NATURAL SCIENCE REVIEW **3** 200701 (2026)
nsr.jinr.int

Multifunctional Information and Computing Complex of JINR

A. I. Balandin[1], N. A. Balashov[1], O. Yu. Derenovskaya*[1], A. G. Dolbilov[1], A. P. Gavrish[1], A. O. Golunov[1], N. I. Gromova[1], A. V. Evlanov[1], I. A. Kashunin[1], V. V. Korenkov[1], N. A. Kutovskiy[1], V. V. Mitsyn[1], A. N. Moibenko[1], I. S. Pelevanyuk[1], D. V. Podgainy[1], O. I. Streltsova[1], S. V. Shmatov[1], T. A. Strizh†[1], V. V. Trofimov[1], A. S. Vorontsov[1], N. N. Voytishin[1], and M. I. Zuev[1]

[1] *Joint Institute for Nuclear Research, Dubna, Russia*

Abstract

The Multifunctional Information and Computing Complex (MICC) of the JINR Meshcheryakov Laboratory of Information Technologies (MLIT) is a key element of the JINR network and information and computing infrastructures. The MICC is regarded as JINR's unique basic facility and plays a decisive role in scientific research, which entails advanced computing power and storage systems. Its uniqueness is ensured by the consolidation of all state-of-the-art information technologies for data processing and storage, united by the network infrastructure with a bandwidth of up to 4×100 Gbps. It consists of distributed data processing and storage systems based on both grid and cloud technologies and the hyperconverged computing infrastructure with liquid cooling. Multifunctionality, high reliability, and availability in $24 \times 7 \times 365$ mode, scalability and high performance, information security and an advanced software environment are the main requirements that the MICC meets. The reliability and availability are ensured by the enhanced high-speed telecommunication system and the modern local network infrastructure, as well as by the reliable engineering infrastructure that provides guaranteed power supply and cooling for server hardware. This infrastructure is a staple for computing the experiments at the NICA accelerator complex. The BM@N, MPD, and SPD experiments intensively use all computational components and storage systems. Being part of the Worldwide LHC Computing Grid, the MICC serves as the Tier1 grid site for the CMS experiment at the LHC and as the Tier2 grid site that provides support for the experiments at the LHC and other world's large-scale experiments in high-energy physics. The integrated cloud environment of the JINR Member States focuses on supporting users and experiments in Russia, China, the USA, etc. (e.g., NICA, NOvA, Baikal-GVD, JUNO). The HybriLIT platform comprising the Govorun supercomputer provides capabilities for elaborating mathematical models and algorithms and performing resource-intensive computations, including on graphics accelerators that enable the development of the ecosystem for machine and deep learning tasks, Big Data analysis, and quantum computing on simulators.

Keywords: grid technologies, cloud technologies, Govorun supercomputer, distributed data storage, LHC, NICA, Baikal-GVD

DOI: 10.54546/NaturalSciRev.200701

---

*Corresponding author e-mail address: `odenisova@jinr.ru`
†Corresponding author e-mail address: `strizh@jinr.ru`

Received: 23 January 2026; Revised: 25 March 2026; Accepted: 25 March 2026; Published: 15 April 2026

Contents

1. Introduction

The research topic “Information Technologies & High-Performance Computing” is one of the priority vectors of the development of the Joint Institute for Nuclear Research (JINR) [1]. This is due to the increasingly considerable dependence of scientific research on the state and development of information technology. In turn, the development of scientific research at JINR defines the requirements for the computing infrastructure. The Multifunctional Information and Computing Complex (MICC) [2, 3] represents a key element of this infrastructure and plays a decisive role in scientific research, which entails advanced computing power and data storage systems. The MICC is placed at the JINR Meshcheryakov Laboratory of Information Technologies (MLIT) and considered as a large research infrastructure project [4] (Figure 1).

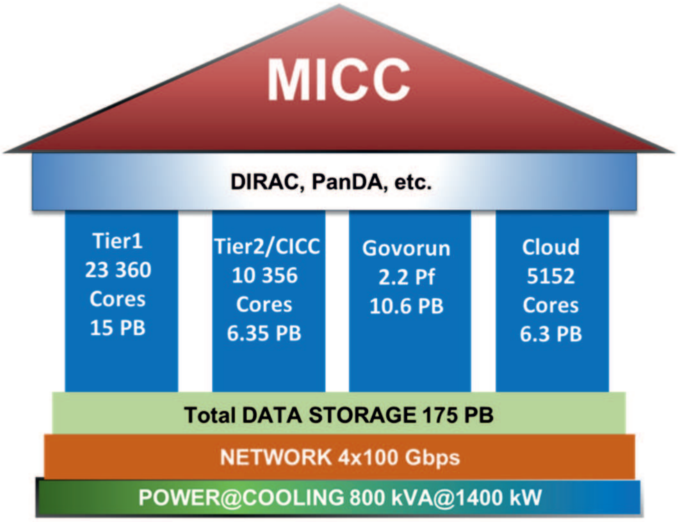

Figure 1. MICC figurative representation.

The concept of the MICC hardware and software complex (Figure 1) relies on the synergy of distributed high-throughput computing (HTC) and high-performance computing (HPC). HTC calculations are provided by facilities based on grid and cloud technologies. The MICC grid infrastructure consists of the Tier1 and Tier2 sites [5] within the WLCG (Worldwide LHC Computing Grid) project [6] created to support the experiments at the Large Hadron Collider (LHC). Now it crucially impacts on other large-scale experiments ongoing or planned at JINR at the NICA accelerator [7]. The JINR cloud infrastructure [8] is built on the

OpenNebula platform [9], providing users with an extensive service set and serving as the essentials for computing of the experiments within the JINR neutrino program. The HybriLIT heterogeneous platform integrates CPU- and GPU-enabled machines for HPC operations. The main component of this hardware and software platform is the Govorun supercomputer, which is the first hyperconverged and 100% "hot water" liquid-cooled supercomputer in the world [10]. HybriLIT also includes the advanced educational, testing and research ecosystem for machine learning and deep learning (ML/DL) tasks, as well as for the development of quantum algorithms by quantum simulators. Various types of solid state drives (SSDs), hard disk drives (HDDs), and tape library storages are used to store and manage data. The integral parts of the MICC are the data transmission system based on high-speed or high-bandwidth networking, the engineering infrastructure, and the monitoring system [11].

On 26 March 2015, a solemn presentation of a Tier1 grid site created to process data collected by the CMS experiment [12] at the LHC [13] took place during the meeting of the Committee of Plenipotentiaries of the Governments of the JINR Member States. Tier1 has opened up unprecedented opportunities for physicists from JINR, the JINR Member States, and the RDMS (Russian Dubna Member States) CMS collaboration [14] to fully participate in the processing and analysis of CMS experiment data. This moment was the starting point for the creation of a Multifunctional Information and Computing Complex on the basis of the JINR Central Information and Computing Complex, initially designed for data storage, processing, and analysis.

It is noteworthy that the JINR research program for the next decade is aimed at conducting ambitious and large-scale experiments at the Institute's basic facilities and within international cooperation. This program is related to the implementation of the NICA megaproject, the construction of new experimental facilities, the JINR neutrino program, participation in the CERN and other international large-scale experiments, as well as programs in condensed matter physics, nuclear physics, and life sciences. The implementation of the above projects entails decent and commensurate investments in systems that ensure the processing and storage of ever-increasing data volumes. The experience in recent years has demonstrated that progress in obtaining research results directly depends on the performance and efficiency of computing and storage resources provided. Therefore, the further development and enhancement of the MICC performance, as well as the provision of state-of-the-art IT solutions to the complex's users and the improvement of its operational efficiency, are the primary objectives of MLIT.

These tasks for the modernization of JINR computing infrastructures go hand in hand with the challenges facing the entire high-energy physics (HEP) research community [15–17] on the way to the upgrade of the LHC. Starting in 2030, the LHC will operate in its high-luminosity phase (High-Luminosity LHC, HL-LHC) [18], reaching peak luminosities up to $7.5 \times 10^3\ \mathrm{cm}^{-2} \cdot \mathrm{s}^{-1}$. This upgrade will increase the accumulated statistics by more than an order of magnitude (around 350–400 PB of raw data annually for ATLAS [19] and CMS [20]), with a projected integrated luminosity of $\mathcal{L}_{\mathrm{int}} \sim 3000\ \mathrm{fb}^{-1}$. Analogous tasks are also relevant for other large-scale experiments, such as JUNO [21] and Baikal-GVD [22] in the field of neutrino physics and other sciences, for example, the SKA project [23] in astronomy, which require a number of resources comparable to the LHC and enable using the successful experience of the WLCG project.

To meet these requirements, it is needed to evolve solutions on which, for almost two decades, within the WLCG project, the HEP computing infrastructure has demonstrated the unique capabilities of a distributed computing infrastructure for the LHC experiments in managing many hundreds of petabytes of data, scaling computing resources and data storage resources, and providing transparent access to this system for users. It is worth recalling the role of the global distributed computing infrastructure created within the WLCG project to process data obtained during the LHC experiments. The well-built WLCG computing infrastructure allows the LHC experiments' participants to collectively publish around three hundred scientific papers annually. It helps to produce new knowledge and generate new insights in the field of particle physics, which contributes to understanding the fundamental properties of matter. These efforts were recognized with several prestigious prizes. Thus, the international team of the ATLAS and CMS collaborations, including computer scientists from JINR, was awarded the 2013 High-Energy and Particle Physics Prize of the European Physical Society for "the discovery of a Higgs boson, as predicted by the Brout–Englert–Higgs mechanism".

This was particularly highlighted at the historic seminar on 4 July 2012, dedicated to the observation of a new particle similar to the Higgs boson at the CMS and ATLAS experimental facilities. CERN Director-General R. Heuer praised grid technologies in his talk, underlining their invaluable significance for global science. He identified three key components that had contributed to this triumphant result, namely, the CERN accelerator complex, the ATLAS and CMS experimental facilities, and, surely, the LHC grid infrastructure. It was the LHC grid infrastructure that made it possible to process and store enormous volumes of data coming from the collider experiments, which ultimately resulted in a scientific discovery of global importance. The JINR grid site also made a considerable contribution to this result, remaining the leading grid site in the RDIG (Russian Data Intensive Grid) consortium and being among the world's elite Tier2 grid sites for many years. The 2013 Nobel Prize in Physics was deservedly awarded to Professors Peter Higgs and François Englert "for the theoretical discovery of a mechanism that contributes to our understanding of the origin of mass of subatomic particles, and which recently was confirmed through the discovery of the predicted fundamental particle at CERN".

It would be useful to mention that in 2025, the Breakthrough Prize in Fundamental Physics was awarded to thousands of international researchers from the ATLAS, CMS, ALICE, and LHCb experimental collaborations at the LHC, in which researchers from JINR participate. It was given for "detailed measurements of Higgs boson properties confirming the symmetry-breaking mechanism of mass generation, the discovery of new strongly interacting particles, the study of rare processes and matter/antimatter asymmetry, and the exploration of nature at the shortest distances and most extreme conditions at CERN's Large Hadron Collider".

Currently, work to elaborate a strategy for the development of HEP computing is actively underway [24, 25]. The evolution strategy is based on three points: i) the use of the existing HEP computing infrastructure and its enhancement as a common computing system for HEP and other sciences, ii) the development of tools and services for creating an HEP data lake, and iii) the work in the field of creating common software and software methods. JINR has been actively involved in the WLCG project since its first day of functioning, being its full participant, therefore, the development and modernization of the MICC computing infrastructure is part of the development of this project.

It should be particularly noted that in 2018, the project of the Govorun supercomputer was initiated by scientists from MLIT and the JINR Bogoliubov Laboratory of Theoretical Physics (BLTP). It was aimed to drastically accelerate complex theoretical and experimental research, including the development of computing for the NICA megaproject. The implementation of this project on the most state-of-the-art computing architectures and IT solutions provided users with the opportunity not only to efficiently perform parallel computing, but also to develop ML/DL methods and algorithms for solving a wide range of tasks, including experimental data processing using a neural network approach. The widespread implementation of these approaches, including for experimental data processing in HEP [26, 27], neutrino experiments [28], astronomy [29], etc., is due to a multitude of factors. The main ones include the development of computing architectures, especially when using DL methods in training convolutional neural networks, as well as the development of libraries that implement a wide variety of algorithms and frameworks allowing for the rapid construction of various neural network models.

