



Heterogeneous Platform

HybriLIT 2025

Towards Efficient
Scientific Computing

Meshcheryakov Laboratory
of Information Technologies, JINR, Dubna



HybriLIT: Towards Efficient Scientific Computing

The brochure showcases the capabilities of the “Govorun” supercomputer, which serves as the main computational component of the HybriLIT heterogeneous platform. It also introduces advanced high-performance computing technologies and GPU-accelerated workflows, including modern tools for resource-intensive computations used in simulation and experimental data processing — including Big Data — in physics, biology, and research aligned with the key projects of the Joint Institute for Nuclear Research.

The publication was prepared by the HybriLIT heterogeneous computing team and is dedicated to the 70th anniversary of JINR and the 60th anniversary of the Laboratory of Information Technologies.

HybriLIT Team

Project leads:

George Adam
Dmitry Podgainy
Oksana Streltsova

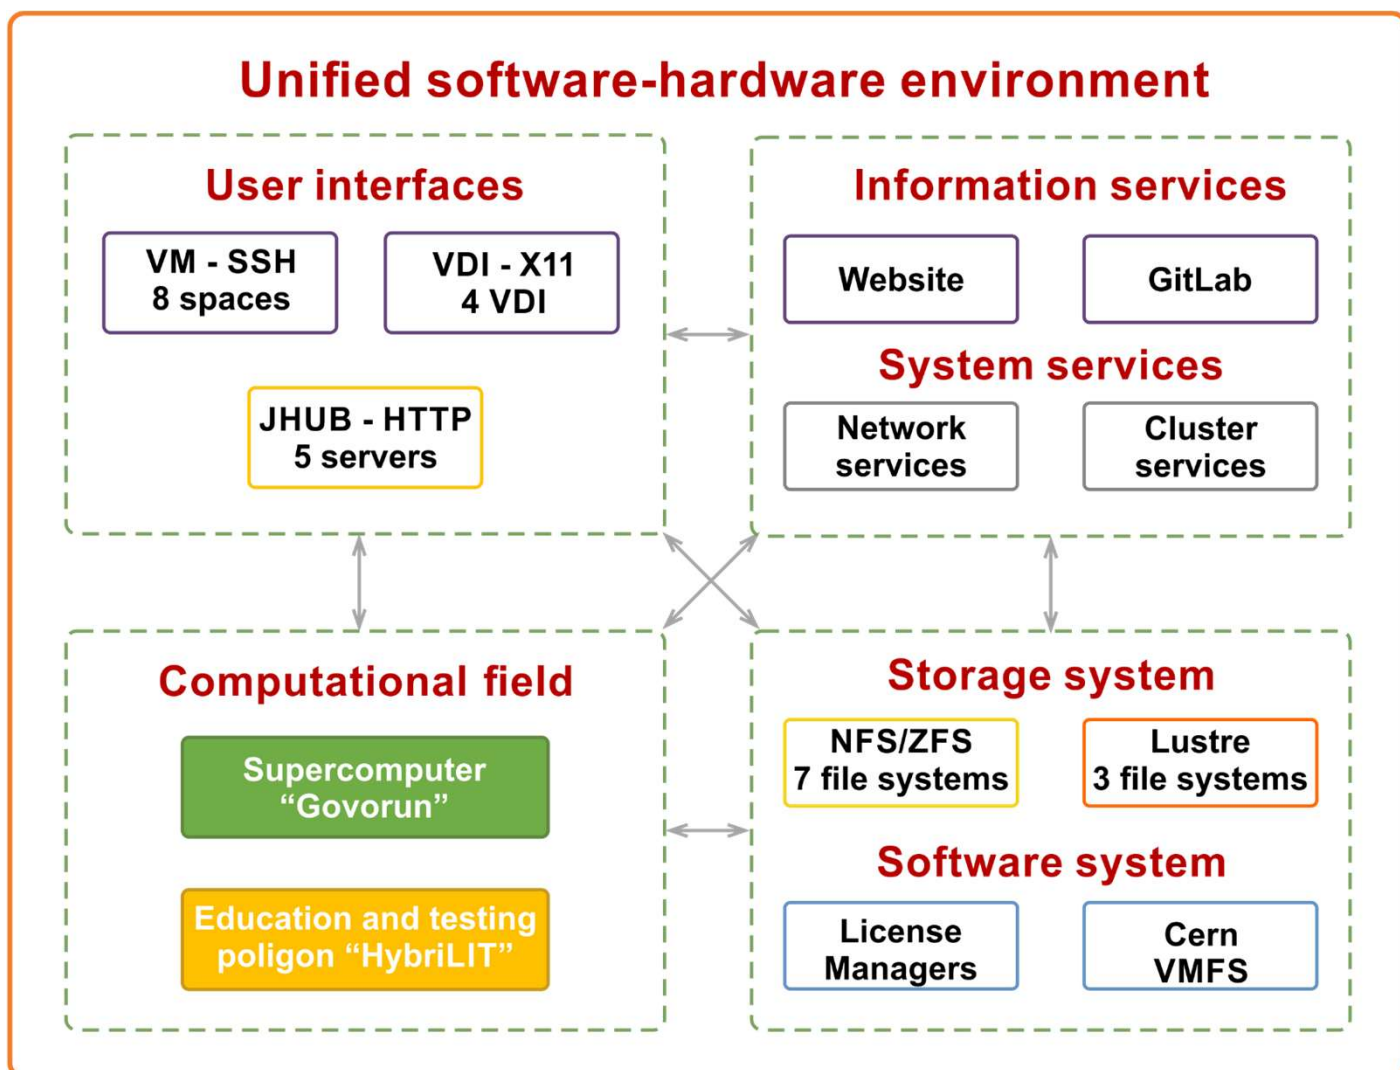
HybriLIT Group:

