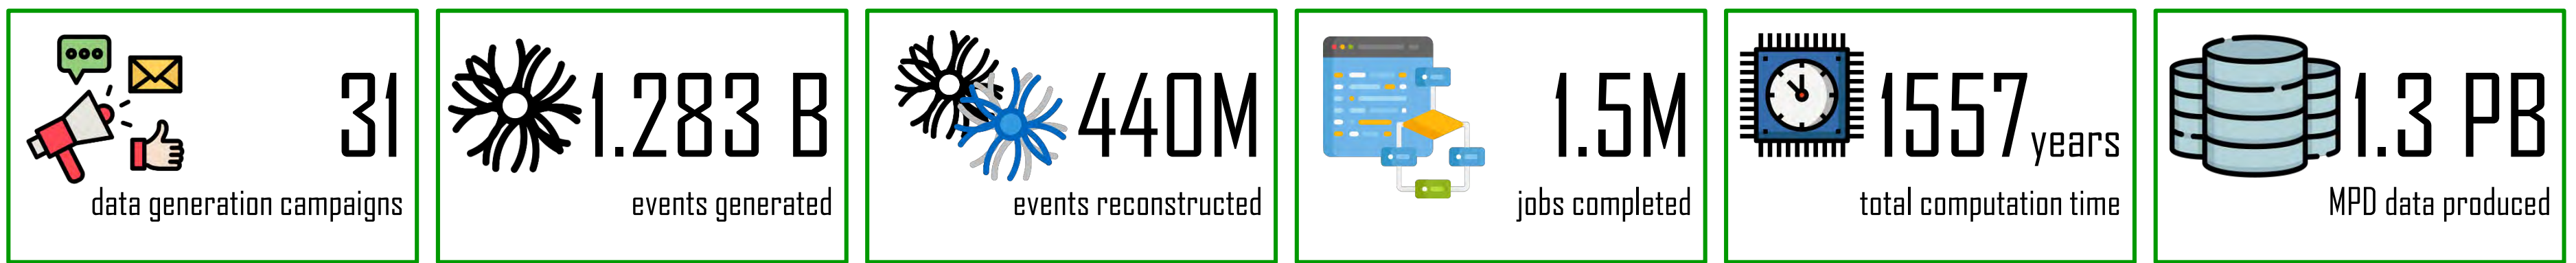


Distributed heterogeneous computing infrastructure for the MPD experiment


V.V. Mitsyn¹, A.A. Moshkin², I.S. Pelevanyuk¹, D.V. Podgayny¹, O.V. Rogachevsky², V.G. Riabov², M.I. Zuev¹
¹ Meshcheryakov Laboratory of Information Technologies
² Veksler and Baldin Laboratory of High Energy Physics



