Multifunctional Information and Computing Complex status and trends

Artem Petrosyan, MLIT JINR

16th Collaboration Meeting of the BM@N Experiment at NICA

VBLHEP JINR, Dubna, Russia

May 14, 2026

Laboratory of Information Technologies

LIT is not only ~350 employees, is designed to imply a powerful data-center, consists of a set of components, that, all together, form a Multifunctional Information and Computing Complex

Server halls located on the 2nd and 4th floors of the LIT building, there are also several technical floors filled with engineering equipment like power supply, climate control, water cooling, etc.

Power supply

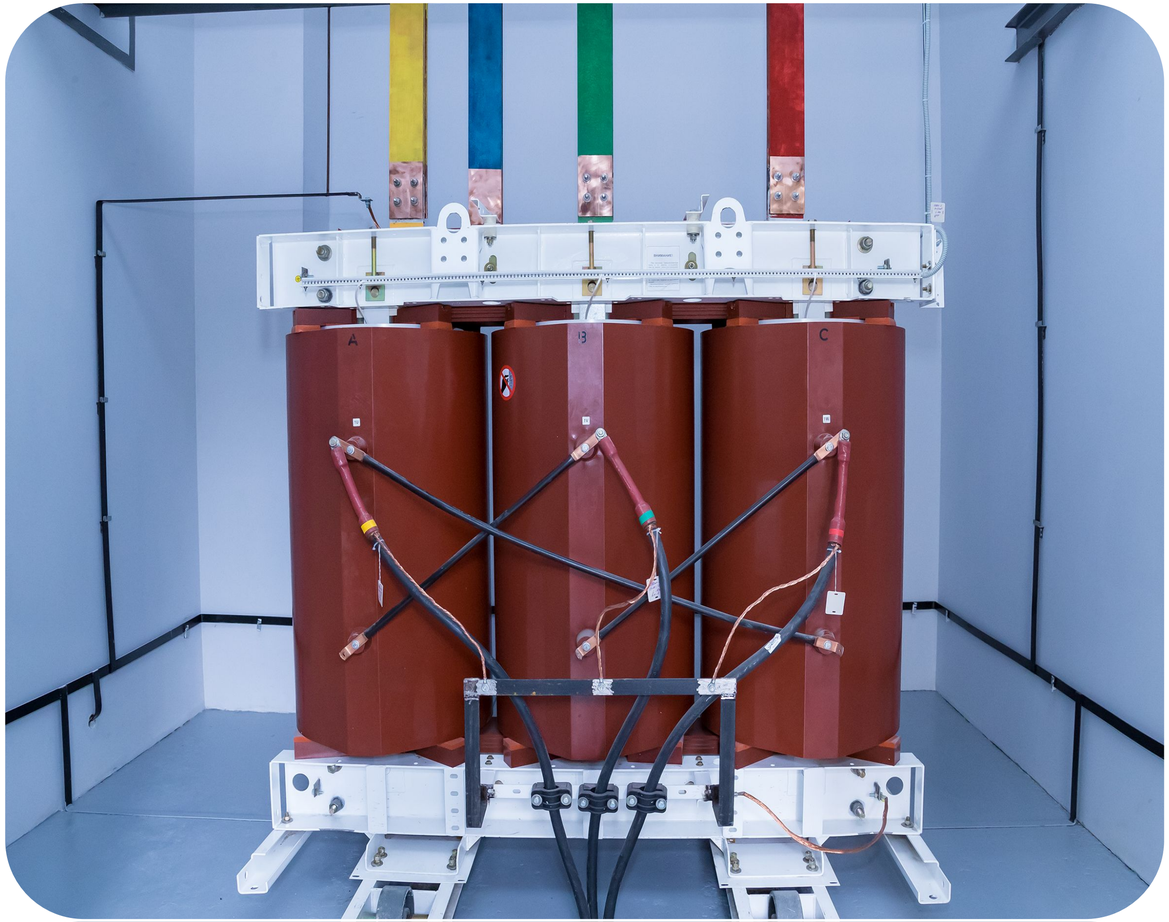

Dry transformers at the power supply input to the LIT building (2x 2500 kVA)

Main switchboard

Uninterrupted power supply system

APC Galaxy 7000

300 KVA

APC Symmetra PX

160 KVA

Riello

MHP 250

250 kVA

Riello

MHP160

160 kVA

Batteries

Uninterrupted power supply system

Containers with diesel generator units (DGU) to ensure uninterruptible power supply for the Multifunctional Information and Computing Complex

