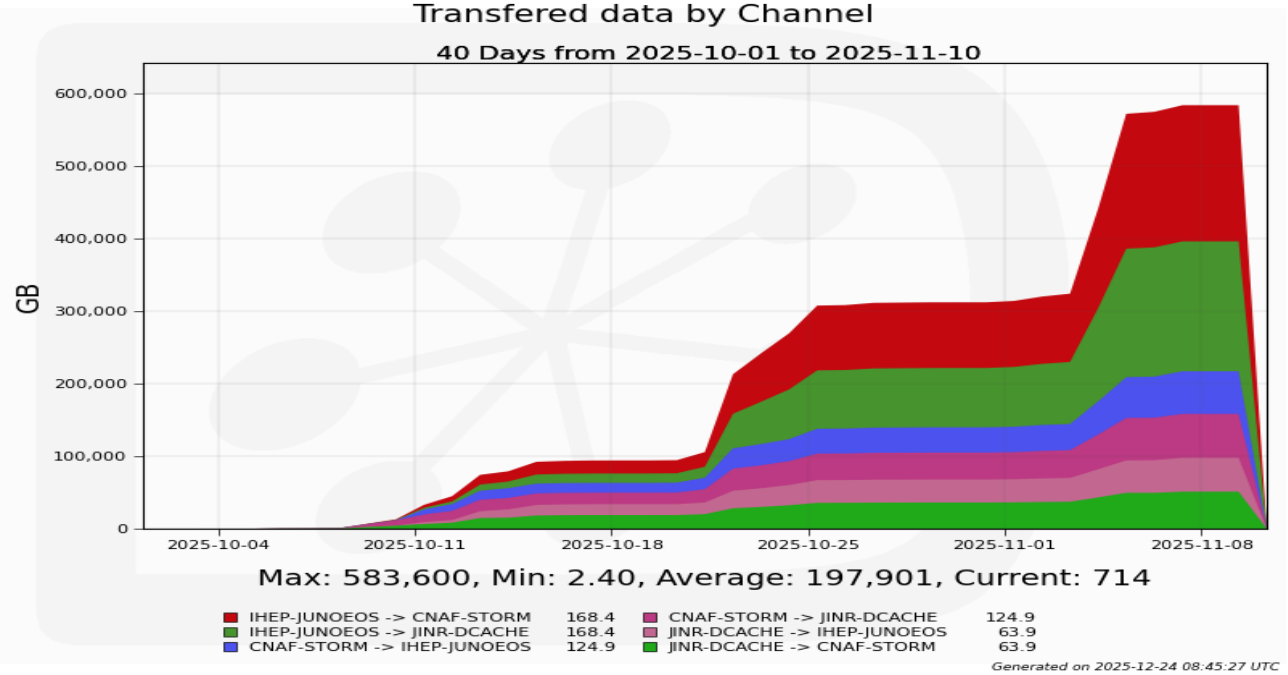

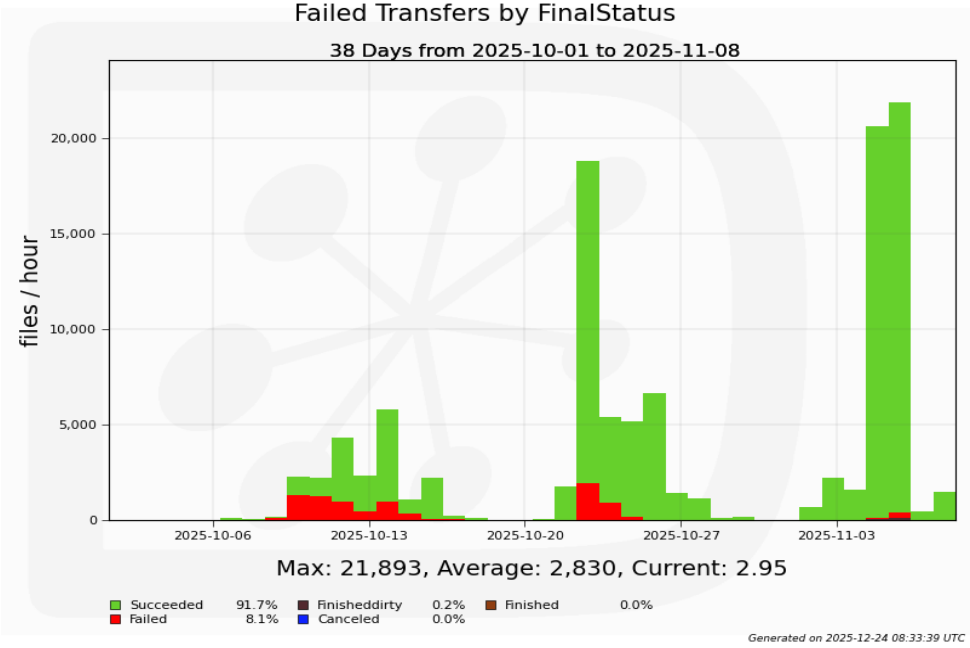

Uneven Resource Utilization after data taking

- **Total pledge resources (red line) are ~4500 cores**
- **Three ReProductions takes more beyond the pledge line**
- **Few resource usage between ReProd25 and ReProd26**
- **Possible Solutions besides waiting for MC productions:**
  - Flatten production workflows, Utilize opportunistic resources during peak times