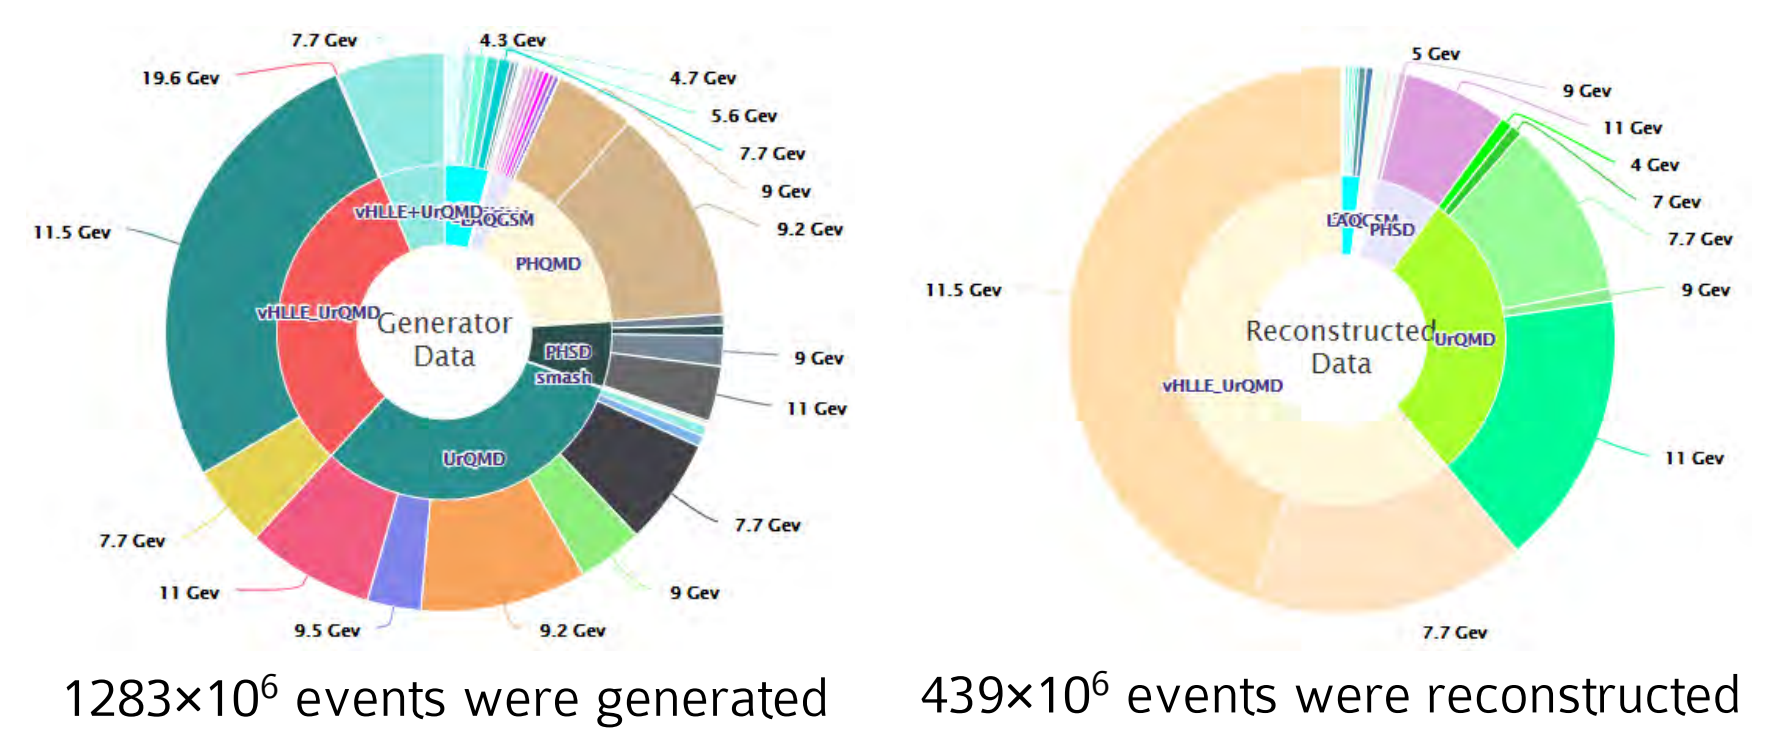
MPD Physics

The main goal of the NICA/MPD project is to study heavy ion collisions at energies $\gamma/\text{sNN}=4-11$ GeV to study the properties of hadrons in a superdense nuclear medium, as well as the equation of state of nuclear matter and the properties of phase transitions, including the search for possible deconfinement signals, a critical point and partial restoration of chiral symmetry. The physical objectives of the experiment are highly ambitious and difficult, both because of the large uncertainties in theoretical predictions and as a result of the very high accuracy required to measure small effects. We plan to carry out high-precision scanning of the phase diagram by varying the collision energy and the size of colliding systems with a change in the impact parameter and/or atomic weights of particle beams.

At the moment, with the help of data simulation, the possibility of implementing the physical program at the MPD experimental facility is being evaluated.