Main tank – 1 ton
Additional tank – 2 tons
Consumption – 335 l/h at 1 MW
Operating time with a full tank ~ 8 h

Cooling system

Dry cooling towers of the MICC cooling system

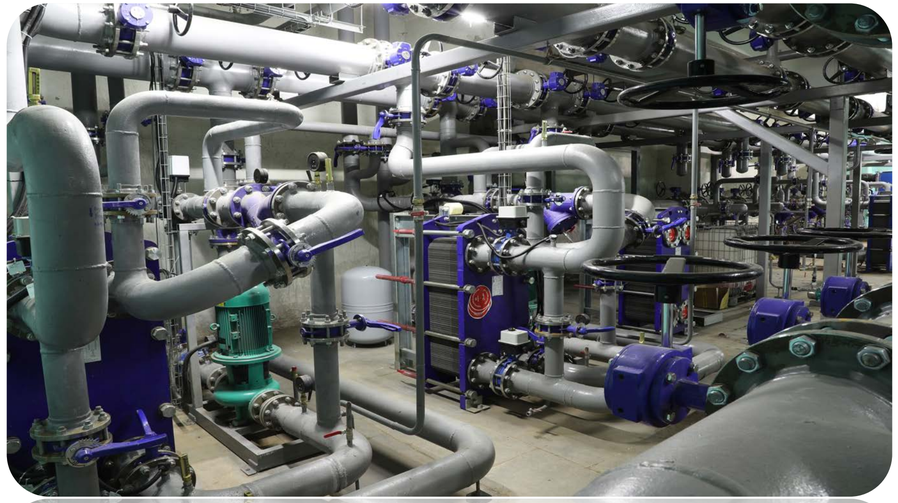

Cooling of the MICC hall (underground floor)