It is known that a modern computing complex should meet requirements such as multifunctionality, high performance, high reliability, fault tolerance and availability, information security, scalability, and have a developed software environment support system for various user groups. To fulfill these requirements in 24 × 7 × 365 mode, it is necessary to ensure the constant upgrade and enlargement of the complex's capabilities. The rapid development of information technology does not allow one to fully identify specific solutions that will define the MICC development for the coming years, however, the trends for this upgrade are clear enough and aimed at:

— modernization and development of the Tier1 and Tier2 sites, which provide support for experiments using the grid environment and collaborating with physics groups at JINR;

— modernization and development of the IT infrastructure of the NICA project;

— empowering the cloud component to enlarge the range of services provided to users and create an integrated cloud environment for experiments involving JINR (NICA, BESIII [30], NOvA [31], JUNO, etc.) and its Member States using containerization technology;

— extension of the performance and capacity of storage systems to satisfy the needs of international collaborations with JINR's participation;

— modernization and development of the computer cluster for non-grid users of the JINR Laboratories and its Member States, serving as user interfaces to grid infrastructures;

— increasing the performance of the HybriLIT heterogeneous platform and the Govorun supercomputer [32];

— development and enhancement of the JINR telecommunication and network infrastructure;

— stage-by-stage modernization of the MICC engineering infrastructure.

2. Network infrastructure

The development of information technology at JINR and the MICC project is directly related to the further maturing of the JINR network infrastructure. Proper networking ensures the creation and functionality of distributed data processing and storage systems based on multiple components–devices as a whole. It cannot be considered as a separate component.

The network infrastructure is an intricate complex of multifunctional network hardware and specialized software, which is the foundation for the JINR information and computing infrastructure that has been created and is constantly developing.

The JINR network infrastructure consists of four functional components: i) the external optical telecommunication data transmission channel, ii) the distributed multisite cluster network (cluster backbone), iii) the backbone of the JINR local area network (campus backbone), including the local area networks of the JINR subdivisions, and iv) the MICC local area network [33, 34].

The MICC and JINR network has direct connections to a number of scientific, educational, and public networks. To ensure connectivity between the MICC and other WLCG sites, dedicated external telecommunication channels (JINR–CERN duplicated links) are used. At the end of 2025, the telecommunication channels have the following data transfer speed parameters:

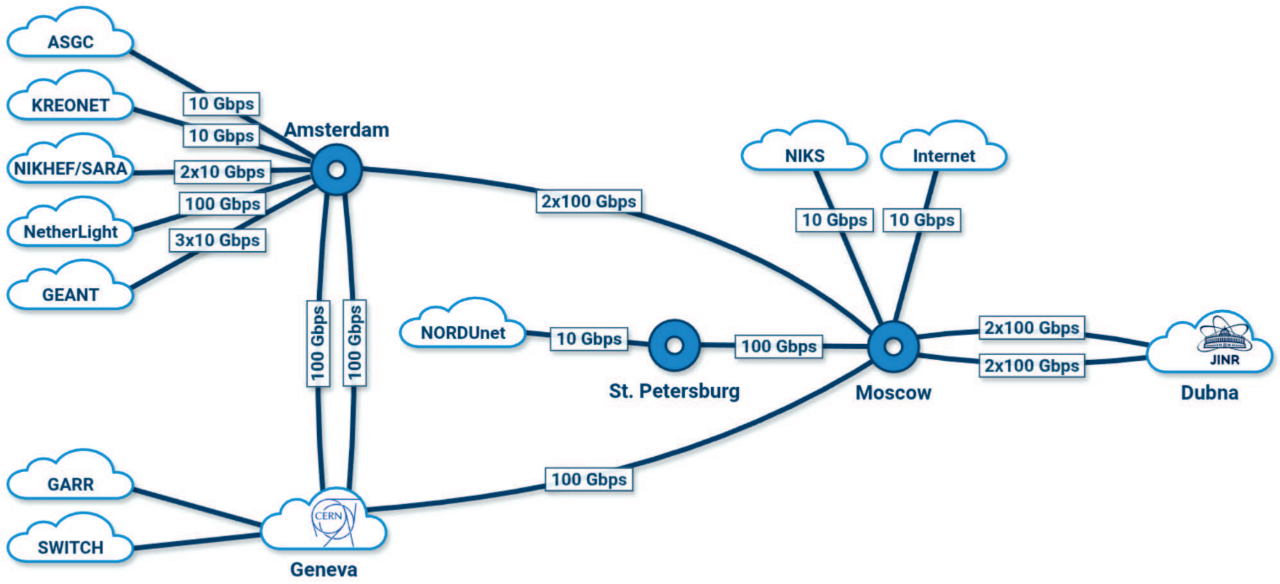

Figure 2. Scheme of JINR external telecommunication channels (as of the end of 2025).

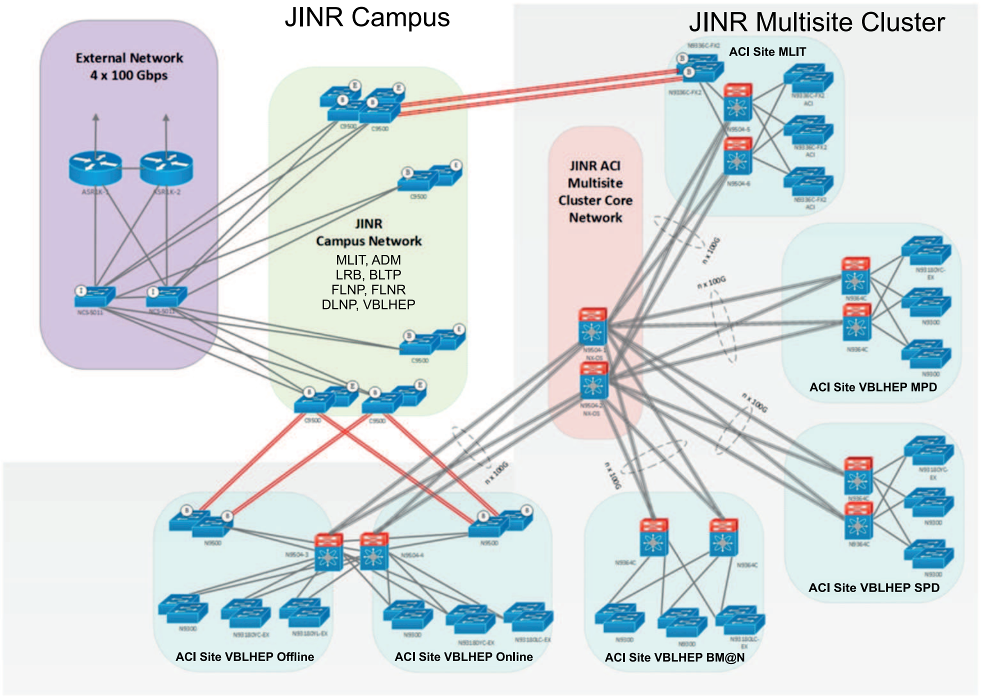

Figure 3. Scheme of the JINR campus and multisite cluster network.

2 × 100 Gbps with the LHCOPN network [35], 2 × 100 Gbps with the LHCONE network [36], 30 Gbps with the GEANT network [37], 10 Gbps with the National Research Computer Network of Russia (NIKS) [38], 100 Gbps with the networks of Moscow and St. Petersburg; 10 Gbps for public Internet (Figure 2).

The main optical transport medium of the campus backbone comprises three optical rings and provides data transmission with a speed of 2 × 100 Gbps. The cluster backbone links the geographically distant sites of JINR. It provides reliable data transfer with a speed of 4 × 100 Gbps from the computing equipment of the NICA complex to the MICC components for further processing, storage, and analysis (Figure 3).

In 2018, the Cisco company, a worldwide leader in information technology, awarded the 2018 Best Networking Project Prize to the JINR backbone development project, aimed to establish a high-speed local area network at JINR with a bandwidth of 100 Gbps, connecting all the JINR Laboratories and subdivisions, including the network infrastructure of the NICA megaproject.

3. Engineering infrastructure

The MICC computing facilities are placed in one computing hall of 900 m^2 of floor-space. There are eight separate IT equipment modules:

— Module 1 and Module 2 (Figure 4) with an area of 36.4 m^2 each composed of 40 server racks (20 kW per rack);

— the Tier1 module with a floor-space of 29.33 m^2 consisting of 16 server racks (35 kW per rack);

— Module 4 (Figure 5) covering an area of 36.12 m^2, which permits placing 20 server racks (35 kW per rack);

— the Govorun supercomputer (Figure 6) with an area of 1.97 m^2, 4 racks (100 kW per rack);

— the tape library assembled from IBM TS3500 and IBM TS4500 units placed on 13 m^2 of floor-space (Figure 7), 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.

The "cold corridor" scheme serves as the basis for cooling Modules 1, 2, 4, and Tier1. It is based on inter-row air conditioning, with Tier1 cooling powered by two separate chillers. A distinguishing feature of the Govorun supercomputer is the 100% liquid cooling of the CPU and storage components with its own dry cooling tower, duplicated by common dry cooling towers.

Thus, the current power consumption in the hall on the 2nd floor is 1 MW, and its capacity for placing equipment is almost exhausted. In 2025, work on reconstructing the machine hall on the 4th floor of the MLIT building started. This will allow one to place computing equipment

Figure 4. Modules 1 and 2 in the machine hall on the 2nd floor of the MLIT building.

Figure 5. Module 4 using the cooling of the cold corridor of the module with inter-row air conditioners.

Figure 6. The Govorun supercomputer in 2025.

and data storage systems for both the experiments at NICA and the neutrino program facilities (Baikal-GVD and JUNO) given the start of data taking. It is planned to place over 100 racks for equipment in the modernized MICC hall (600 kW) and create a zone for robotic tape libraries for long-term data storage.

The MICC engineering infrastructure is designed to ensure the reliable, uninterrupted, and fault-tolerant operation of the information and computing system and the network infrastructure. The holistic approach to building the MICC engineering infrastructure enabled one to work out algorithms for the operation of specialized equipment and the interaction of separate systems in both normal operation mode and emergencies. This provides an uninterrupted working capacity regardless of external factors.

Figure 7. Robotic tape libraries. IBM TS3500 and IBM TS4500 units.

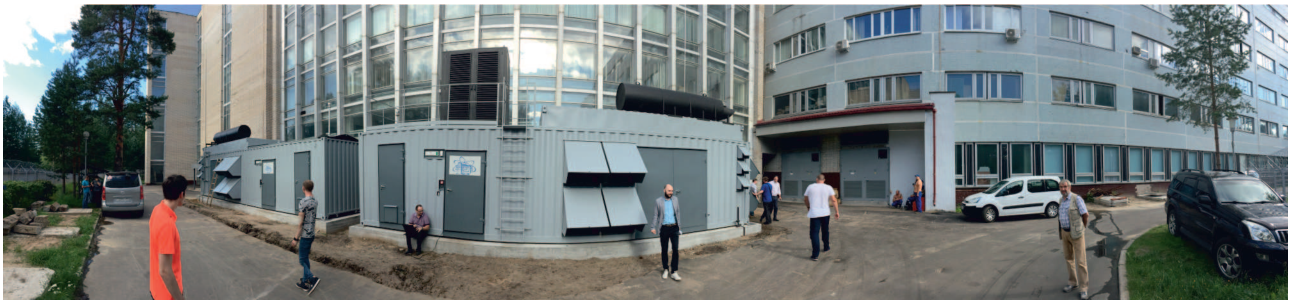

Figure 8. Diesel generator units near MLIT.

It should be emphasized that an inherent feature of the engineering infrastructure is scalability, which grants the power to the growth of computing facilities for the nearest three or four years.

In 2017–2019, the MLIT power supply system was modernized. 1000 kVA transformers were replaced by 2500 kVA ones, and the mode of simultaneous functionality was put into operation for uninterruptible power supplies (UPS) and diesel generator units (DGU) (Figure 8). This meets the most stringent first-class reliability standards for the MICC functioning in 24 × 7 × 365 mode.

The guaranteed and uninterruptible power supply system created during the first years of the MICC project implementation ensures:

— UPS to connected consumers;

— automatic diesel generator start;

— automatic switching of the load from the main external power supply to the diesel generator and then back again;

— issuing an alarm signal to the operator's console in the case of emergencies and DGU equipment failure.

The 2N redundancy system is employed to increase the fault tolerance of the power supply system.

The main component of the MICC climate control system is refrigeration equipment, that is, a complex of interconnected units utilizing various air and liquid cooling schemes. This combination and its coordinated work guarantee the preset temperature regime.

The cooling system also comprises the following components: the free cooling of the equipment with the cooled air of the machine hall; the cooling of the cold corridor of the module by inter-row air conditioners; the liquid cooling of computer elements. According to the type of heat removal, the MICC climate control system refers to a mixed type of execution that combines systems with the evaporation of a refrigerant and systems with an intermediate coolant.

The climate control system is based on a centralized cooling system with dry (for the winter period) and wet (for the summer period) cooling towers with a capacity of up to 1.8 MW and a possibility of extension up to 3 MW. The implementation of the MICC climate control system involves connecting refrigeration equipment of any type and operating principle, including energy-efficient refrigeration systems with free cooling, which enhance the efficiency of energy usage in the computing complex and reduce energy costs.

The further development and modernization of the cooling system rely on novel technological solutions applied in modern computing centers to create the required indoor microclimate and satisfy the needs for the development of all MICC components.