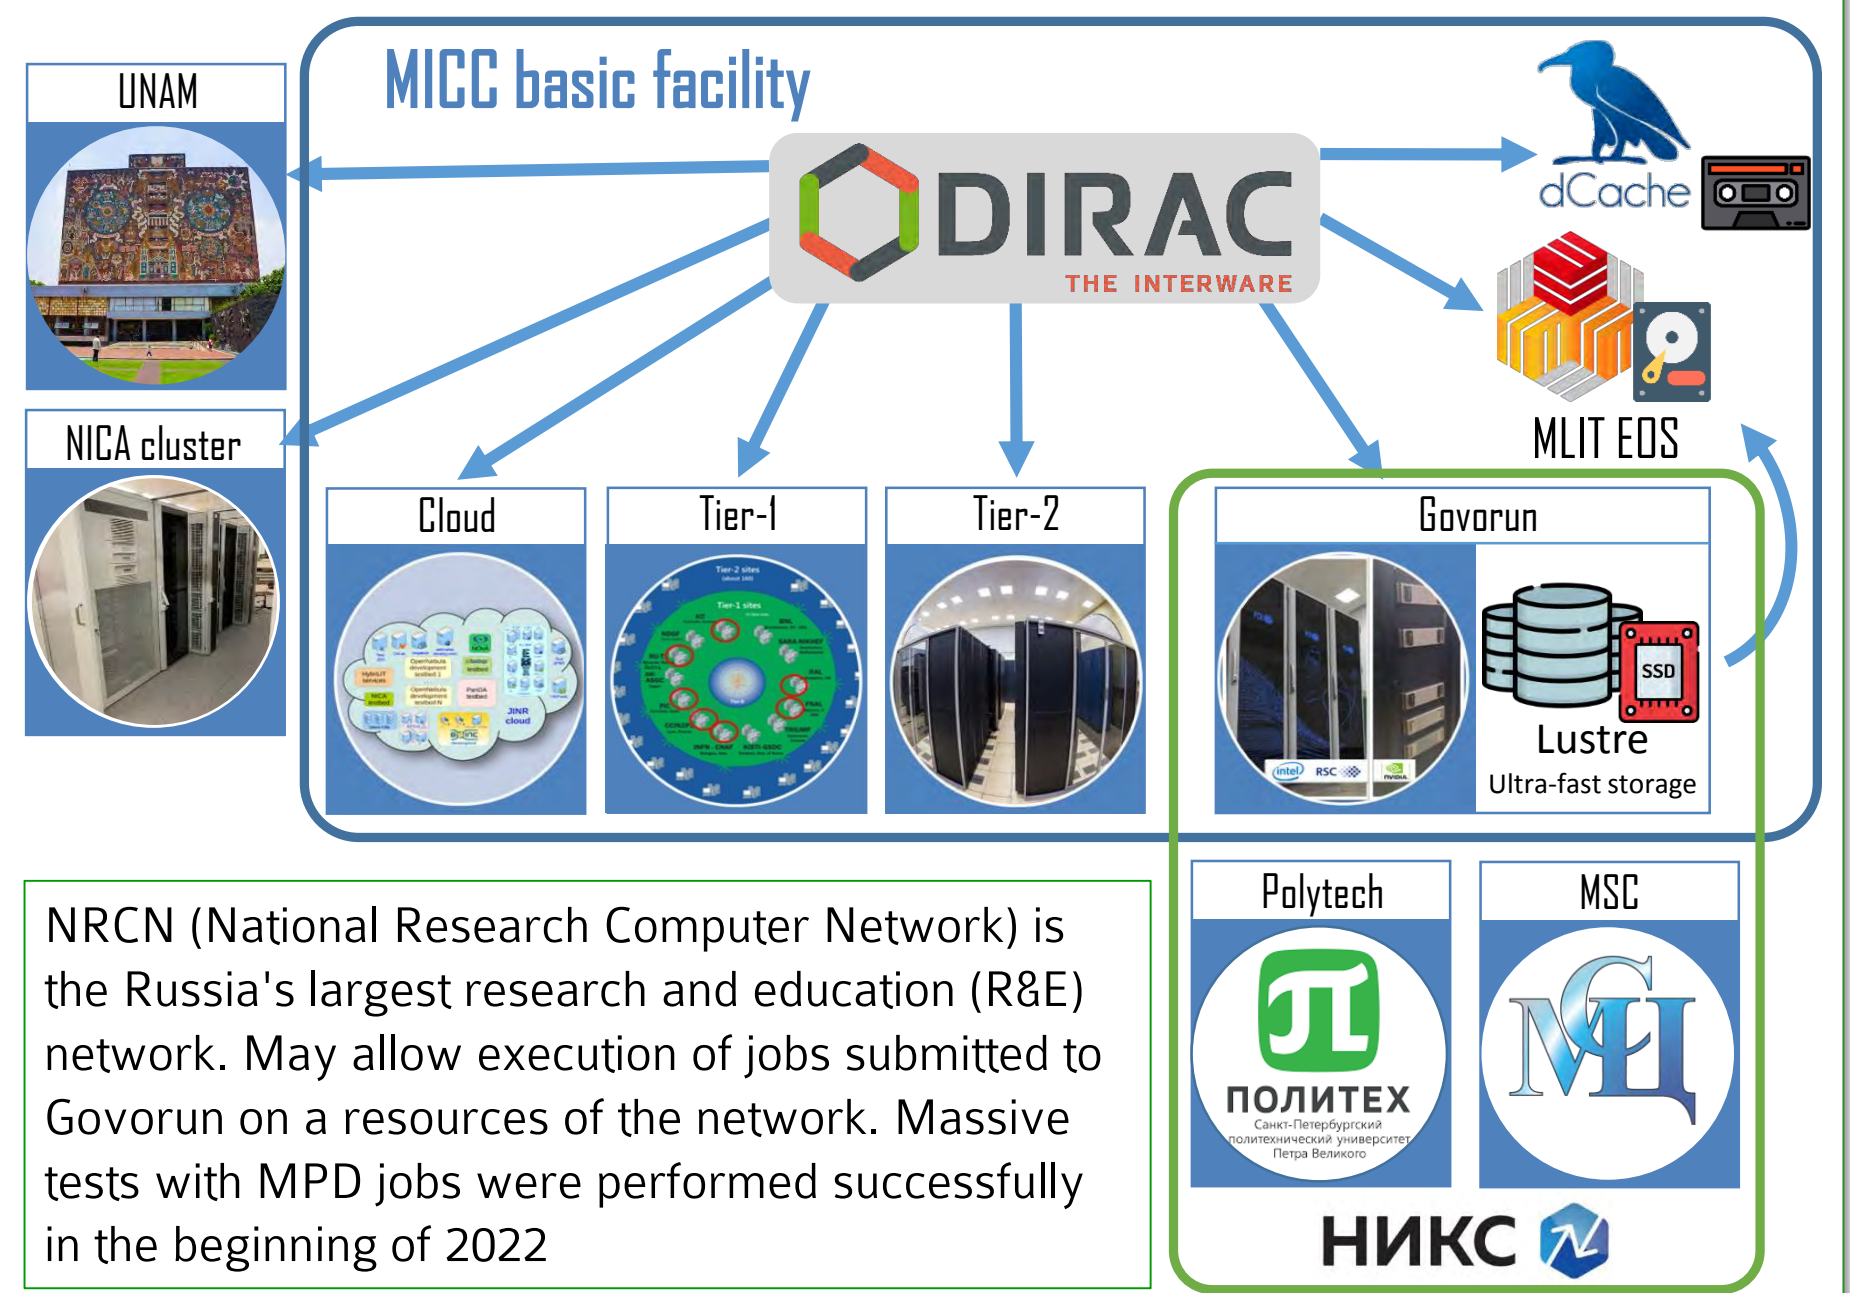


31 mass production requests were completed



Distributed Infrastructure

Since 2016, a distributed heterogeneous computing infrastructure, created using the DIRAC platform, has been being developed at JINR. It is used to process data from JINR experiments. At present, the system combines all MICC resources, the clouds of the JINR Member States, the resources of the National Research Computer Network of Russia and the cluster of the National Autonomous University of Mexico (being part of the MPD collaboration). Using the DIRAC platform makes it possible to get unified access to all integrated computing and storage resources, to carry out performance analysis and account for consumed resources.