Anastasia Anikina	Igor Pelevanyuk
Alexander Ayriyan	Adiba Rahmonova
Maksim Bashashin	Tatiana Sapozhnikova
Dmitry Belyakov	Sara Shadmehri
Tatevik Bezhanyan	Alexey Streltsov
Ján Buša Jr.	Shushanik Torosyan
Konstantin Chizhov	Martin Vala
Mohamed Gaber	Lucia Valova
Margarit Kirakosyan	Alexey Vorontsov
Aleksandr Kokorev	Tatiana Zaikina
Maria Lyubimova	Elena Zemlyanaya
Mikhail Matveev	Maxim Zuev
Andrey Moshkin	



Heterogeneous Computing Platform HybriLIT

The HybriLIT heterogeneous computing platform is part of the Multifunctional Information and Computing Complex of the Meshcheryakov Laboratory of Information Technologies at JINR.



The main components of the Platform include: *the computing environment*, represented by the Govorun supercomputer and the training and testing ground; *the data storage system*, represented by several network file systems (NFS/ZFS and Lustre); *the software distribution system*, implemented using license managers (FlexLM/MathLM) and a read-only network file system (CernVM-FS); *user interfaces* providing access to the Platform's resources in various modes; *system services* that ensure the operation of computing nodes within the cluster and the supercomputer; and *information services* designed to support users.

<http://hlit.jinr.ru/>

Heterogeneous Computing Platform HybriLIT

2014



Full peak performance:
140 TFlops for single precision
50 TFlops for double precision

2018




#18 в Top50
Full peak performance:
1 PFlops for single precision
500 TFlops for double precision
9th in the current edition of the IO500 list (July 2018)

2019



#10 в Top50
Full peak performance:
1.7 PFlops for single precision
860 TFlops for double precision
288 TB DSS
with I/O speed >300 Gb/s
17th in the current edition of the IO500 list (July 2020)

2023




Full peak performance:
3.4 PFlops for single precision
1.7 PFlops for double precision

2025




Russian DC Awards 2020 in "The Best IT Solution for Data Centers"




Full peak performance:
2.2 PFlops for double precision
58 PFlops for half precision

Supercomputer "Govorun"



PUE ~1,06

Power usage effectiveness



The "Govorun" supercomputer is equipped with a precision liquid-cooling system developed by RSC. The supercomputer is supplied with water cooled to 45 °C. As it circulates through the entire system, the water heats up to 50 °C and then returns to a heat exchanger, where it is cooled by transferring thermal energy to the dry cooling tower's water circuit.


The cooling system features smooth performance regulation, allowing the cooling capacity to be increased or decreased according to the actual load. This makes it possible to significantly reduce power consumption when the supercomputer is operating under partial load.

<http://hlit.jinr.ru/engineering-infrastructure/>

Total Peak Performance

2.2 PFLOPS double precision

58 PFLOPS half precision



more than

8500

CPU cores

21 servers with Intel Xeon Phi

Intel Xeon Phi 7290 (72 cores @1.50 GHz), 96 GB RAM

76 servers with Intel Xeon Scalable Gen2 (RSC Tomado TDN511)

2x Intel Xeon Platinum 8268 (24 Cores @2.90 GHz), 192 GB RAM

32 servers with Intel Xeon Scalable Gen2 (RSC Tomado TDN511S)

2x Intel Xeon Platinum 8368Q (38 Cores @2.60 GHz), 2 TB RAM

Peak performance: 800 TFLOPS double precision



5 servers with NVIDIA V100

2x Intel Xeon E5-2698 v4 (20 cores @2.20 GHz),

8x NVIDIA V100 16 GB, 512 GB RAM

5 c servers with NVIDIA A100

2x AMD EPYC 7763 (64 Cores @2.45 GHz),

8x NVIDIA A100 80 GB, 2 TB RAM

2 RSC ExaStream AI servers with Hopper H100 GPUs:

2x Intel Xeon Platinum 8468 (48 Cores @2.1 GHz),

8x NVIDIA H100 80 GB, 1 TB RAM

Peak performance: 58 PFLOPS half precision


Storage system

10.6 PB

http://hlit.jinr.ru/supercomputer_govorun/#_HardwareSG

Data Storage Systems

The “Govorun” supercomputer includes an RSC Storage on-Demand network storage system, which functions as a unified, centrally managed system and provides several data storage tiers — very hot data, hot data, and warm data.