4. Grid infrastructure

For more than 20 years, the MICC grid site has been part of the WLCG global grid infrastructure created to provide distributed computing resources to process, store, and distribute data collected every year by the experiments at the Large Hadron Collider.

This work is performed in cooperation with the CMS, ATLAS, ALICE [39] collaborations and world's scientific centers, where Tier1 sites for the CMS experiment and Tier2 sites for other experiments are located. Apart from Tier1 JINR, there are a total of six Tier1 sites (Karlsruhe Institute of Technology in Germany, Port d'Informació Científica in Spain, IN2P3 in France, INFN-CNAF in Italy, FNAL in the USA, Rutherford Appleton Laboratory in the UK) and around 160 Tier2 sites in more than 40 countries.

Since the beginning of the 2000s, MLIT JINR, jointly with world's centers, has been working on both the creation and the expansion and modernization of computing clusters built on distributed grid technologies to process and store data from the experiments at the LHC. Tier2 JINR created at MLIT has ensured data processing and analysis for all experiments at the LHC, as well as the operation of virtual organizations of other international experiments, in which JINR specialists participate. Today, Tier2 also provides data processing for all experiments at NICA and is the most productive in the RDIG consortium.

Since the beginning of 2015, a full-scale WLCG Tier1 site for the CMS experiment has been operating at MLIT JINR [40]. 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. There exist only seven such centers for CMS worldwide, and the JINR site is consistently ranked first in terms of performance, demonstrating almost 100% availability and reliability (Figure 9).

In 2015, the Tier1 Site to Process Data from the Large Hadron Collider at CERN project won the Project of the Year competition, organized by the Association of Directors of Russian IT Companies, GlobalCIO, in the Development and Modernization of Infrastructure category.

It is important to note that the first simulation data for the MPD experiment at NICA was produced with the resources of the JINR Tier1 and Tier2 grid sites in August 2019. This was made possible through the integration of the MICC resources using the DIRAC platform [41]. Thus, the JINR grid sites began to be regularly used to model data for the experiments of the NICA complex.

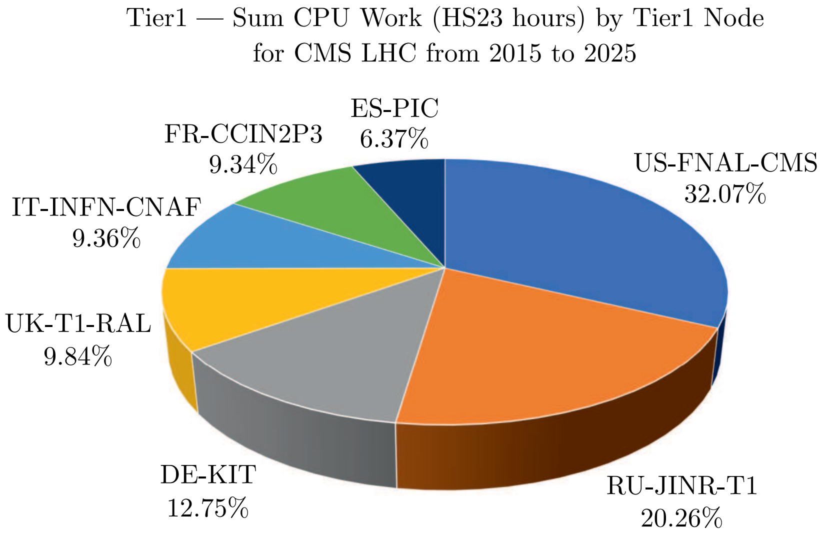

Figure 9. Distribution of the normalized CPU time used to process data from the CMS experiment at the LHC by global Tier1 sites for 2015–2025. JINR's share is 20.26% of the total.

The infrastructure and services of the Tier1 (JINR-T1) and Tier2 (JINR-LCG2) sites ensure the operation of:

— computing service for experimental data processing and Monte Carlo (MC) data production;

— data storage service for experimental and simulation data storage and handling;

— data transfer service for transmitting data between various Tier1 sites, as well as between Tier2 and Tier1 sites;

— service for management of data transfer and data handling;

— service of access to user home directories;

— service of access to user software versions;

— grid support service;

— distributed computing control system;

— information service (monitoring, information sites).

The computing services are ensured by a set of special basic services. These contain the Slurm Workload Manager [42], which is a free open source task scheduler for Linux and Unix-like systems used by many world's supercomputers and computing clusters. To organize computing in the grid environment, the Advanced Resource Connector (ARC) [43], middleware for grid computing, is used. It provides a common interface for transferring computational tasks to various distributed computing systems and thus can embrace grid infrastructures of different size and complexity. The set of services and utilities that provide the interface is called the ARC Computing Element (ARC-CE), open source software released under the Apache 2.0 license. Like Slurm, ARC-CE is used by most grid sites in the world.

The main functions of the Tier1 site include:

— receiving, archiving and custodially storing a significant share of CMS experiment data from Tier0 CERN;

— consistent and continuous data processing, reconstructed data (RECO) and analysis object data (AOD) extraction;

— data reprocessing using new software or new calibration and alignment constants, skimming, calibration;

— MC data production, including simulation with physics event generators, simulation of detector response, and digitization processing steps;

— storing a complete copy of RECO and AOD;

— sharing data with the other Tier1 sites for their duplicate storage (replication);

— receiving data from numerous Tier2/3 sites for their duplicate storage (replication);

— providing data to numerous Tier2/3 sites for further physics analysis;

— data processing and analysis, MC data production and analysis for the NICA experiments (MPD, BM@N, SPD) [44–46].

The functions of the Tier2 site include:

— providing capacities for user analysis, calibration studies, and MC data production;

— extracting the reduced versions of data tiers (e.g., MiniAOD and NanoAOD for the CMS experiment);

— storing the reduced versions of data tiers;

— MC data generation;

— sharing originally produced MC data with the Tier1 sites for their duplicate storage (replication);

— sharing all originally produced data with other Tier2 sites for further physics analysis.

Currently, Tier1 embraces 482 compute nodes (23 360 cores) with a performance of 427 920.04 HEP-SPEC06. The launch of tasks for CMS experiment data processing is carried out by 16 core pilots, and all computing resources are available to them. 4000 cores are currently available for data processing and simulation tasks for the experiments at NICA, and their number can be increased upon the experiments' request. Data storage is provided by the 15 023.78 TB dCache system [47], the 100 PB robotic tape storage running the Enstore [48] and CTA (CERN Tape Archive) [49] software, the common EOS system [50] with a capacity of 20 743.20 TB, and the MPD EOS and SPD EOS systems with a capacity of 7030.71 TB each.

The Tier2 site embraces 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. Tier2 JINR does not pledge to support tape resources.

The Rucio software platform [51] with customizable policies is used to manage data and access large amounts of data. Data is distributed across globally distributed and heterogeneous data centers, combining different storage and networking technologies into a unified federated structure. Rucio offers advanced features such as distributed data recovery or adaptive replication and is highly scalable, modular, and expandable.

5. Cloud infrastructure

The most universal component of the MICC is the cloud infrastructure, which is built on the fundamental cloud resource delivery model (Infrastructure-as-a-Service, IaaS). This type of service provides end users with computing resources in the form of virtual machines (VMs), which may serve as the final usable product of the cloud infrastructure or as a platform for building complex, multicomponent systems consisting of numerous interconnected VMs. The ability to dynamically allocate only the required volume of resources (CPU cores, memory and disk space) to VMs, as well as the ability to create and delete VMs on demand, ensures flexible resource redistribution among various users and systems in the cloud according to their current needs.

The JINR cloud infrastructure is used to solve a wide range of tasks, including hosting various web services, testing new versions of different systems and debugging developed software

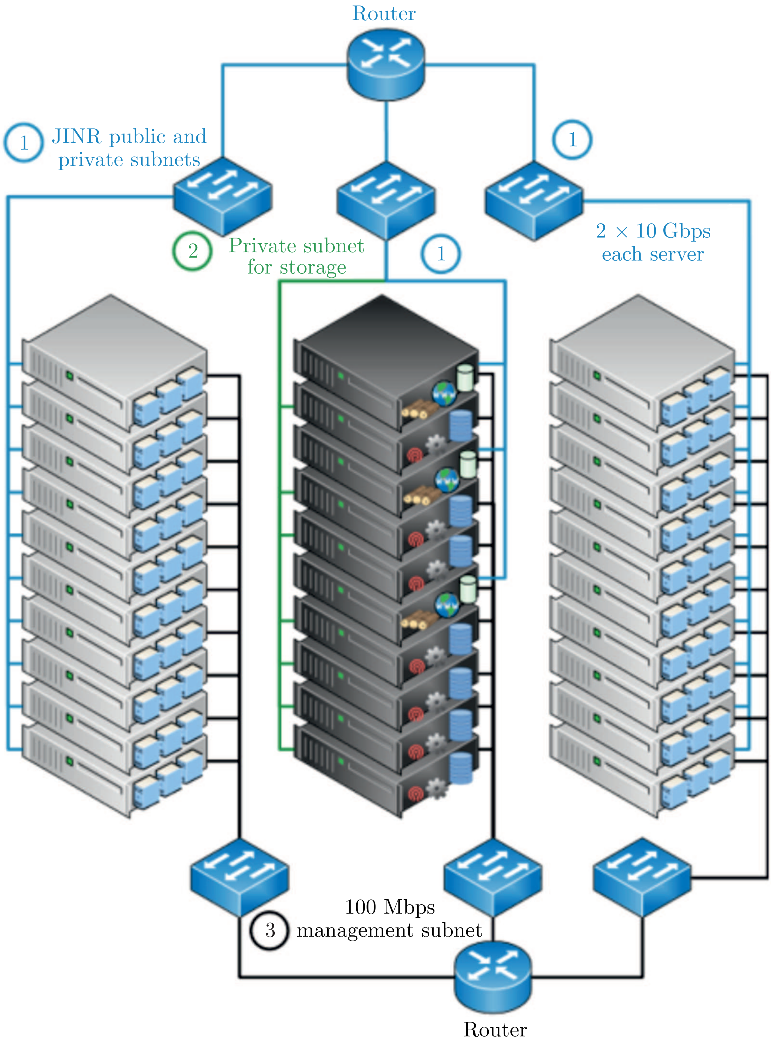

Figure 10. JINR cloud infrastructure architecture.

tools for various operating systems. It also hosts numerous information and computing services, such as a batch cluster on top of the HTCondor system [52] for distributed computing, the JupyterHub interactive programming service [53], the GitLab service for managing software development projects, and others.

A defining characteristic of the cloud infrastructure is the separation of management levels: users act as full system administrators of their VMs, while MICC operators take full responsibility for the entire underlying hardware and the virtualization environment. The operators ensure the reliable functioning of VMs by carrying out a full range of technical maintenance work on servers and storage systems (including failure monitoring, node replacement, firmware updates, etc.), as well as by participating in network diagnostics and troubleshooting along with the networking service. High availability is achieved through live migrations of running VMs between the nodes, which allows planned server and equipment maintenance to be performed without service interruption. Furthermore, VM images and disks are located on a shared network storage, i.e., in the case of a compute node failure, the virtual machine's operation is promptly resumed on another server, preserving the state of its disks.

Currently, the cloud service is built on the OpenNebula platform [12] using the KVM (Kernel-based Virtual Machine) hypervisor (the general architecture diagram is shown in Figure 10). The infrastructure includes 180 servers, providing approximately 5700 CPU cores and 69 TB of memory (RAM) for more than 800 VMs deployed within it. To ensure network connectivity, each compute node is connected by a 10 Gbps channel.

The cloud infrastructure comprises three independent storage systems based on the distributed fault-tolerant Ceph storage [54]. The network architecture of the storage systems is designed considering increased bandwidth requirements: SSD cluster nodes are connected via 100 Gbps and HDD cluster servers use two 10 Gbps (2×10 Gbps) links. These storage systems support various types of services:

- SSD-based storage with a volume of ~ 615 TB. It is used to provide block devices connected as VM disks, requiring high performance for input/output (I/O) operations.
- HDD-based storage with a volume of ~ 2.2 PB. It provides disks for VMs with lower I/O requirements, functions as an object storage service with S3 protocol access for integration with external systems (e.g., GitLab Docker container registry) and is used for storing data backups and VM images.
- HDD-based storage with a volume of ~ 3.2 PB. It is dedicated to storing data from neutrino experiments via network file system access.

To ensure a high level of data preservation and fault tolerance, triple redundancy is applied to all data in all three storage systems.

The management and configuration of all nodes of the cloud infrastructure, including the storage systems, is implemented in accordance with the Infrastructure-as-Code (IaC) paradigm. Configuration files are stored in a version control system based on Git, operating system deployment is carried out using the Foreman service, and final configuration tracking and maintenance are handled via the Puppet system. This automated and declarative approach ensures high efficiency in managing the cloud infrastructure, guarantees configuration consistency for homogenous nodes, and minimizes the likelihood of human error.