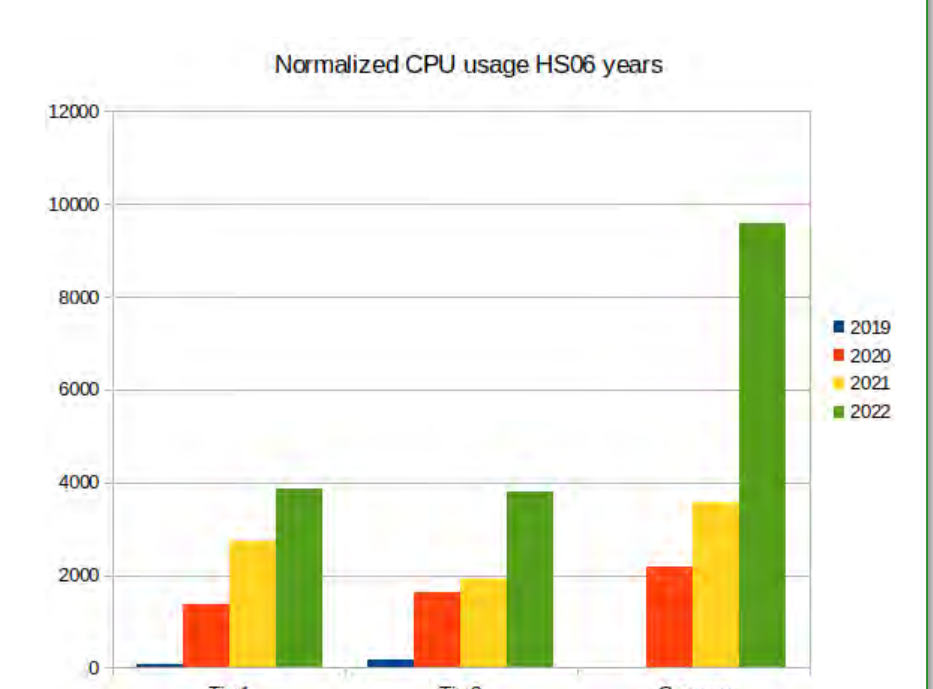
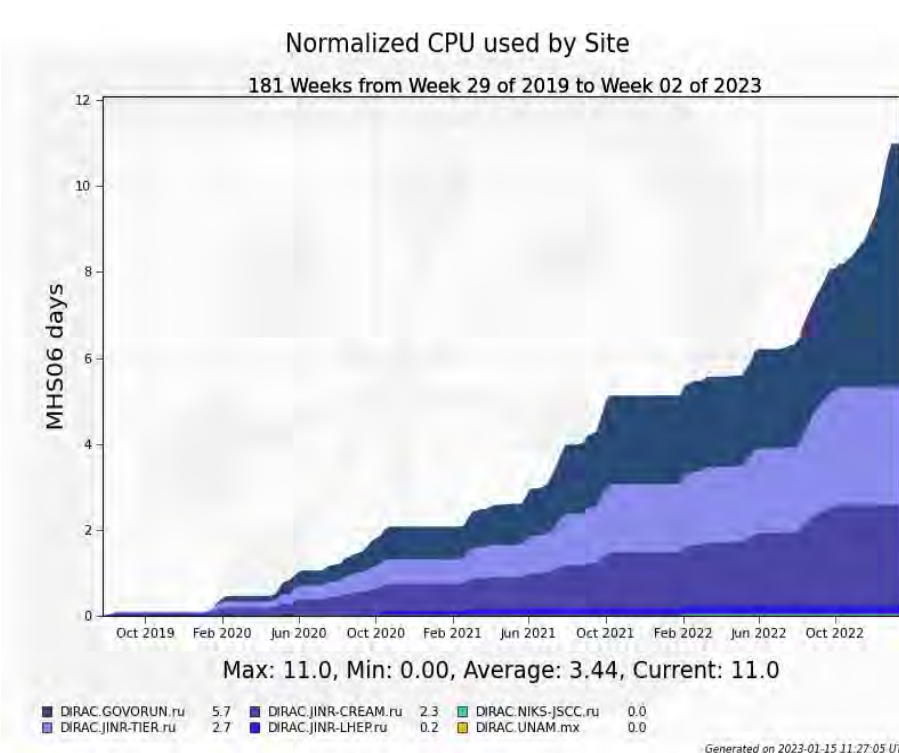


NRCN (National Research Computer Network) is the Russia's largest research and education (R&E) network. May allow execution of jobs submitted to Govorun on a resources of the network. Massive tests with MPD jobs were performed successfully in the beginning of 2022