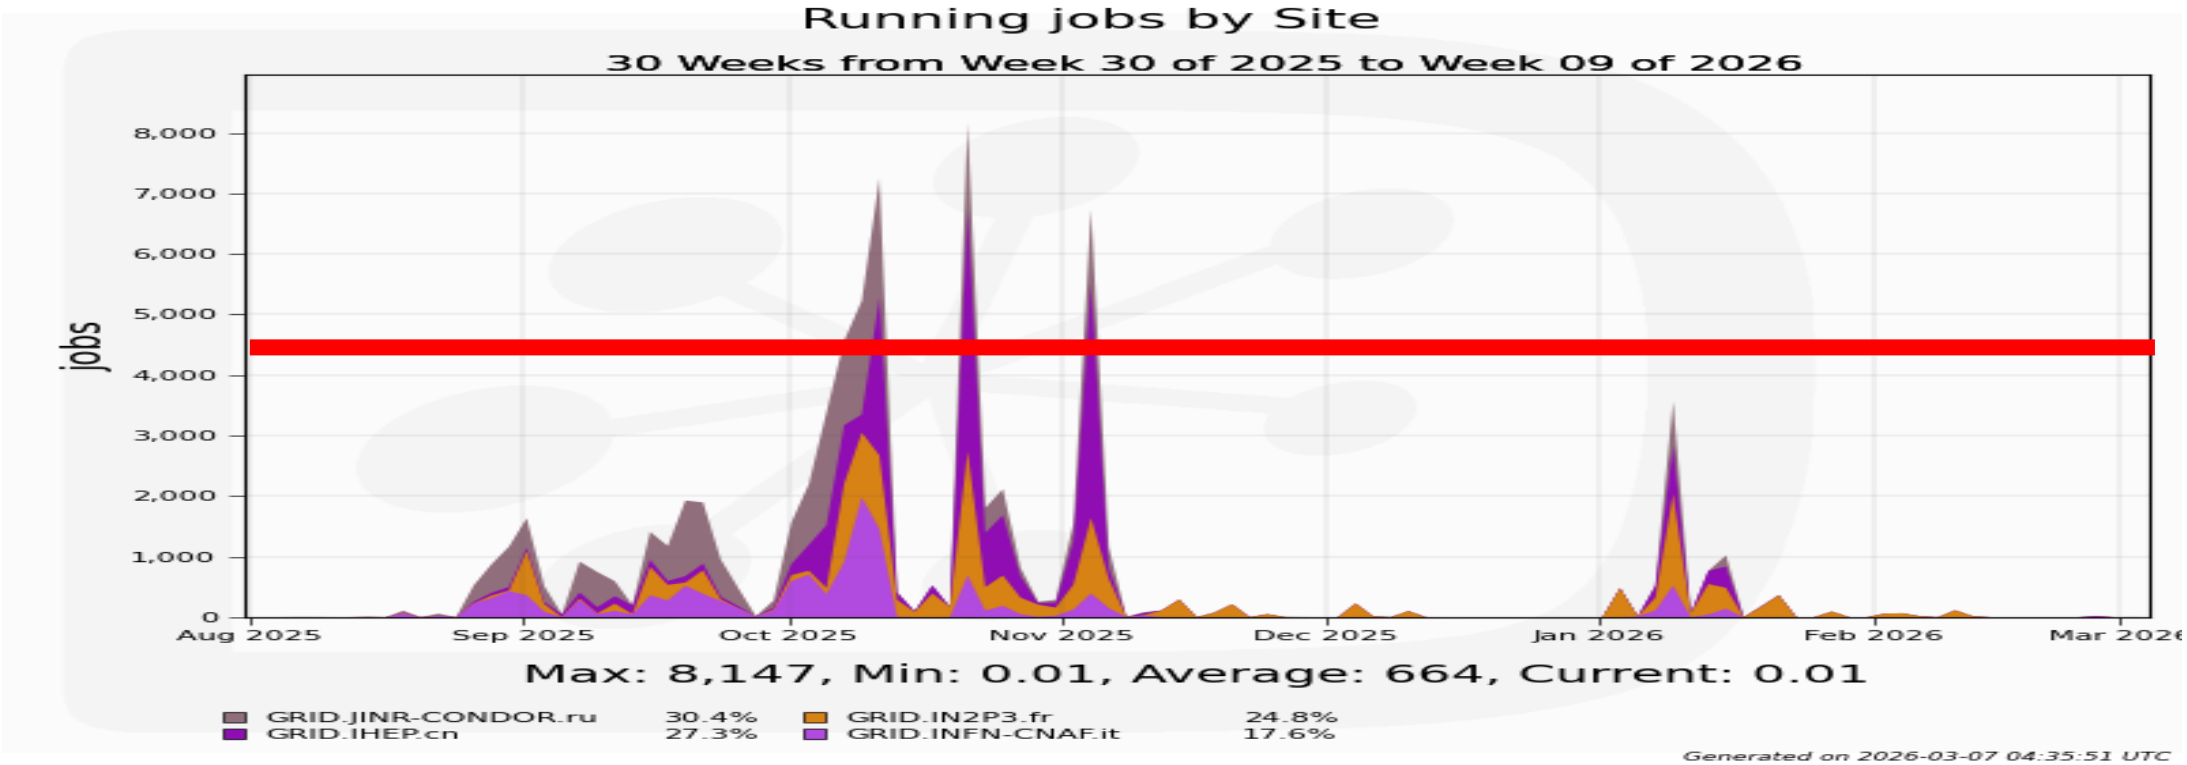

MC production

Simulation of millions of optical photons present challenges in both memory and CPU time

Detector Simulation:

- IBD: ~60 /day
- Radioactivity: MHz
- Atmospheric neutrino: hundreds per year
- Muon flux in LS: 4 Hz

High Energy:
Lots of light

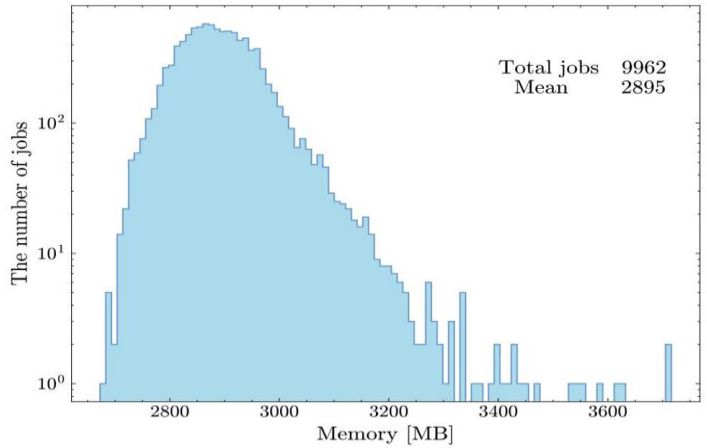

Memory consumption of Muon simulation

MC production with ProdSys

- Simulate signal + background to mimic real data
- 90% success rate, most failures from muon and atmospheric neutrino
- productions excluding high energy reached ~99%