The inventory of the JINR cloud hardware is maintained in a system based on the iTop platform [55]. Cloud user support is currently also provided through this service via a ticketing system.

A key distinguishing feature of the cloud infrastructure is the ability to manage virtual resources using an application programming interface (API). Specifically, this programmatic interface enabled the integration of cloud infrastructures from various scientific and educational organizations of the JINR Member States into a unified distributed information and computing environment (DICE) [56] based on the DIRAC system.

The cloud infrastructure resources are utilized in various scientific projects, for Institute-wide services and to support the needs of individual subdivisions. In addition to general-purpose use, the cloud infrastructure also serves as the main environment for JINR neutrino program projects through a dedicated set of cloud resources and services referred to as the Neutrino Computing Platform (NCP). This platform comprises computing and storage resources together with the associated services deployed on them, such as the batch system and grid interfaces, tailored to the specific requirements of neutrino experiments.

The major scientific experiments using the MICC cloud infrastructure include neutrino experiments such as Baikal-GVD, JUNO, NOvA, Borexino, and TAIGA, which are supported via the NCP, as well as the BM@N and SPD experiments within the NICA project (Figure 11).

The cloud infrastructure's contribution is particularly significant for the JUNO collaboration, where the NCP operates as a Tier1 site within the experiment's distributed computing

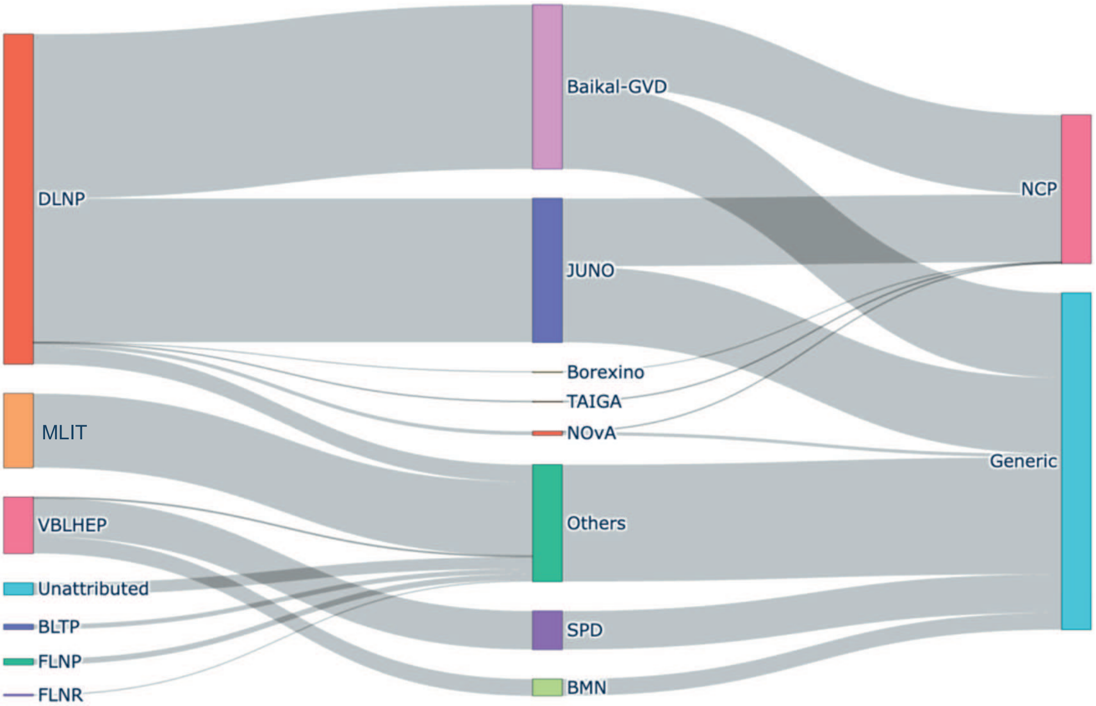

Figure 11. Distribution of the computing resources of the cloud infrastructure from the leading Laboratories (left) through their experiments (center) to available resource pools (right) in 2025. The upper pool refers to the Neutrino Computing Platform, while the lower pool represents the shared generic resources available to all the JINR Laboratories and experiments. The widths of the connections indicate the relative share of resource usage for each Laboratory–experiment–resource-pool combination.

infrastructure. It provides 2000 CPU cores in dedicated mode with the capability to scale up to approximately 3700 cores in opportunistic mode. In addition, storage support for the JUNO experiment is provided through a dedicated dCache instance deployed within the MICC infrastructure. For the Baikal-GVD experiment, the infrastructure is one of the key elements of the data processing chain, providing a separate dedicated cluster of 484 cores specifically for data processing (operating independently of the batch system), while additional 3700 cores are available opportunistically for batch processing.

6. HybriLIT heterogeneous platform

The HybriLIT heterogeneous computing platform is part of the MICC. Its history began in 2013 with the decision to expand the JINR MICC by adding a heterogeneous computing cluster. This decision followed contemporary trends in global computing technology, which only intensified due to the active development and implementation of artificial intelligence methods. In 2014, the HybriLIT computing cluster was commissioned. It featured the most advanced computing architectures of the time, namely, first-generation Intel Xeon Phi MIC multicore coprocessors and NVIDIA K20X and K40 graphics accelerators. This allowed for the integration of two key accelerator development trends, enabling cluster users to utilize the main HPC tools available on the market.

Since 2018, the platform's main computing heart has been the Govorun supercomputer, designed to perform resource-intensive and massively parallel computing. The creation of the

Govorun supercomputer at JINR is an essential technological achievement being of great importance for the implementation of the JINR research program and international cooperation.

The Govorun supercomputer was created on top of the experience gained during the operation of the HybriLIT heterogeneous cluster. HybriLIT has shown its relevance in solving tasks of quantum chromodynamics (QCD) on lattices, radiation biology, applied research, etc. The continuous growth in the number of users and the expansion of the range of tasks to be solved entailed not only a significant increase in the computing capabilities of the cluster, but the elaboration and implementation of novel technologies, which resulted in the creation of a new computing system, the Govorun supercomputer. The Govorun supercomputer is built as a high-performance, scalable liquid-cooled system with a hyperconverged and software-defined architecture. The current configuration of the Govorun supercomputer involves computing modules containing GPU and CPU components, as well as a hierarchical data processing and storage system.

With the launch of the Govorun supercomputer, the HybriLIT cluster continued to operate successfully, but its role changed. Thanks to the developed and implemented software and information environment, the cluster and supercomputer resources were combined into a heterogeneous platform named HybriLIT. Using the resources of the cluster, which became known as the education and testing polygon, platform users were able to utilize available application packages, develop their own applications, and perform calculations using various computing architectures (central processors and graphics accelerators), without the need to migrate data or recompile programs. At present, the HybriLIT heterogeneous computing platform is a multicomponent system consisting of the Govorun supercomputer, the education and testing polygon, network data storage systems, and a number of specialized services. The platform is designed for application development, high-performance computing, data processing, and storage [57].

For the Govorun supercomputer, RSC Group [58], which is the leading Russian developer and integrator of the full cycle of supercomputer solutions and has a number of its own innovative developments, elaborated an updated, ultradense, scalable, and energy-efficient cluster solution, which is a set of components for building advanced computing systems of various sizes with 100% liquid cooling in hot water mode [59]. It comprises high-performance compute nodes integrated with a high-speed Intel Omni-Path switch with similar hot water cooling.

The GPU component of the Govorun supercomputer contains 96 graphics accelerators from NVIDIA of various generations. The peak performance of the GPU component amounts to 1.4 PFlops for double-precision computations and 58 PFlops for half-precision computations (for artificial intelligence tasks).

The CPU component contains 163 hyperconverged compute nodes with a total of 8552 computing cores. The peak performance of the CPU component is 800 TFlops for double-precision computations.

The hierarchical data processing and storage system of the Govorun supercomputer is based on 18 RSC Tornado TDN511S blade servers and 10 Tornado TDN851 AFS DDSS nodes. The total storage capacity is 10.6 PB, and the data access speed is 300 Gbps.

The total peak performance of the Govorun supercomputer reaches 2.2 PFlops for double-precision computations, 58 PFlops for half-precision calculations (for artificial intelligence tasks), and a read/write speed of 300 Gbps for the hierarchical data processing and storage system (Figure 12).

The technology of direct liquid cooling of RSC Group was chosen for the CPU component of the supercomputer. Thanks to the introduction of these technologies, the Govorun supercom-

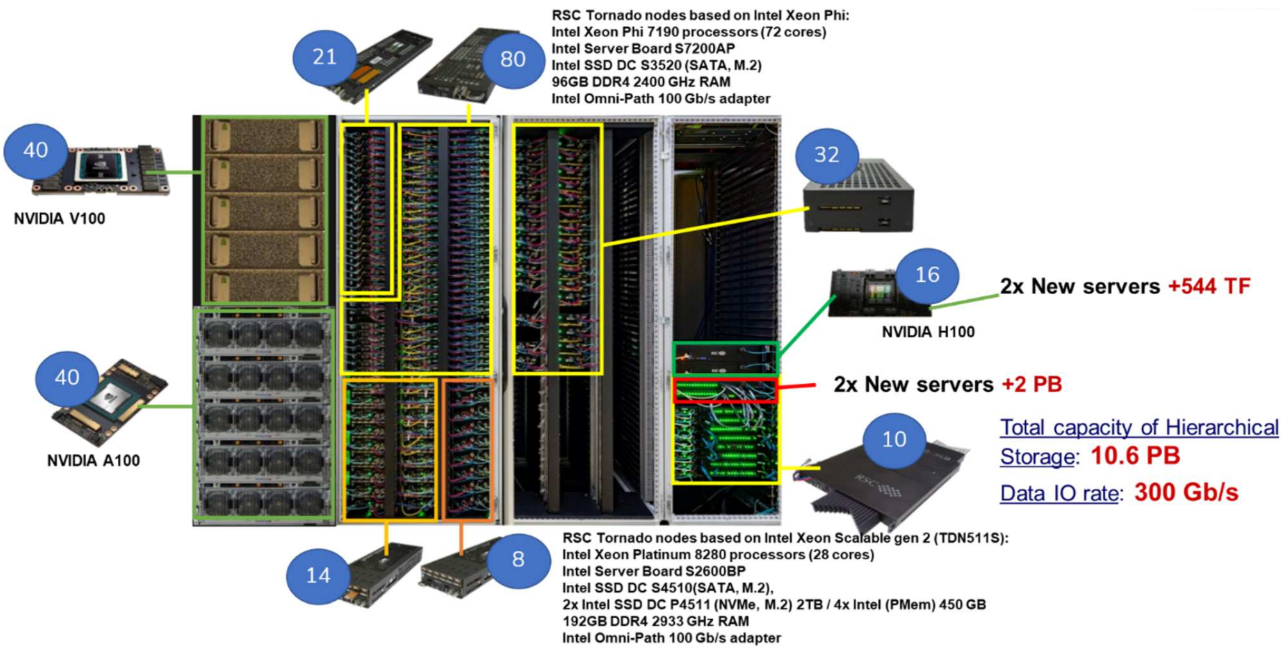

Figure 12. Components of the Govorun supercomputer.

puter managed to achieve a record density of placement of compute nodes per rack (153 nodes vs 25 nodes for air cooling), and the operation in hot water cooling mode made it possible to use year-round free cooling mode. In addition to high-energy efficiency, this approach enabled one to significantly simplify the infrastructure of the supercomputer center, i.e., the cooling system of the Govorun supercomputer was created using only dry cooling towers that cool the liquid using ambient air. Due to liquid cooling, the average annual PUE indicator of the system, reflecting the level of energy efficiency, is less than 1.06. That is, less than 6% of the total electricity consumed is spent on cooling, which is an outstanding result for the HPC industry. The given system is the first system in the world with 100% liquid cooling; all components, namely, compute nodes, network switches, and the data storage system, are cooled.