Refrigeration machines of the MICC hall

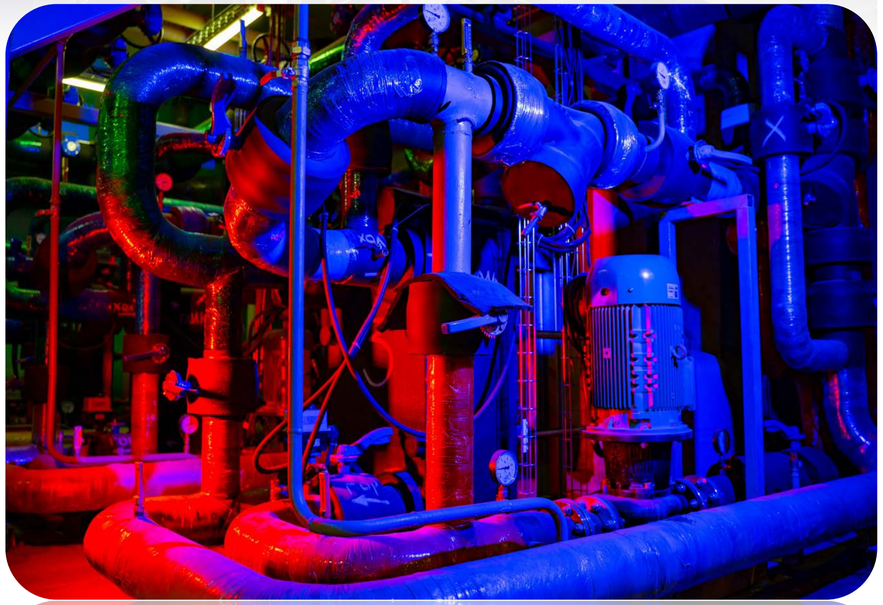

Dry Cooling Tower Pump and Heat Exchanger

Inter-row air conditioning

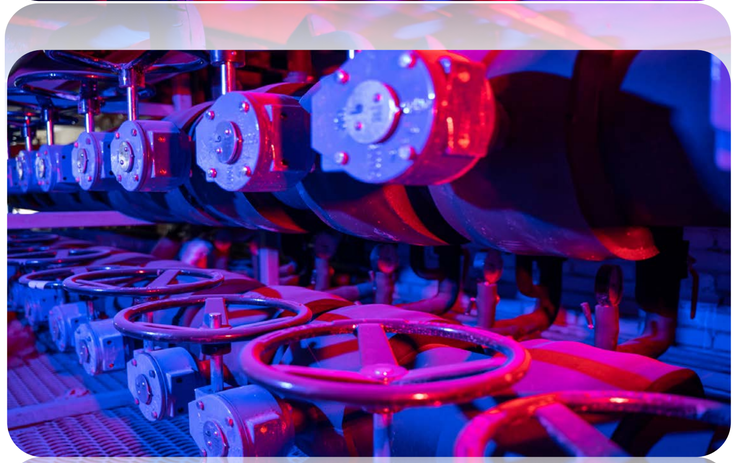

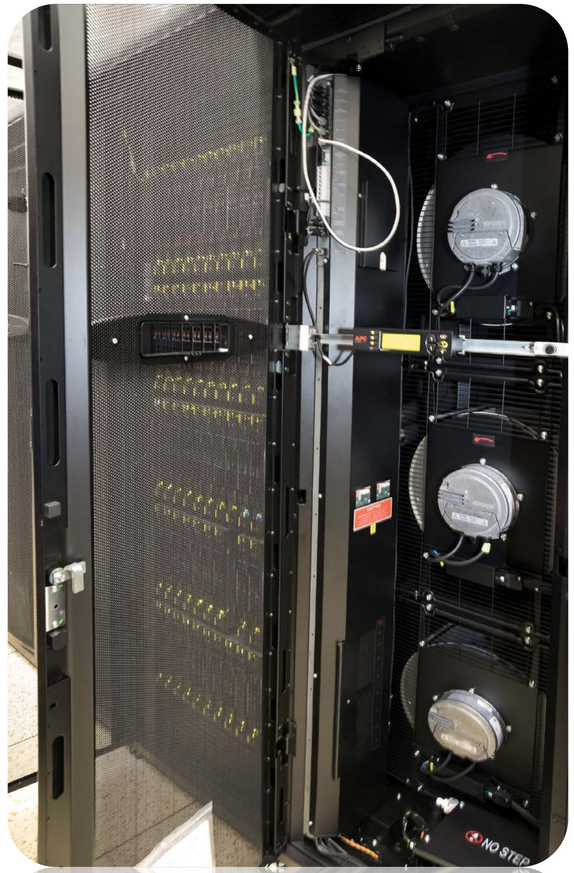

2nd floor

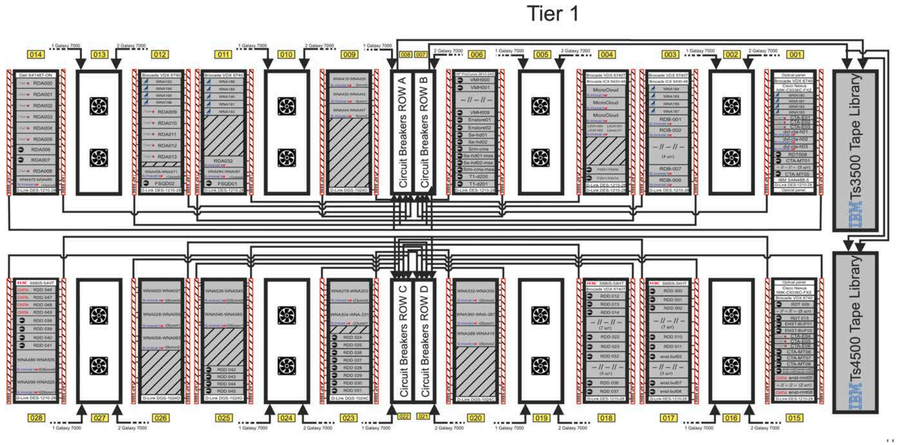

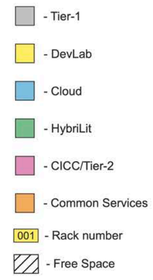

The computing capacities of the MICC are located in one computing hall with an area of 800 m^2 on the 2nd floor of the LIT building. At present, it consists of eight separate modules with computing equipment with a total permissible capacity of 2 MW.