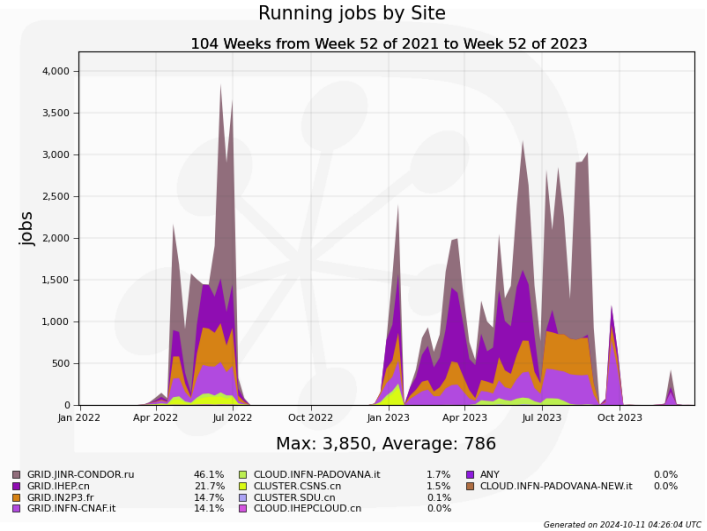

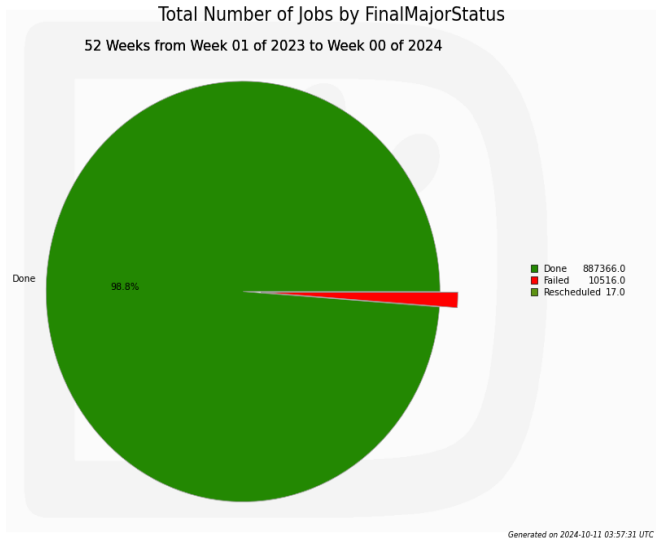

ProdSys: JUNO Data production system

- **Implemented as a data-driven pipeline system**
  - ⇨ create production jobs and submit JUNO production tasks in grid env
  - ⇨ chain and manage workflow and dataflow in tasks automatically
- **Data connection through DIRAC File Catalogue (DFC)**
  - ⇨ Frontend: accept requests, create workflow and dataflow
  - ⇨ DIRAC Transformation system (TS): Transform JUNO workflow and dataflow into a pipeline
  - ⇨ Jobs and file transfers submitted to DIRAC WMS+DMS