Another technology underlying the Govorun supercomputer is a hyperconvergent approach to building a computing complex, which enables creating computing environments whose hardware and software configuration is optimized for specific user tasks, without changing the hardware of the compute nodes themselves. Hyperconvergence allows orchestrating computing resources and data storage elements, as well as creating computing systems on demand, taking into account the needs of user applications, with the help of the RSC BasIS [58] software. The notion "orchestration" means the logical disintegration of a compute node into separate components, such as compute nodes, data storage elements, with their subsequent integration into the configuration. Thus, computing elements (CPU cores and graphics accelerators) and data storage elements (SSDs) form independent sets of resources (pools). Due to orchestration, the user can allocate for his task the required number and type of compute nodes (including the required number of graphics accelerators), the required volume and type of data storage systems, as well as automatically configure the required software, including parallel file systems. After the task is completed, the compute nodes and storage elements are returned to their corresponding pools and are ready for the next use. This feature allows users to effectively solve user tasks of different types, to enhance the level of confidentiality of working with data and avoid system errors that occur when crossing the resources for different user tasks. The storage-on-demand system implemented on the hyperconverged nodes of the first modification

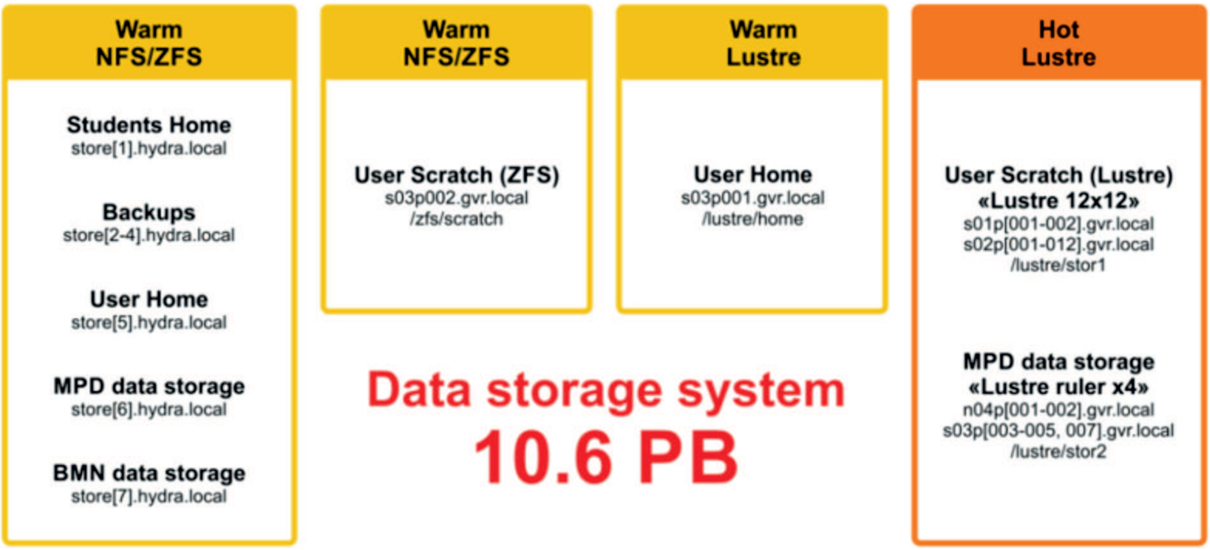

Figure 13. Data storage system of the HybriLIT platform.

under the management of the Lustre file system [60] allowed the Govorun supercomputer to take the 9th place in the IO500 world rating (June 2018) for HPC storage systems (Figure 13).

Thanks to its hyperconvergence, the Govorun supercomputer has a flexible architecture that enables creating software-defined HPC subsystems, which qualitatively distinguishes it from other supercomputers having, as a rule, a "rigid" architecture and designed to effectively solve highly specialized classes of tasks.

The experience of operating the first modification of the Govorun supercomputer revealed the need not only to enlarge computing resources, which was defined by its demand for solving JINR tasks and the growth in the number of users, but also the need to create tools for working with Big Data, primarily for the NICA megaproject. In this regard, a hierarchical data processing and storage system with a software-defined architecture was developed and implemented on the Govorun supercomputer. According to the speed of accessing data, the system is divided into layers, namely, very hot data, the most demanded data, to which it is currently required to provide the fastest access, hot data, and warm data. Each layer of the developed data storage system can be used both independently and as part of data processing workflows. For the high-speed data processing and storage system, the Govorun supercomputer received the prestigious Russian DC Awards 2020 in the Best IT Solution for Data Centers nomination. The further enhancement of the hierarchical data processing and storage system led to the development and commissioning of a geographically distributed data storage system implemented on top of the Lustre parallel file system. This system is employed for fast data copying and simultaneous computing on the Govorun supercomputer, located at MLIT JINR, and on the NCX computing cluster, located at VBLHEP JINR [61].

The tasks of mass generation and data reconstruction within the NICA MPD experiment actively use the hierarchical data processing and storage system of the Govorun supercomputer. At the same time, at different stages of event reconstruction and simulation workflows, there is a need for different access rates to data; for example, for long-term storage tasks, access speed is not an important factor, however, for reconstruction tasks, speed plays a relevant role. In addition, for a number of MPD tasks, there was a need for a large amount of RAM, which resulted in the introduction of hyperconverged nodes with a large amount of memory in the supercomputer architecture. Thus, methodologically, to ensure all workflows for the tasks of the NICA megaproject, a system that combines both computing architectures of different types and the developed hierarchical data processing and storage system was created on the Govorun supercomputer (Figure 13). The computing resources and the hierarchical data processing and storage system of the Govorun supercomputer were integrated into a DIRAC-based distributed heterogeneous environment that includes the resources of JINR and its Member States. To date, the NICA MPD experiment collaboration has completed about 1 300 000 computational tasks using the resources of the Govorun supercomputer for modeling and reconstructing approximately 625 million physics events. This volume of computations required the usage of over 11 million core hours on the Govorun supercomputer, which is equivalent to performing calculations on one computing core during 1270 years.

The experience of using different computing resources of JINR and other MPD collaboration institutes has shown that at present, the use of the Govorun supercomputer resources is the most efficient.

One of the most resource-intensive tasks, utilizing all the computing components and data storage systems of the Govorun supercomputer, is the computational study of QCD under extreme environmental conditions using lattice modeling, conducted by BLTP. In particular, the Govorun supercomputer's resources were used to perform calculations to determine the equation of state for quark–gluon plasma in the presence of an external magnetic field and a nonzero chemical potential, to study the influence of rotation on the properties of gluon plasma, and to investigate the influence of a magnetic field on the confinement/deconfinement transition, the chiral transition at a finite temperature and zero baryon density, and other processes using numerical simulations of QCD on a lattice with physical quark mass. It should be noted that many new results in this area of research were obtained by BLTP scientists using the Govorun supercomputer's resources.

The resources of the Govorun supercomputer are also used for research in the field of nuclear physics, namely, for the study of the structure of light exotic, heavy, and superheavy nuclei and reactions with them, the relativistic molecular and periodic quantum-chemical calculation of superheavy elements and their compounds, the study of changes in the Periodic Law in the region of extremely heavy elements, the study of the electronic structure of elements at the end of the 7th and the beginning of the 8th periods, the study of radiation safety of heavy-ion accelerators at FLNR JINR using Monte Carlo simulation, the modeling of the radiation environment of the DC-140 accelerator complex using the FLUKA software package, etc.

The technologies implemented on the Govorun supercomputer enabled the development of IT solutions, such as the ML/DL/HPC ecosystem, which provide opportunities not only for solving tasks in the field of machine and deep learning, but also for the convenient organization of calculations and analysis of results (Figure 14). One example of such solutions is the information and computing system being developed for processing, analyzing, and visualizing data from radiobiological studies underway at LRB JINR. In particular, the MOSTLIT web service was elaborated and deployed to automate the analysis of radiation-induced foci in cell nuclei [62]. The algorithmic core of the service is trained neural network models for segmenting cell nuclei and detecting foci in cell nuclei. The model was trained on the GPU component of the Govorun supercomputer using data annotated on the CVAT (Computer Vision Annotation Tool) platform deployed in the ML/DL/HPC ecosystem.

The flexible architecture of the Govorun supercomputer allows one not only to carry out calculations, but also to use the supercomputer as a research polygon for working out software–hardware and IT solutions for tasks solved at JINR. One example of such a polygon is the quantum computing polygon deployed on the GPU component of the Govorun supercom-

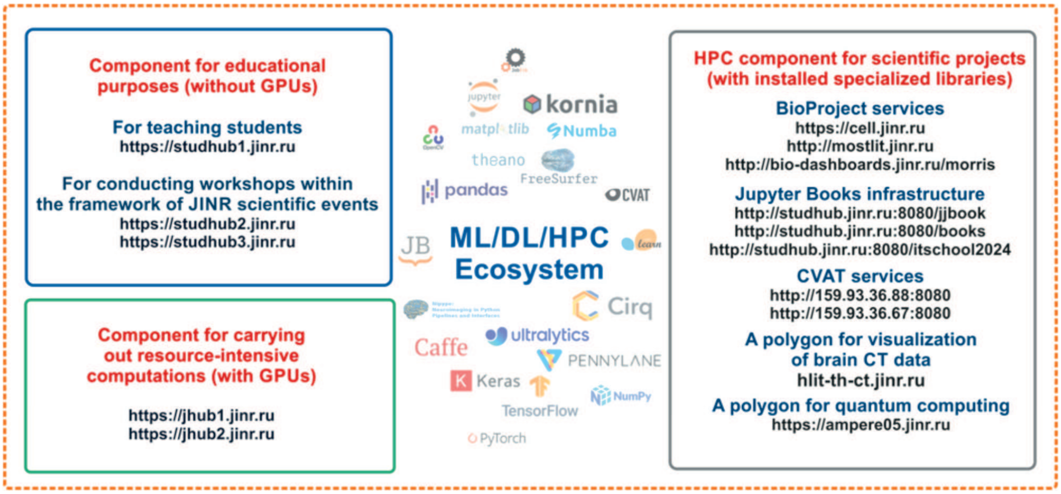

Component for educational purposes (without GPUs)

For teaching students
https://studhub1.jinr.ru

For conducting workshops within the framework of JINR scientific events
https://studhub2.jinr.ru
https://studhub3.jinr.ru

Component for carrying out resource-intensive computations (with GPUs)

https://jhub1.jinr.ru
https://jhub2.jinr.ru

HPC component for scientific projects (with installed specialized libraries)

BioProject services
https://cell.jinr.ru
http://mostlit.jinr.ru
http://bio-dashboards.jinr.ru/morris

Jupyter Books infrastructure
http://studhub.jinr.ru:8080/jjbook
http://studhub.jinr.ru:8080/books
http://studhub.jinr.ru:8080/itschool2024

CVAT services
http://159.93.36.88:8080
http://159.93.36.67:8080

A polygon for visualization of brain CT data
hlit-th-ct.jinr.ru

A polygon for quantum computing
https://ampere05.jinr.ru

Figure 14. ML/DL/HPC ecosystem of the HybriLIT platform.

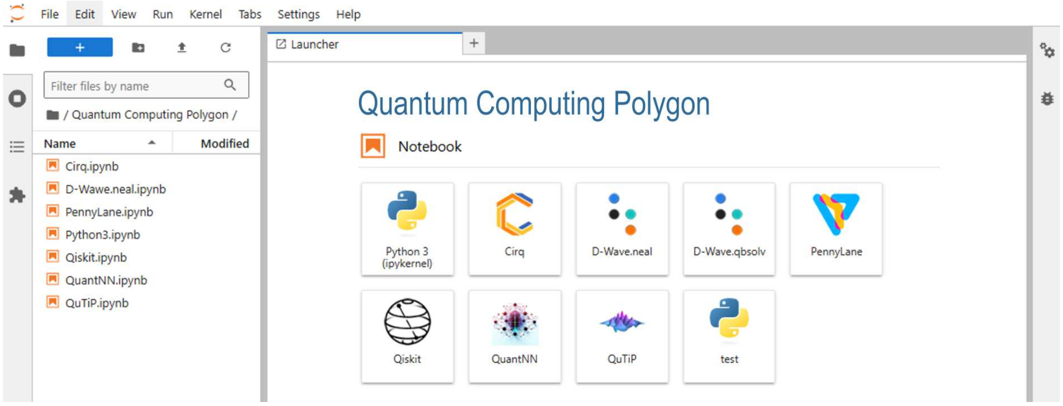

Figure 15. Quantum computing polygon.

puter [63]. JupyterLab was chosen as the polygon's computing environment, allowing users to visually work with quantum circuits and perform computations in a web browser. A wide range of quantum computing simulators, Cirq, Qiskit, PennyLane, etc., is installed on the polygon (Figure 15). The creation of such a polygon is driven by the fact that quantum computing is approaching so-called quantum supremacy in a number of practical applications. However, the implementation of direct quantum computing is currently limited by technological capabilities, including the problem of decoherence and high noise levels. This imposes certain restrictions on the architecture of implemented quantum computers. Therefore, there is a need for the efficient emulation of quantum computing using classical resources. It is noteworthy that the task of emulating quantum computing is itself extremely resource-intensive and requires supercomputer resources. At present, the resources of the quantum computing polygon are used, among other things, to investigate classical quantum computing emulation methods, which can reduce their resource intensity. In particular, a new promising area related to the integration of ML methods, such as neural network optimization, with tensor networks is being studied. Modern ML frameworks installed on the Govorun supercomputer allow for the optimization of tensor network parameters using gradient methods, enhancing the accuracy of representing complex quantum states. Thus, the elaboration of methods for modeling the evolution of quantum systems based on tensor networks using neural network optimization is crucial for overcoming

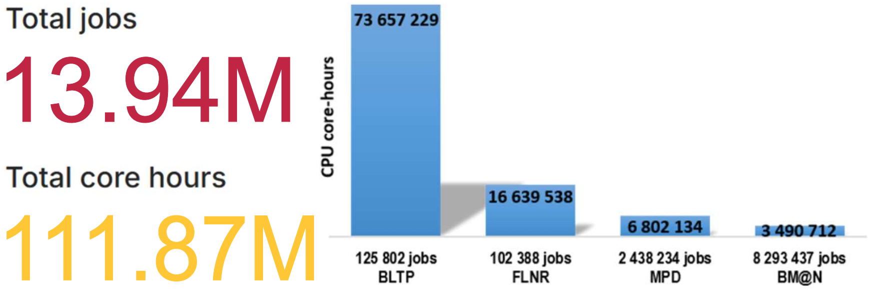

Figure 16. Statistics of the most resource-intensive projects on the Govorun supercomputer.

computational barriers in quantum research and accelerating the development of practically applicable quantum algorithms and technologies.