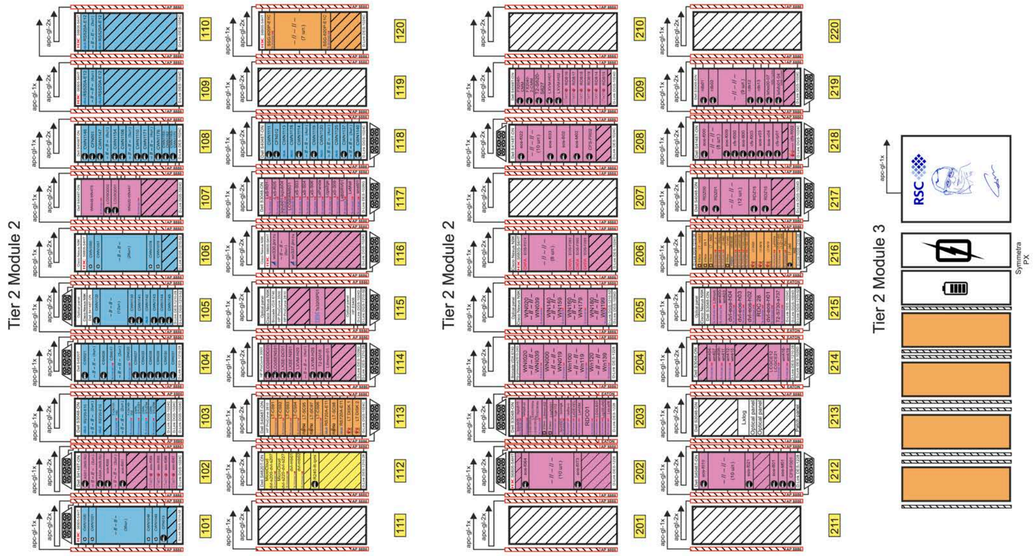

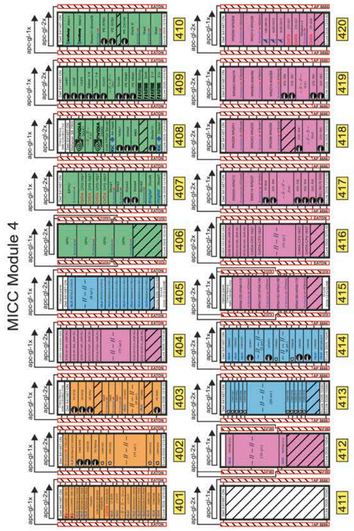

Along with computing and storage servers modules contain mission-critical servers of standard business computing: administrative systems, financial accounting and databases, etc., and network equipment.