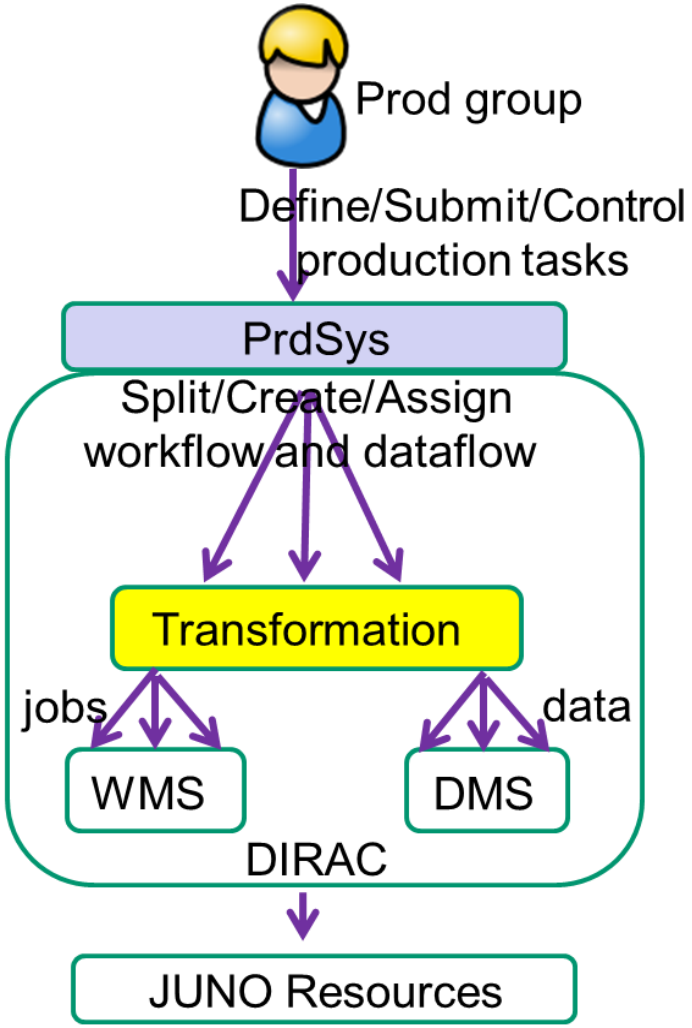

Parallelized MC production

- **To solve CPU and Memory problems**
  - Opticks@GPU speeds up simulation of high energy events
    - JUNO high energy events simulation is highly limited by optical photons, with 99% CPU time consumed by photon simulation,
    - GPU processing with Opticks very beneficial with overall simulation, also reduce CPU memory requirements
  - Multi-threaded DetSim implemented to decrease memory per core
    - Greatly reduce memory footprint below 2GB/core