Since the commissioning of the Govorun supercomputer in July 2018, about 14 million computational tasks have been carried out on its resources, which corresponds to about 112 million core hours (Figure 16). The results obtained using the resources of the Govorun supercomputer are reflected in 509 user publications, two of them in the Nature Physics journal.

Thus, the operation experience of the Govorun supercomputer has shown the relevance and effectiveness of using both novel hyperconverged computing architectures and the hierarchical data processing and storage system being part of it. At present, the resources of the Govorun supercomputer are used by scientific groups from all the Institute's Laboratories. The number of users of the Govorun supercomputer is 357 people.

7. Data storage systems and services

Data storage systems are an essential part of the MICC that ensures data preservation and efficient data operations. The use of a wide range of different storage systems within the MICC is defined by a variety of use cases (Figure 17). Some storage systems have stronger coupling

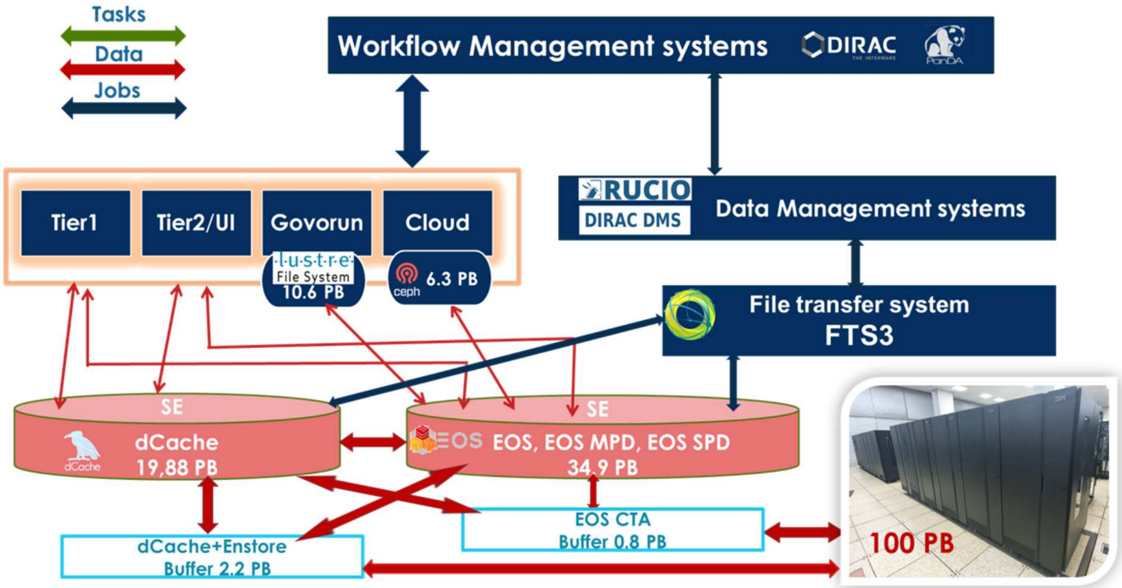

Figure 17. JINR data storage system.

with the corresponding computing resources, like Ceph for the JINR cloud infrastructure or Lustre for the Govorun supercomputer. They are described in previous sections. Other storage systems are declared as general-purpose systems, which means that they are capable of storing data regardless of where the data is processed. In other words, these systems enable efficient data access and processing when using heterogeneous computing resources. These include the EOS and dCache disk-based storage systems, as well as the Enstore and CTA tape-based storage systems. The total JINR storage capacity is about 178 PB (Table 1).

Table 1. JINR storage capacity.

Storage system	Total volume (PB)
dCache (all instances)	23.88
EOS (all instances)	34.9
Lustre on Govorun	10.6
Ceph on cloud	6.3
Enstore (buffer)	2.2
CTA (buffer)	0.8
Tapes (total)	100

dCache [47] is a distributed storage system traditionally used at the MICC grid sites and optimized for the storage of large volumes of experimental data. dCache CMS Tier1 is sampled from CERN and shows satisfactory performance. Local file access is provided via the native dCache DCAP protocol, as well as through standard Grid data access protocols. User authentication and authorization are performed using the Kerberos protocol and X.509 certificates associated with virtual organizations (VOs). Enstore is used for operations within the dCache system.

EOS [50] is a scalable distributed storage system designed for the storage and access of very large data volumes, reaching tens to hundreds of petabytes. It provides the high-performance and convenient storage of experimental data and supports its efficient processing on both computing clusters and user workstations. Within the MICC infrastructure, EOS is mounted on all compute nodes and is accessible as a local file system. Therefore, authorized users can read and write data locally or using standard tools such as *scp* or *rsync*. This access model is consistently applied across all MICC computing resources. In addition, global access to EOS is enabled through WLCG middleware using the GridFTP or xRootD protocols. The existing EOS system shows up to 160 Gbps.

Currently, EOS is used for data storage by a wide range of experiments and user groups, including BM@N [44], MPD [45], SPD [46], Baikal-GVD [22], DANSS, Daya Bay, DsTau, FOBOS, JUNO [21], and PANDA. In total, more than 30 collaborations and user groups at JINR make use of this shared storage infrastructure. To minimize potential interference between different user groups and enable the application of experiment-specific configurations and usage policies, dedicated EOS instances were deployed for the MPD and SPD experiments at NICA.

With the start of full-scale data taking at the experiments of the NICA accelerator complex, a significant expansion of long-term data storage based on robotic tape libraries will be required due to the expected increase in data volumes produced by these experiments and related scientific programs. In addition to the existing tape storage infrastructure used for the CMS experiment on the Tier1 site, in the period from 2024 to 2030, it is planned to deploy a dedicated long-term data storage for the experiments at the NICA complex, the neutrino program, and other user groups.

A system for the long-term storage of user and collaboration data is being developed on the basis of CTA [49]. This system provides reliable long-term data storage using heterogeneous storage media, primarily robotic tape libraries. Its core component is EOS, extended with dedicated services for interacting with tape robotic systems and managing metadata associated with stored files. Currently, the following experiments started using it: BM@N, MPD, SRC, JUNO, and Baikal-GVD.

The compute nodes of the clusters are equipped with limited local disk space, intended primarily for the operating system, auxiliary software, and temporary user files with a short retention period. This local storage is also used during the execution of user tasks in batch or grid mode. The persistent storage of local users' home directories and software is provided by the AFS distributed global file system.

The globally accessible system for managing and distributing large software packages of collaborations and user groups is based on the CERN-developed CVMFS (CERN Virtual Machine File System). CVMFS is used to deploy software required for experimental data processing, allowing applications to run seamlessly across the MICC computing resources. Files and directories are hosted on dedicated servers and mounted in a unified namespace at /cvmfs. Currently, the MICC provides CVMFS-hosted software versions for NICA, BM@N, MPD, SPD, DsTau, ER, JJNANO, JUNO, and Baikal-GVD.

To enable safer and more reliable updates of core software related to Tier1 and Tier2, testbeds were deployed and configured for most important systems, including Slurm, dCache, EOS, and CTA. The testbeds allow for the thorough testing of software updates in a controlled environment, including the verification of interactions between the components, thereby minimizing the risk of conflicts or errors in the production systems.

8. Integration of computing resources

A key element in the organization of distributed heterogeneous computing infrastructures, particularly the WLCG, is specialized middleware that enables the coordinated operation of diverse information and computing systems. This is especially important for the implementation of the experiments within the NICA megascience project (BM@N, MPD, SPD). For example, according to estimates presented in the Technical Design Report "MPD Data Acquisition System" (2018), the raw data rate from the MPD detector is expected to reach at least 6.5 Gbps. For the SPD experiment, preliminary estimates indicate data rates of up to 20 Gbps. The processing, transfer, storage, and analysis of such large data volumes will require substantial computing and storage resources.

In addition to the computing and storage facilities of the MICC, JINR operates a variety of heterogeneous computing resources, including the NICA cluster and the NICA online farm. These resources can be employed to meet the computing requirements of the experiments at the NICA collider. However, they differ in terms of hardware architecture, access and authorization mechanisms, and operational models. The main challenge, therefore, is to integrate these resources into a unified computing environment while preserving their existing modes of operation and without adversely affecting the performance of other tasks.

One option for middleware integration is DIRAC Interware [64], a set of tools designed to integrate heterogeneous computing and data storage resources into a unified system. The key features of DIRAC as an integration tool at JINR include the following:

— the ability to manage computing tasks and data within a single system;

— a built-in set of tools and methods for integrating the major MICC resources, both computing and storage;

— the support for multi-VO mode, in which a single DIRAC instance can be used simultaneously by different scientific groups;

— an active community of DIRAC developers, users, and administrators.

Resource integration with DIRAC is based on the use of standard data access protocols (xRootD, GridFTP, etc.) for storage resources and the use of pilot tasks for the integration of computing resources. Users are provided with a unified environment for tasks submission, data management, workflow automation, and monitoring.

One of the first steps in deploying DIRAC within the MICC was the integration of the cloud infrastructures of JINR and its Member States. This required the development of a dedicated module that enables DIRAC to initiate the creation of virtual machines in the OpenNebula system, which serves as the basis for both the JINR computing cloud and the clouds of the Member States [65]. The developed module was included in the official DIRAC repository and was also used by the BESIII and JUNO experiments. The integration of the cloud infrastructures of the JINR Member States into the DIRAC-based distributed platform provides new opportunities for the Member States to participate in computing for the experiments of the NICA megascience project. Subsequently, all major JINR computing and storage resources available for distributed operation were integrated into the JINR DIRAC system. The integrated resource architecture is shown in Figure 18.

In August 2019, the first batch of data simulation tasks for the MPD experiment was submitted to the Tier1 and Tier2 resources via DIRAC. Afterwards, the Govorun supercomputer was integrated into the distributed computing system. In the summer of 2020, the NICA cluster and the cluster of the National Autonomous University of Mexico (UNAM) were added. EOS was selected as the standard disk storage system for MPD, while access to the tape library was initially provided through the integrated dCache/Enstore system and later changed to CTA tape library. It is noteworthy that the UNAM cluster became the first computing

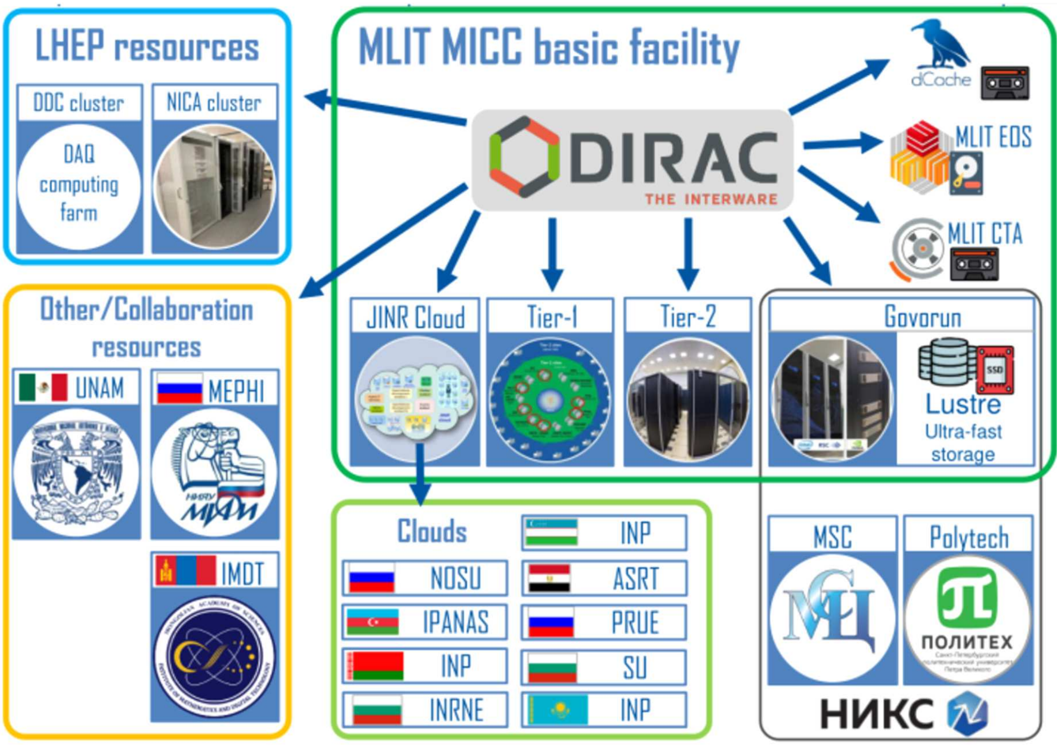

Figure 18. Diagram of resources integrated with the DIRAC platform at JINR.

resource located outside Europe or Asia to be used within the DIRAC infrastructure at JINR. These developments led to the adoption of the DIRAC-based system as a standard solution for centralized large-scale production within the MPD collaboration [66].