4th floor

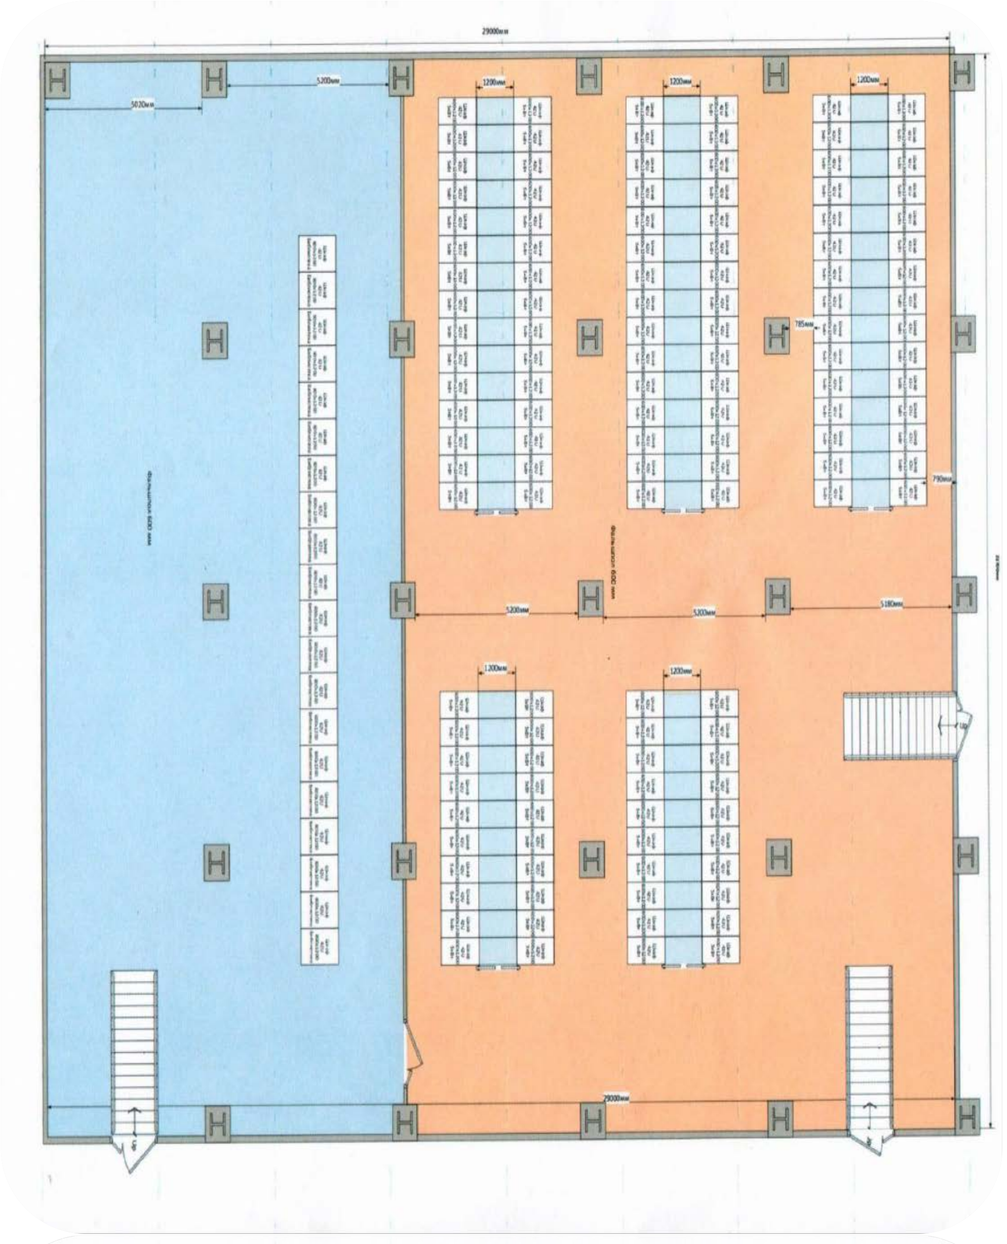

There is ongoing work on preparation of the 4th floor server room with power capacity at ~600 kW

Planned equipment:

- robotic tape library zone
- 130 server racks

2nd floor in numbers

- Module 1 and Module 2 with an area of 36.4 m2 each composed of 40 server racks (20 kW per rack)
- The Tier1 module with a floor-space of 29.33 m2 consisting of 16 server racks (35 kW per rack)
- Module 4 covering an area of 36.12 m2, which permits placing 20 server racks (35 kW per rack)
- The Govorun supercomputer with an area of 1.97 m2, 4 racks (100 kW per rack)
- The tape library assembled from IBM TS3500 and IBM TS4500 units placed on 13 m2 of floor-space, which provide a total tape capacity of 100 PB
- Two modules that host critical services of the JINR administrative system and the main network services for the MICC, the JINR local and wide area networks

Tiers in numbers

- Tier1
  - 482 compute nodes (23 360 cores) with a performance of 427 920.04 HEP-SPEC06
  - Data storage is provided by the 15 023.78 TB dCache system, the 100 PB robotic tape storage running the Enstore and CTA (CERN Tape Archive) software, the common EOS system with a capacity of 20 743.20 TB, and the MPD EOS and SPD EOS systems with a capacity of 7030.71 TB each
- Tier2
  - 485 compute nodes (10 356 cores) with a total performance of 166 788.4 HEP-SPEC06
  - Data storage is provided by the 4826.32 TB dCache system, 1527.77 TB ALICE@EOS, and EOS as a common distributed data storage system for all MICC components

T1 for CMS and NICA

Since the beginning of 2015, a full-scale WLCG Tier1 site for the CMS experiment has been operating at MLIT JINR.

The importance of developing, modernizing and expanding the computing performance and data storage systems of this center is dictated by the research program of the CMS experiment, in which JINR physicists take an active part within the RDMS CMS collaboration.

JINR Tier1 is regularly ranked on top among world Tier1 sites that process data from the CMS experiment at the LHC.

Since 2021, Tier1 center provides resources also for NICA experiments.