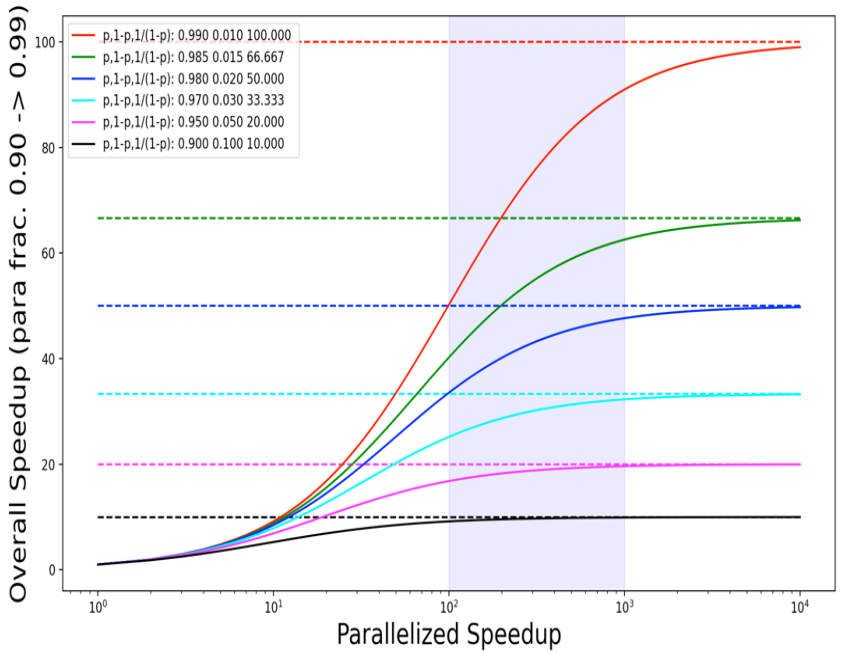

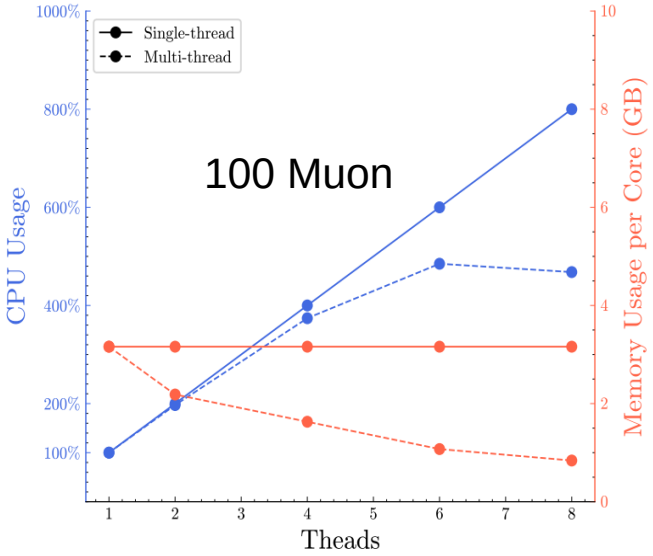

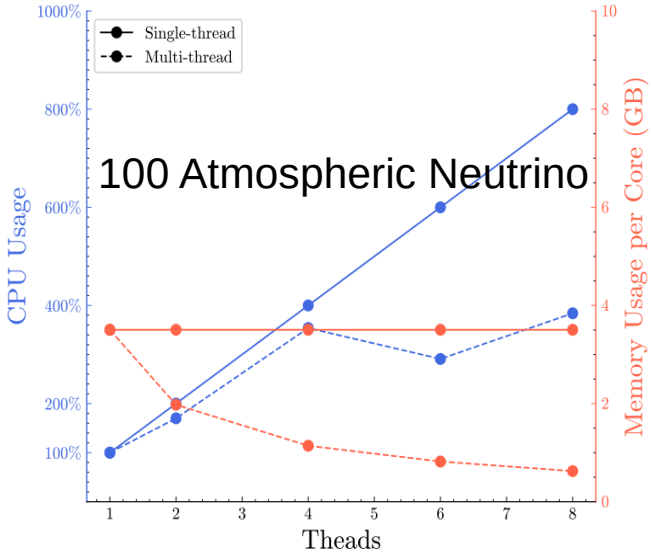

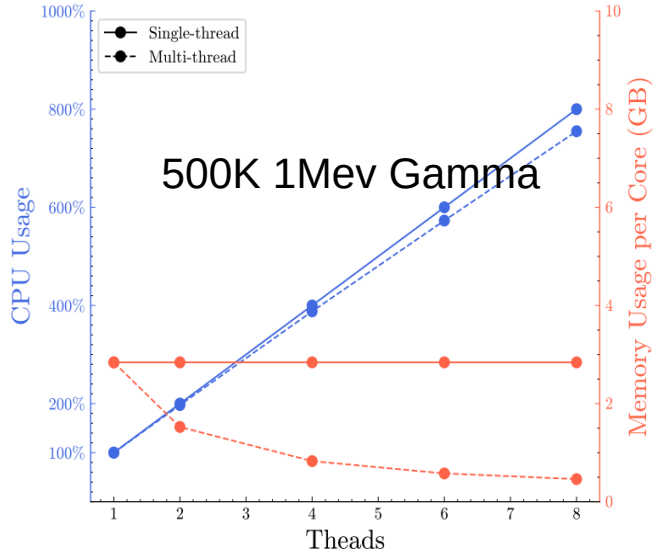

Parallelized MC production in large scale

- **Multi-threaded DetSim is ready to run in large production**
  - ⇨ DCI can support DetSim to run in 4-8 threaded with multi-core pilots
  - ⇨ With ProdSys multi-threaded jobs can be distributed in large scale
  - ⇨ Although monitoring and accounting still need improvements
- **Opticks simulation is ready to run in single GPU node**
  - ⇨ Optical-only simulation > 1000× faster vs. single-threaded Geant4
  - ⇨ JUNO muon simulation (with NVIDIA RTX Gen4) > 250× overall speedup
  - ⇨ But Opticks still remain at "isolated breakthrough" stage with low GPU utilization
- **How to make such CPU-GPU Hybrid mode into large scale production?**
  - ⇨ Still need efforts from software and computing