The MPD use case and the experience gained during the introduction of new approaches made it possible to propose the DIRAC-based distributed computing platform at JINR as a standard solution for large-scale data processing. Accordingly, the BM@N experiment tested DIRAC for tasks related to raw data simulation and processing and successfully completed the full data processing cycle of BM@N Run 8 within this platform [67]. The Baikal-GVD [68] and SPD experiments used the established distributed infrastructure for Monte Carlo simulation tasks. Table 2 presents the total amount of computing resources consumed by different VOs within DIRAC.

Table 2. Amount of computing resources consumed by different VOs within DIRAC.

VOs	Normalized time (MHS06 days)	Tasks completed (millions)	Walltime (years)	Average task duration (hours)
BM@N	3.9	1.01	459	3.98
MPD	16.8	2.33	2268	8.53
SPD	3.9	0.47	610	11.37
Baikal-GVD	1.2	0.3	168	4.91
Total	25.8	4.11	3505	7.47

The SPD collaboration decided to develop a system based on a combination of PanDA as the workload and workflow management system, Rucio as the data management system, and CRIC as the information system [69]. This ecosystem addresses data processing tasks similar to those handled by DIRAC, but on a considerably larger scale, as demonstrated by its use in the ATLAS experiment. This choice is particularly relevant given that, at peak luminosity, the SPD experiment is expected to collect up to 20 PB of experimental data annually [70]. As part of the preparation of the computing infrastructure for the SPD experiment at NICA, a prototype of a geographically distributed data processing system based on this ecosystem was developed.

9. MICC monitoring system

To ensure the reliable and uninterrupted functioning of the complex, advanced monitoring systems enabling one to anticipate hardware failures and minimize the response time are essential. For these purposes, a multilevel monitoring system that operates 24 × 365 and allows monitoring climate control and power supply systems, local network equipment, telecommunication channels and compute nodes, running tasks, disk and tape storage systems was created and is being expanded at the MICC.

The monitoring system has the properties of scalability (implementation of additional load distribution nodes), heterogeneity (allows introducing any hardware in the monitoring system) and is based on various technologies, such as NagVis, Icinga2, Grafana, InfluxDB2, and systems developed at MLIT.

The system core is built on the Icinga2 software. Icinga2 allows for the flexible configuration of thresholds and hardware parameter polling intervals. Additionally, an alert system is used for early response.

Thresholds are set in accordance with the requirements of engineers managing a specific type of hardware. To organize the monitoring of a specific type of hardware, a custom plugin is being developed. It involves data acquisition and processing. The plugin collects a variety of parameters, which are then converted according to the plugin's logic into a resulting parameter sent to the monitoring system. Typically, this parameter is a numerical value. In addition to this parameter, the monitoring system also receives a status. The status represents the result of the plugin's operation. It embraces the following types: OK (everything works properly), Warning (warning errors), Critical (critical errors), and Unknown (the plugin failed to work). This approach optimizes the monitoring system's disk space by filtering out unnecessary data and simplifies its tracking by specialists. Different thresholds are used to determine the *Warning* and *Critical* states. In our computing complex, early warning works on the basis of the *Warning* threshold. In this case, under certain circumstances, there may be various notifications, including by e-mail or Telegram. In the case of critical emergencies, operator calls may be made.

The hardware polling time is also based on engineer feedback. Different types of hardware have access protocols and hardware limitations that define polling intervals. For example, air conditioner readings are monitored once per minute, while readings from network equipment are monitored every five minutes. This is needed for their stable operation. There are the most critical elements, such as the voltage drop indicator, when it is necessary to poll the hardware with a frequency of fractions of a second. This can only be achieved using tracking hardware (specialized controllers). However, Icinga2 enables including output data from these hardware devices and displaying it on a web page screen within a few seconds.

Monitoring data visualization is based on NagVis, Grafana, and PHP data visualization plugins, providing a powerful tool for analyzing and creating various reports and presentations. An example of visualization of the server infrastructure of the Tier1 grid site is displayed in Figure 19.

The InfluxDB2 database is used to store monitored hardware characteristics, which optimizes the process of data transmission to the Grafana visualization system and organizes data replication to enhance the security of the storage system.

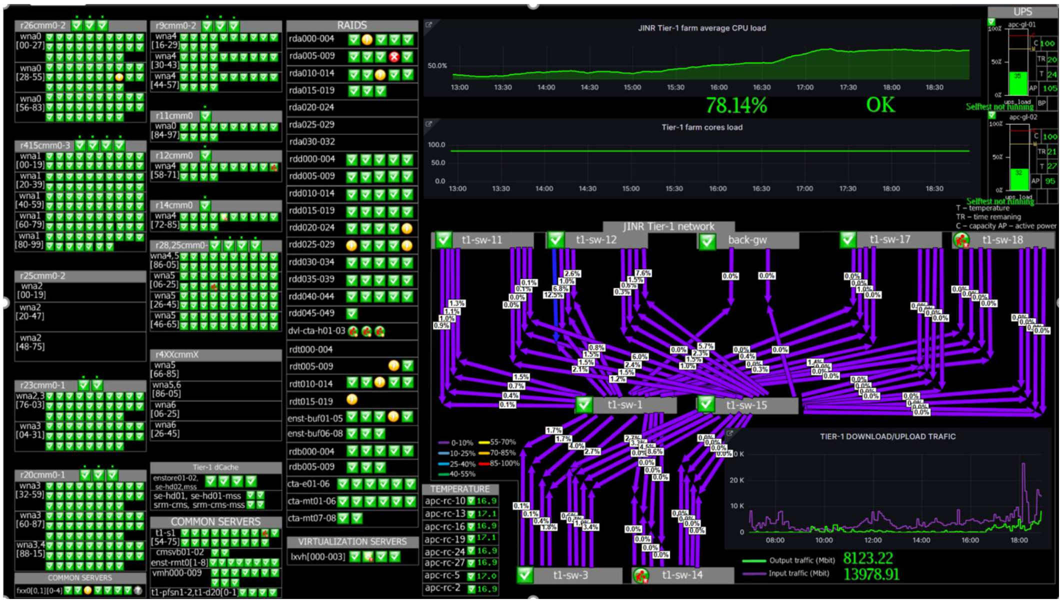

Figure 19. Information screen of the server infrastructure of the Tier1 grid site.

litmon.jinr.ru/grafana-server/d/h1YLNAEMk/ru-jinr-t1-accounting?orgId=1

account: cms_mcore
time_period: day
clock_type: HS06_cpuclock

Figure 20. Information screen of the CMS task accounting system of the Tier1 grid site.

Figure 21. MICC operation center.

The general monitoring system also embraces an accounting system that provides statistical data on resource usage by user tasks (for any time interval): the astronomical task execution time and CPU time in HEP-SPEC06 hours, the number of tasks, and the efficiency of the computing cluster. The system allows for the accounting of resources and their use both within the distributed data processing system and locally. An example of accounting for the CMS experiment's tasks of the Tier1 grid site is shown in Figure 20.

In addition to computing servers, a shared EOS-based data storage was integrated into the accounting system, enabling the prompt monitoring of storage system usage by users and experiments.

To ensure the operational control of the MICC components, the MICC operation center (OC) was developed (Figure 21). The main function of the OC is the round-the-clock surveillance of the state of hardware components, services, the engineering and network infrastructure.

10. Conclusion

The JINR research program for the next decade focuses on ambitious and large-scale experiments both on the basis of the JINR facilities and within international collaboration. This program is related to the implementation of the NICA megaproject, the construction of new experimental facilities, the JINR neutrino program, the modernization of the LHC experimental facilities, and programs in condensed matter physics and nuclear physics. The implementation of the above projects entails decent and commensurate investments in systems that ensure the processing and storage of ever-increasing data volumes. It is known that progress in obtaining research results directly depends on the performance and efficiency of computing and storage resources provided.

The JINR MICC provides scientists with a full range of services for large-scale computing. Its uniqueness lies in the combination of all state-of-the-art information and computing technologies, from grid-based distributed computing and cloud technologies to supercomputer technologies, in a single center. Today, all MICC components are integrated into a unified distributed computing system based on the DIRAC platform, allowing for their use for Monte Carlo simulation and in the full cycle of experimental data processing and analysis.

The next steps in the MICC future development include, among other things, the creation of a long-term data storage center on the MICC resources at MLIT (Tier0) for the NICA megaproject. The process of modeling, processing, and analyzing experimental data obtained from the BM@N, MPD, and SPD detectors will be implemented in a distributed computing environment based on the MICC and the computing centers of VBLHEP and collaboration member countries. Other ongoing efforts include the scaling of computing and storage resources for the JUNO and Baikal-GVD experiments.

The enlargement of the performance of computational facilities and the capacity of storage systems to meet the requirements of the experiments at the LHC in the WLCG project has enabled the Tier1 grid site for the CMS experiment to hold a leading position in the world among seven similar centers.

In connection with the above tasks, it will be necessary to modernize the engineering infrastructure of the JINR MICC, including the design and creation of a new computer room at MLIT in accordance with modern requirements.

Acknowledgements

The authors are grateful to all MLIT staff for their contribution to the development and continuous operation of the JINR Multifunctional Information and Computing Complex. We commend the JINR Directorate for their support and constant attention.

Conflicts of interest

The authors declare no conflicts of interest.

References

[1] JINR Long-Term Development Strategy up to 2030 and beyond: Science & Technology. Dubna: JINR, 2020, 235 p.

[2] V. Korenkov, The JINR Multifunctional Information and Computing Complex, in: Proceedings of the IEEE Xplore:2020 International Scientific and Technical Conference Modern Computer Network Technologies (MoNeTeC), Moscow, Russia, 2020, pp. 1–4. https://doi.org/10.1109/MoNeTeC49726.2020.9258311.

[3] A. Baginyan, A. Balandin, N. Balashov, A. Dolbilov, A. Gavrish, A. Golunov, N. Gromova, I. Kashunin, V. Korenkov, N. Kutovskiy, V. Mitsyn, I. Pelevanyuk, D. Podgainy, O. Streltsova, T. Strizh, V. Trofimov, A. Vorontsov, N. Voytishin, M. Zuev, Current status of the MICC: An overview, CEUR Workshop Proceedings of the 9th International Conference "Distributed Computing and Grid Technologies in Science and Education" 3041 (2021) 1–8. https://ceur-ws.org/Vol-3041/1-8-paper-1.pdf.

[4] Seven-Year Plan for the Development of JINR for 2024–2030. Dubna: JINR, 2023, 76 p.

[5] A. Baginyan, A. Balandin, A. Dolbilov, A. Golunov, N. Gromova, I. Kashunin, V. Korenkov, V. Mitsyn, I. Pelevanyuk, S. Shmatov, T. Strizh, V. Trofimov, A. Vorontsov, N. Voytishin, JINR grid infrastructure: Status and plans, Physics of Particles and Nuclei 55 (3) (2024) 355–359. https://doi.org/10.1134/S1063779624030079.

[6] Worldwide LHC Computing Grid (WLCG), https://wlcg.web.cern.ch/, accessed 2025-12-01.

[7] Nuclotron-based Ion Collider fAcility (NICA), https://nica.jinr.ru/ru/, accessed 2025-12-01.

[8] A. V. Baranov, N. A. Balashov, N. A. Kutovskiy, R. N. Semenov, JINR cloud infrastructure evolution, Physics of Particles and Nuclei Letters 13 (2016) 672–675. https://doi.org/10.1134/S1547477116050071.

[9] OpenNebula: Enterprise Cloud and Virtualization Platform, https://opennebula.io/, accessed 2025-12-01.

[10] E. I. Alexandrov, D. V. Belyakov, M. A. Matveyev, D. V. Podgainy, O. I. Streltsova, Sh. G. Torosyan, E. V. Zemlyanaya, P. V. Zrelov, M. I. Zuev, Research of acceleration calculations in solving scientific problems on the heterogeneous cluster HybriLIT, Bulletin of Peoples' Friendship University of Russia. Series: Mathematics. Information Sciences. Physics 4 (2015) 30–37.

[11] I. Kashunin, V. Mitsyn, V. Trofimov, A. Dolbilov, Integration of the cluster monitoring system based on Icinga2 at JINR LIT MICC, Physics of Particles and Nuclei Letters, 17 (3) (2020) 345–352. https://doi.org/10.1134/S1547477120030073.