1500 CPU for BM@N

BM@N jobs this year

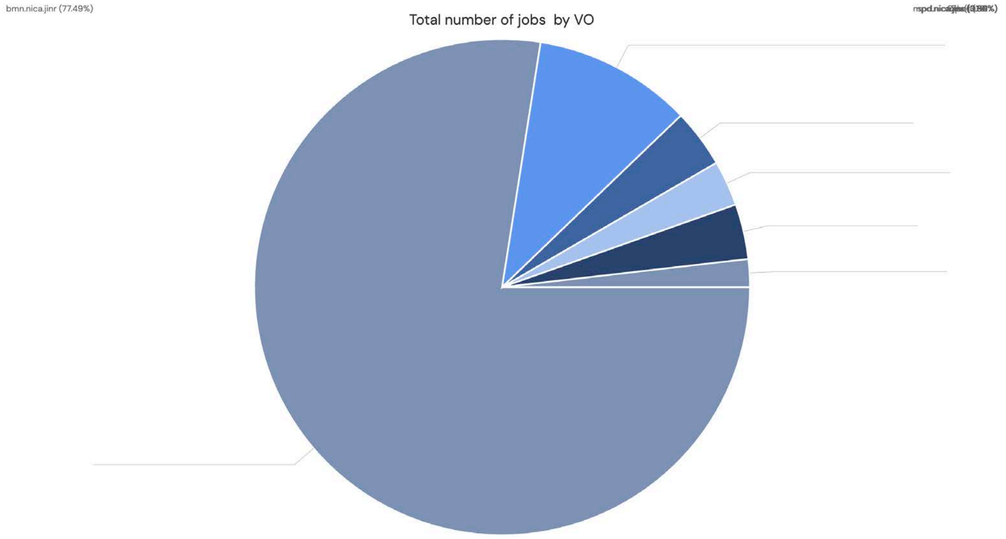

BM@N cpu time this year

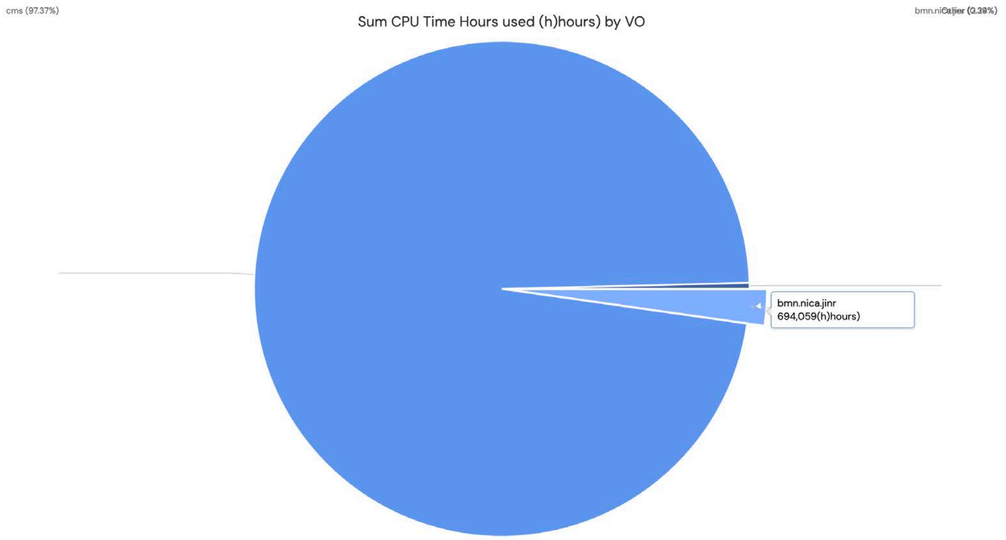

T2 for various experiments

BM@N jobs this year

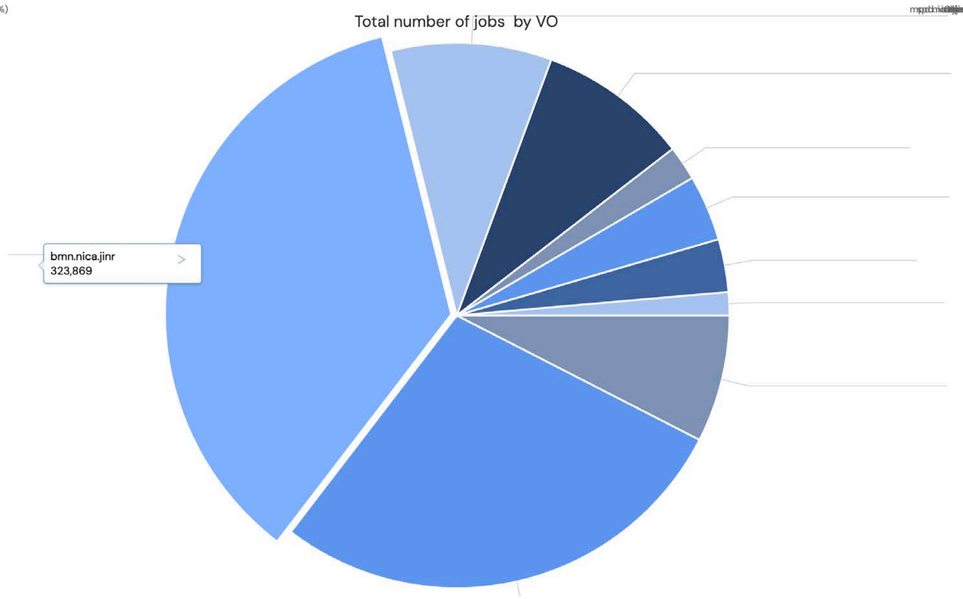

BM@N cpu time this year

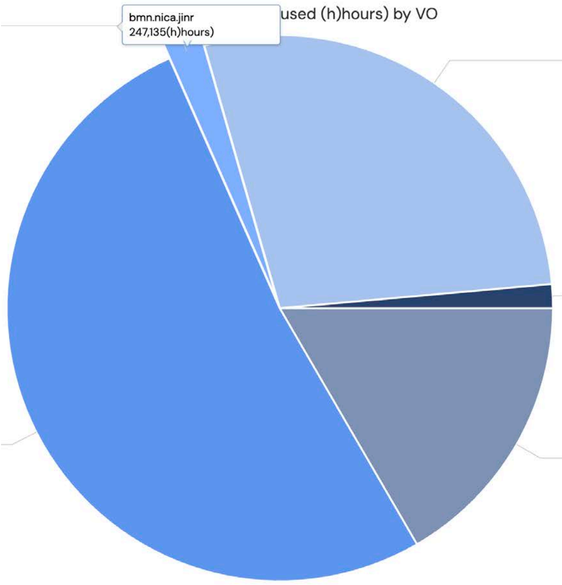

Tier2 at JINR is the oldest constantly evolving computing facility in LIT. Tier2 provides computing power and data storage and access systems for the majority of JINR users and user groups, as well as for users of virtual organizations (VOs) of the grid environment (LHC, NICA, etc.).