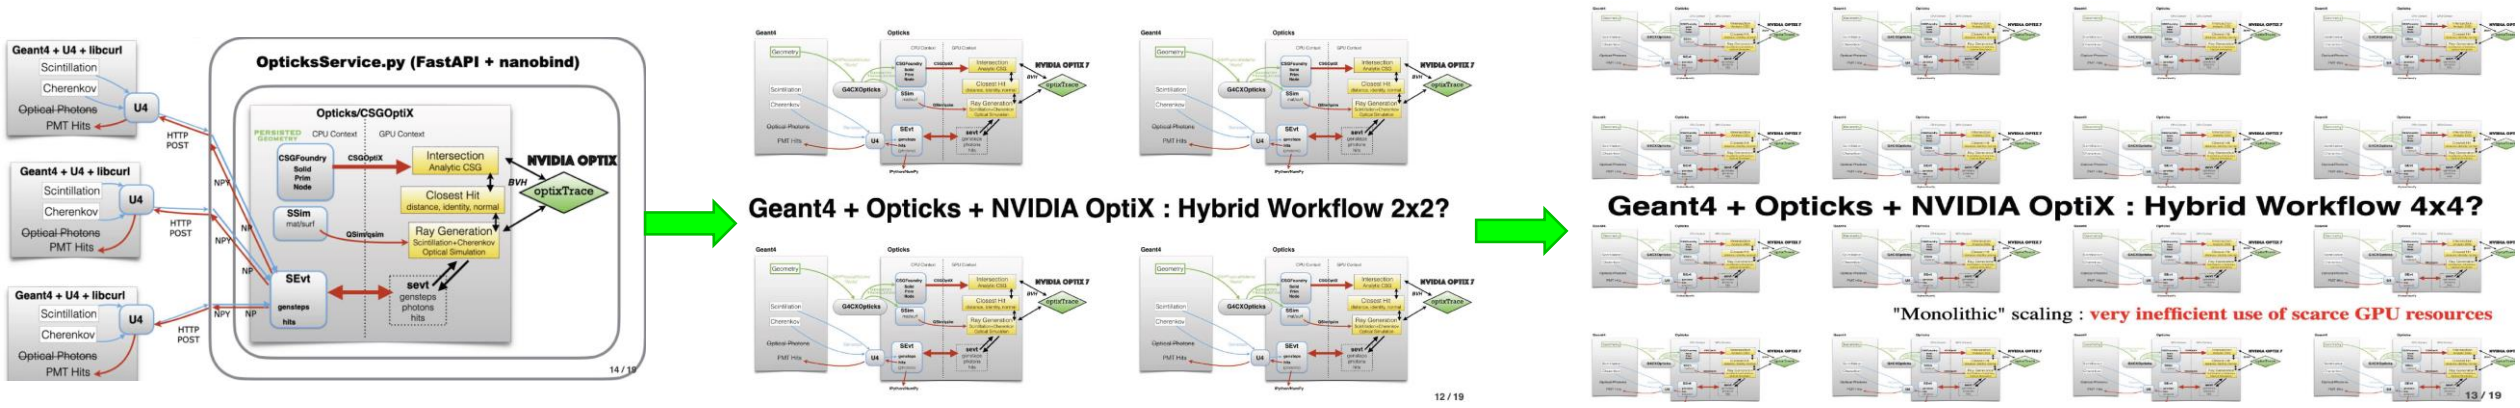

Summary

- JUNO data-taking began in August 2025
- Distributed computing infrastructure successfully supported early operations
  - Global transfers operated smoothly
  - Three reproduction campaigns completed
  - More reproduction in 2026
- Key bottlenecks identified and improvements still needed
  - Scaling and expansion ongoing
- Plan to support multi-threaded and Opticks simulation in large scale