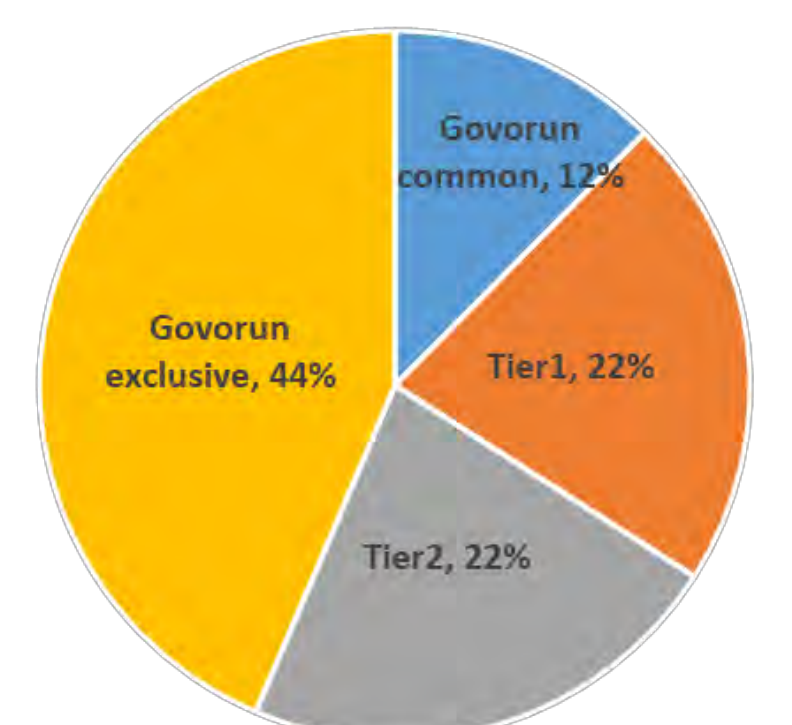

Usage

Since 2019, the MPD collaboration has run 31 data generation campaigns. 1.5 million jobs with a total duration of 1,535 years were completed. In total, 1.3 billion events were generated, 439 million of which were also reconstructed. The total amount of MPD data in the distributed network reached 1.3 PB with a total storage capacity of over 2 PB.

	2019	2020	2021	2022	Total
Executed jobs (k)	70	445	505	514(+1%)	1534
Normalized CPU time (kH506 days)	88	1980	3060	5710(+86%)	10838
Walltime (years)	16	323	444	774(+74%)	1557



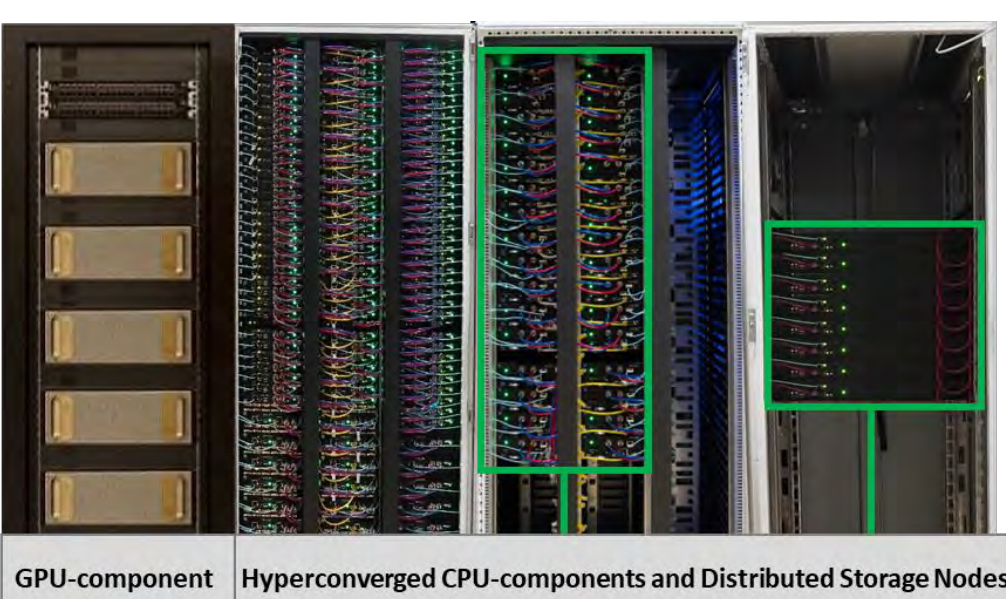
Computing work of MPD using DIRAC 2022



Modernized supercomputer "Govorun" for the MPD experiment

In 2022, the next stage of modernization of the "Govorun" supercomputer was carried out. The CPU components, implemented as part of a hyperconvergent approach to building a computing complex, which is the basis of SC "Govorun", were extended. As a result of the upgrade of the CPU component, the performance of SC "Govorun" increased by 239 Tflops and reached the level of 1.1 Pflops. The volume of the hierarchical data processing and storage system of SC "Govorun" increased by 10 PB. The hyperconvergence of new compute nodes has already made it possible to use them for the tasks of mass generation and reconstruction of data within the NICA MPD experiment.

Computation field: **+32 hyperconverged compute nodes** Hierarchical Storage: **+8 distributed storage nodes** Performance: **+239 Tflops DAOS: +1.6 PB Lustre, EOS: +8 PB**



SC "Govorun" current status:
138 hyperconverged compute nodes
40 GPU accelerator
Total peak performance: 1.1 PFlops DP, 2.2 PFlops SP
Total capacity of hierarchical storage: 8.6 PB
Data IO rate: 300 Gb/s

+1,152 new computational cores for the MPD experiment
Nearly 31 million events were generated in less than month!
 The performance "new CPU"/"old CPU" increased by more than **1.5 times**
+0.4 PB for MPD mass production storages integrated into the DIRAC File Catalog
+1 PB for MPD EOS storage