1000 CPU for BM@N

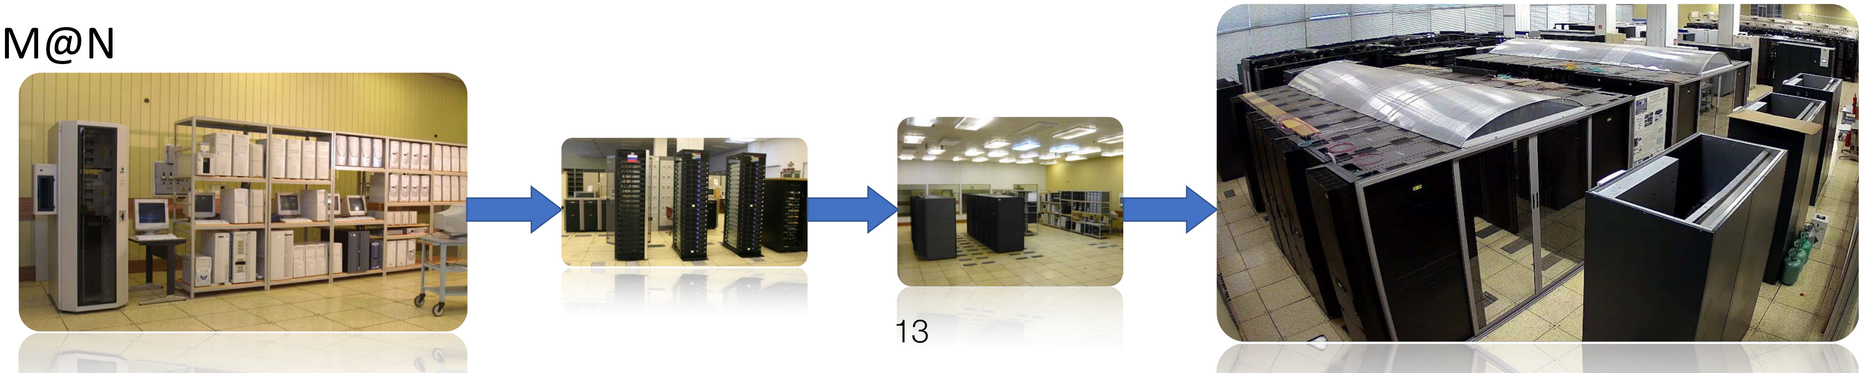

Support for the JINR Neutrino Program

MLIT contribution:
engineering infrastructure (electricity, UPS, cooling, network, racks, manpower)

DLNP contribution:
computing and storage resources (CPUs/GPUs&disks)

Computational resources for the **JINR neutrino program** are being used via the cloud infrastructure of the MICC. The **NOvA, Baikal-GVD and JUNO** experiments are the **major consumers** of the cloud infrastructure.