[12] S. Chatrchyan et al. (CMS Collab.), The CMS experiment at the CERN LHC, Journal of Instrumentation 3 (2008) S08004. https://home.cern/science/experiments/cms, accessed 2025-12-01.

[13] O. S. Burning et al. (Eds.), LHC Design Report Vol. 1: The LHC Main Ring, CERN-2004-003-V-1. Geneva: CERN, 2004, 548 p. https://home.cern/science/accelerators/large-hadron-collider, accessed 2025-12-01.

[14] V. B. Gavrilov, I. A. Golutvin, O. L. Kodolova, V. V. Korenkov, L. G. Levchuk, S. V. Shmatov, E. A. Tikhonenko, V. E. Zhiltsov, RDMS CMS computing: Current status and plans, Computer Research and Modeling 7 (3) (2015) 395–398.

[15] S. Campana, I. Bird, B. Panzer-Steindel, Overview of the WLCG strategy towards HL-LHC computing — April 2020, LHCC Review (2021). https://doi.org/10.5281/zenodo.5499655.

[16] J. Albrecht, A. A. Alves et al., HEP Software Foundation, A roadmap for HEP software and computing R&D for the 2020s, Computing and Software for Big Science 3 (2019) 7. https://doi.org/10.1007/s41781-018-0018-8.

[17] Evolution of Scientific Computing in the Next Decade: HEP and beyond, http://wlcg-docs.web.cern.ch/wlcg-docs/technical_documents/HEP-Computing-Evolution.pdf, accessed 2025-12-01.

[18] High-Luminosity Large Hadron Collider (HL-LHC), https://home.cern/science/accelerators/high-luminosity-lhc, accessed 2025-12-01.

[19] A Toroidal LHC ApparatuS (ATLAS), https://home.cern/science/experiments/atlas, accessed 2025-12-01.

[20] The Compact Muon Solenoid (CMS), https://home.cern/science/experiments/cms, accessed 2025-12-01.

[21] Jiangmen Underground Neutrino Observatory (JUNO), http://juno.ihep.cas.cn/, accessed 2025-12-01.

[22] The Baikal Deep Underwater Neutrino Telescope (Baikal-GVD), https://baikalgvd.jinr.ru/, accessed 2025-12-01.

[23] Square Kilometre Array (SKA), https://www.skao.int/en, accessed 2025-12-01.

[24] V. V. Korenkov, Trends and prospects of the development of distributed computing and Big Data analytics for support of megascience projects, Physics of Atomic Nuclei 83 (6) (2020) 965–968. https://doi.org/10.1134/S1063778820050154.

[25] S. Campana, A. Di Girolamo, P. Laycock, Z. Marshall, H. Schellman, G. A. Stewart, HEP computing collaborations for the challenges of the next decade, 2022. https://arxiv.org/abs/2203.07237.

[26] D. Guest, K. Cranmer, D. Whiteson, Deep learning and its application to LHC physics, Annual Review of Nuclear and Particle Science 68 (2018) 161–181. https://doi.org/10.1146/annurev-nucl-101917-021019.

[27] S. Farrell, D. Anderson, P. Calafiura, G. Cerati, L. Gray, J. Kowalkowski, M. Mudigonda, Prabhat, P. Spentzouris, M. Spiropoulou, A. Tsaris, J. Vlimant, S. Zheng, The HEP.TrkX Project: Deep neural networks for HL-LHC online and offline tracking, European Physical Journal Web of Conferences 150 (2017) 00003. https://doi.org/10.1051/epjconf/201715000003.

[28] A. Radovic, M. Williams, D. Rousseau, M. Kagan, D. Bonacorsi, A. Himmel, A. Aurisano, K. Terao, T. Wongjirad, Machine learning at the energy and intensity frontiers of particle physics, Nature 560 (2018) 41–48. https://doi.org/10.1038/s41586-018-0361-2.

[29] R. Carrasco-Davis, G. Cabrera-Vives, F. Förster, P. A. Estévez, P. Huijse, P. Protopapas, I. Reyes, J. Martínez-Palomera, C. Donoso, Deep learning for image sequence classification of astronomical events, Publications of the Astronomical Society of the Pacific 131 (1004) (2019) 108006. https://doi.org/10.1088/1538-3873/aaef12.

[30] Beijing Spectrometer (BESIII) Experiment, http://bes3.ihep.ac.cn/, accessed 2025-12-01.

[31] NOvA (NuMI Off-axis ν_e Appearance) Experiment, https://novaexperiment.fnal.gov/, accessed 2025-12-01.

[32] HybriLIT Heterogeneous Platform, http://hlit.jinr.ru/, accessed 2025-12-01.

[33] A. Baginyan, A. Balandin, S. Belov, A. Dolbilov, I. Kadochnikov, V. Korenkov, P. Zrelov, JINR Network Infrastructure for Megascience Projects, in: Proceedings of the IEEE Xplore:2020 International Scientific and Technical Conference Modern Computer Network Technologies (MoNeTeC), Moscow, Russia, 2020, pp. 1–5. https://doi.org/10.1109/MoNeTeC49726.2020.9258004.

[34] V. E. Velikhov, V. V. Korenkov, E. A. Ryabinkin, A. G. Dolbilov, Y. V. Gugel, and T. A. Strizh, The Russian Segment (RU-VRF) in WLCG Infrastructure: High Performance Computing Network, in: Proceedings of the IEEE Xplore:2022 International Conference Modern Computer Network Technologies (MoNeTec), Moscow, Russia, 2022, pp. 1–4. https://doi.org/10.1109/MoNeTec55448.2022.9960772.

[35] LHC Optical Private Network (LHCOPN), http://lhcopn.web.cern.ch/, accessed 2025-12-01.

[36] LHCONE, https://lhcone.web.cern.ch/, accessed 2025-12-01.

[37] GEANT, https://geant.org/, accessed 2025-12-01.

[38] National Research Computer Network of Russia (NIKS), https://niks.su/, accessed 2025-12-01.

[39] A Large Ion Collider Experiment (ALICE), https://home.cern/science/experiments/alice, accessed 2025-12-01.

[40] N. S. Astakhov, S. D. Belov, I. N. Gorbunov, P. V. Dmitrienko, A. G. Dolbilov, V. E. Zhiltsov, V. V. Korenkov, V. V. Mitsyn, T. A. Strizh, E. A. Tikhonenko, V. V. Trofimov, S. V. Shmatov, The Tier-1-level computing system of data processing for the CMS experiment at the Large Hadron Collider, Journal of Information Technologies and Computing Systems 4 (2013) 27–36.

[41] V. Korenkov, I. Pelevanyuk, A. Tsaregorodtsev, DIRAC at JINR as a general-purpose system for massive computations, Journal of Physics: Conference Series 2438 (2023) 012029. https://doi.org/10.1088/1742-6596/2438/1/012029.

[42] Slurm Workload Manager, https://slurm.schedmd.com/overview.html, accessed 2025-12-01.

[43] Advanced Resource Connector (ARC), https://www.nordugrid.org/arc/arc6/, accessed 2025-12-01.

[44] BM@N (Barionic Matter at Nuclotron) Experiment, https://bmn.jinr.int/, accessed 2025-12-01.

[45] MPD (Multi-Purpose Detector), https://mpd.jinr.ru/, accessed 2025-12-01.

[46] SPD (Spin Physics Detector), https://spd.jinr.int/, accessed 2025-12-01.

[47] dCache — Distributed Storage for Scientific Data, https://www.dcache.org/, accessed 2025-12-01.

[48] A. Moibenko, J. Bakken, E. Berman, C. Huang, D. Petravick, M. Zalokar, The Status of Fermilab Enstore Data Storage System, in: Proceedings of the Computing in High Energy Physics and Nuclear Physics (CHEP 2004), Interlaken, Switzerland, 2004, p. 1210. https://doi.org/10.5170/CERN-2005-002.1210.

[49] M. Davis, J. Afonso, R. Bachmann, V. Bahyl, J. Camarero Vera, J. Leduc, P. Oliver Cortés, F. Rademakers, L. Wardenær, V. Yurchenko, The CERN Tape Archive beyond CERN: An open source data archival system for HEP, European Physical Journal Web of Conferences 295 (2024) 01048. https://doi.org/10.1051/epjconf/202429501048.

[50] EOS Open Storage, https://eos-web.web.cern.ch/eos-web/, accessed 2025-12-01.

[51] Rucio — Scientific Data Management, https://rucio.cern.ch/, accessed 2025-12-01.

[52] HTCondor Software Suit, https://htcondor.org/htcondor/overview/, accessed 2025-12-01.

[53] Project JupyterHub, https://jupyter.org/hub, accessed 2025-12-01.

[54] Ceph, https://ceph.io/en/, accessed 2025-12-01.

[55] iTop Platform, https://combodo.com/, accessed 2025-12-01.

[56] N. A. Balashov, I. S. Kuprikov, N. A. Kutovskiy, A. N. Makhalkin, Ye. Mazhitova, I. S. Pelevanyuk, R. N. Semenov, D. A. Shpotya, Changes and challenges at the JINR and its Member States cloud infrastructures, Physics of Particles and Nuclei 55 (3) (2024) 366–370. https://doi.org/10.1134/S1063779624030092.

[57] A. Anikina, D. Belyakov, T. Bezhanyan, M. Kirakosyan, A. Kokorev, M. Lyubimova, M. Matveev, D. Podgainy, A. Rahmonova, S. Shadmehri, O. Streltsova, S. Torosyan, M. Vala, M. Zuev, Structure and Features of the Software and Information Environment of the HybriLIT Heterogeneous Platform, in: Proceedings of the 27th International Conference “Distributed Computer and Communication Networks”, Lecture Notes in Computer Science 15460 (2024) 444–457. https://doi.org/10.1007/978-3-031-80853-1_33.

[58] RSC Group, https://rscgroup.ru/, accessed 2025-12-01.

[59] E. A. Druzhinin, A. B. Shmelev, A. A. Moskovsky, V. V. Mironov, A. Semin, Server level liquid cooling: Do higher system temperatures improve energy efficiency?, Supercomputing Frontiers and Innovations 3 (1) (2016) 67–73. https://doi.org/10.14529/jsfi160104.

[60] Lustre, http://www.lustre.org/, accessed 2025-12-01.

[61] D. V. Belyakov, A. A. Kokorev, D. V. Podgainy, The distributed parallel file system Lustre for processing and analyzing data of the NICA megascience project, Physics of Particles and Nuclei (2026) (accepted).

[62] S. Shadmehri, T. Bezhanyan, M. Bondarev, O. Streltsova, M. Zuev, A. Chigasova, A. Osipov, N. Vorobytea, A. N. Osipov, A deep learning model for automated quantification of DNA repair foci in somatic mammalian cells, Physics of Particles and Nuclei 56 (2025) 1623–1627. https://doi.org/10.1134/S1063779625700984.

[63] Yu. Palii, D. V. Belyakov, A. A. Bogolubskaya, M. I. Zuev, D. A. Yanovich, Simulation of the QAOA algorithm at the JINR quantum testbed, Physics of Particles and Nuclei 56 (2025) 989–993. https://doi.org/10.1134/S1063779625700066.

[64] V. Korenkov, I. Pelevanyuk, A. Tsaregorodtsev, Integration of the JINR Hybrid Computing Resources with the DIRAC Interware for Data Intensive Applications, in: Proceedings of the International Conference "Data Analytics and Management in Data Intensive Domains", Springer International Publishing, 2020, pp. 31–46. http://dx.doi.org/10.1007/978-3-030-51913-1_3.

[65] N. Balashov et al., Cloud integration within the DIRAC interware, CEUR Workshop Proceedings 2507 (2019) 256–260.

[66] N. Kutovskiy et al., Integration of distributed heterogeneous computing resources for the MPD experiment with DIRAC interware, Physics of Particles and Nuclei 52 (4) (2021) 835–841. http://dx.doi.org/10.1134/S1063779621040419.

[67] K. V. Gertsenberger, I. S. Pelevanyuk, BM@N Run 8 data processing on a distributed infrastructure with DIRAC, Physics of Particles and Nuclei Letters 21 (4) (2024) 778–781. https://doi.org/10.1134/S1547477124701334.

[68] N. Kutovskiy, I. Pelevanyuk, D. Zaborov, Using distributed clouds for scientific computing, CEUR Workshop Proceedings of the 9th International Conference "Distributed Computing and Grid Technologies in Science and Education" (2021) 196–201. http://dx.doi.org/10.54546/MLIT.2021.78.51.001.

[69] A. Petrosyan, D. Oleynik, A. Zhemchugov et al., Production system of the SPD experiment, Physics of Particles and Nuclei 56 (2025) 1576–1580. https://doi.org/10.1134/S1063779625700893.

[70] V. Abazov et al. (SPD Collab.), Technical Design Report of the Spin Physics Detector at NICA, Natural Science Review 1 (1) (2024). https://doi.org/10.54546/NaturalSciRev.100101.