Recent and ongoing developments

- JINR CA
- Keycloak as application-friendly interface for JINR SSO
- Migration from CentOS7 to AlmaLinux9
- IAM replacement for VOMS
- Dedicated per-experiment EOS/dCache setups
- Central DB service

Transition from VOMS to IAM

VOMS

The Virtual Organization Membership Service enables Virtual Organization access control in distributed services.
It's at the core of the WLCG authorization stack and is used daily to authorize access to storage and computing resources used by thousands scientists worldwide.

⚠ **June 30th 2024: VOMS-Admin EOL**

- IAM (Identity and Access Management) comes to replace VOMS
- IAM can issue X.509 proxy, but also provide many other new features, mostly related to JWT
- https://indigo-iam.github.io/v/v1.14.0/docs/
- We've deployed instance of IAM for BM@N: iam-bmn.jinr.ru
- Igor Pelevanyuk is responsible from the BM@N side for testing and user support while transferring your account from VOMS to IAM
- At the moment we're testing middleware and user operations with the IAM
- Natalia Gromova is IAM VO administrator as she used to be in VOMS
- We expect to finish migration process during the summer this year

Storage

- We provide different types of storage facilities: EOS, dCache, Ceph, tapes
- Most of BM@N data at the moment are stored at main/common EOS
- BM@N quota in April 2026 is 4500.00TB, increased twice during this year from 1900.00TB in January 2026

From common EOS to dedicated setup

- Main EOS was initially built as storage resource for many projects
- After several years of operation we decided to split it in order to make a dedicated installations for large JINR experiments: BM@N, MPD, SPD, neutrino projects.
- At the moment SPD has its own EOS, SPD quota was removed from the main EOS and transferred to the BM@N
- MPD is in the middle of migration process, once it’s finished, quota will be removed as well
- Neutrino experiments moved to dCache instead of EOS, they’ve taken physical servers from the main EOS (1PB) and now adding them to their dCache setup
- Managing servers for the dedicated EOS for BM@N were already purchased and now are on the way, expected delivery date August this year

Migration from EOS to EOS

- BM@N has quite a lot of data at the main EOS
- To migrate to the new setup it’s possible that we’ll have to move these data somewhere else from EOS, because we simply do not have any other EOS of desired capacity
- I suggest to consider using tapes before new equipment is delivered to be prepared for migration to the dedicated EOS

MICC development strategy

- Creation of a dedicated disk storage for each large experiment as well as common user space at the main EOS and encourage to wider use tape libraries
- Organize remote storage centers at the members of collaborations
  - PNPI for SPD as working example
  - Ongoing negotiations with NRC KI
- Create close to physics facilities storage center at VBLHEP as part of the MICC
- Resources usage optimization
  - Managed campaigns of mass productions must be executed at large resources like Tier1 and Tier2 of MICC
  - Usage of Govorun HPC only for tasks where the full potential of the HPC can be reached and avoid using its storage system as long term storage
- Switching to components localized in the Russian Federation

Conclusions

- MICC is a large heterogeneous computing center providing a variety of storage and computing resources and services
- We work closely with both the international communities like WLCG and neutrino experiments as well as with local experiments and users
- While we provide a standard set of services for ordinary users, the set of services for collaborations and the level of their support is the subject of discussion and agreements with each collaboration
- We’re open to to discuss any suggestions for changing or expanding a set of supported services needed by the BM@N experiment and looking forward to find a solution that suits all parties

Thank you!

backup

Services for users

- Interactive farm
- Batch
- Heterogeneous cluster HybriLIT
- Personal quota at common EOS
- Cloud
- Git
- disk.jinr.ru, pm.jinr.ru
- etc.

Services for collaborations

- Interactive farm
- Batch, “Grid” resources, Govorun HPC
- Storage at common and dedicated EOS/CTA, dCache, Ceph, tapes…
- Cloud
- Git
- CernVM FS
- disk.jinr.ru, pm.jinr.ru
- FTS, VOMS/IAM, Dirac, Rucio, PanDA, whatever (discussed with each project individually)
- Central DB service
- etc.