Very Hot Data Storage System

4x RSC Tornado TDN511S

(12x Intel Optane SSD DC P4801X 375GB M.2, IMDT) 4,2 TB

Hot and Warm Data Storage System

14x RSC Tornado TDN511S

76x RSC Tornado TDN511

12x Intel Optane SSD DC P4801X 375GB M.2 Series

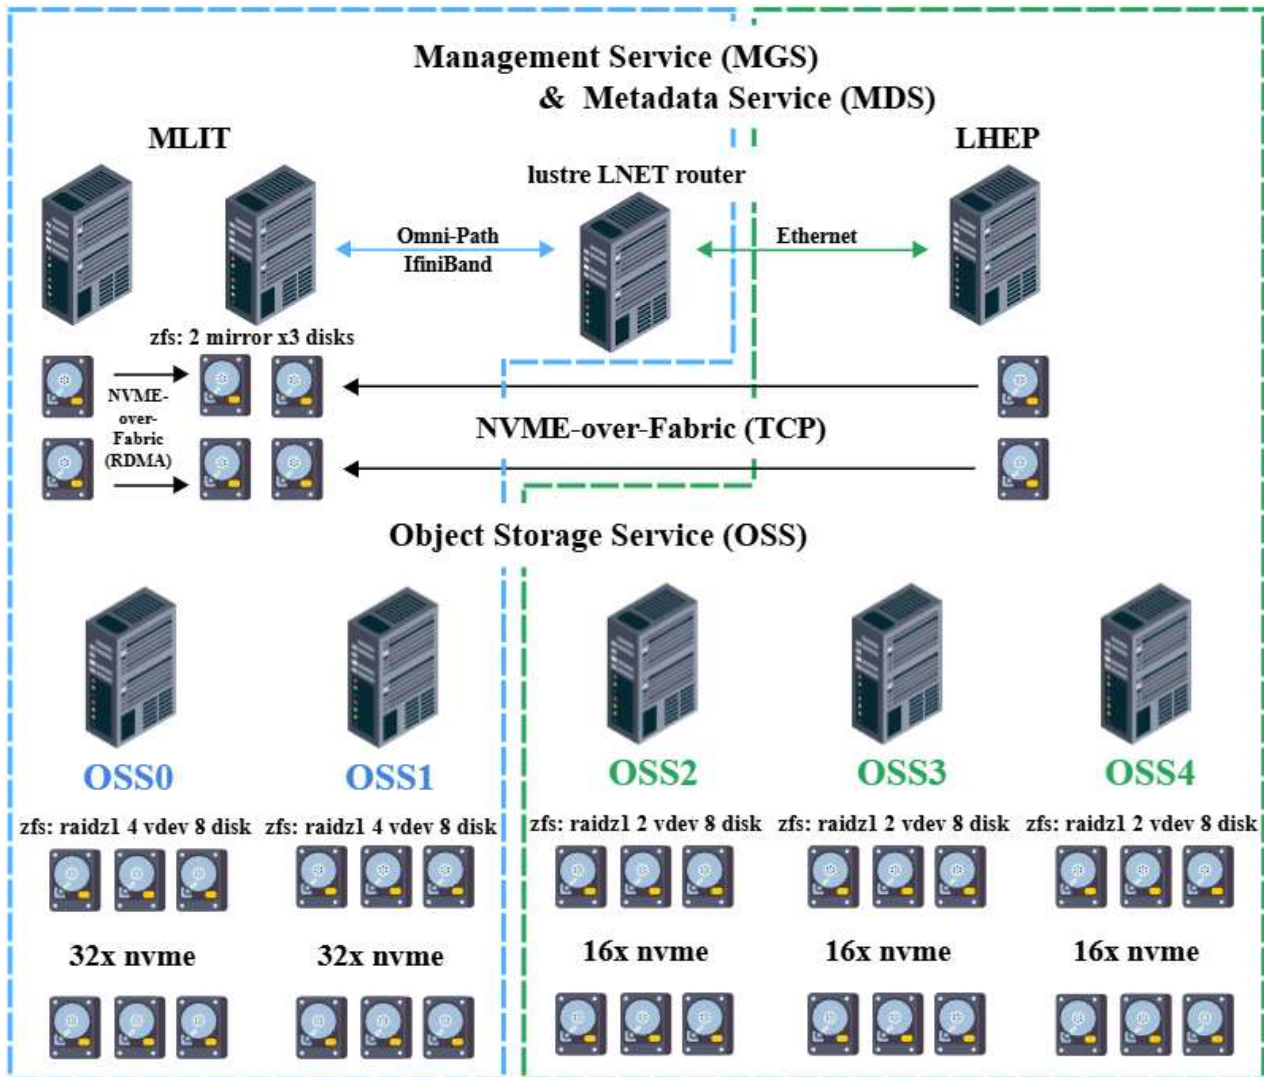
12x Intel SSD DC P4511 (NVMe, M.2)



http://hlit.jinr.ru/supercomputer_govorun/#_DataStorage

Distributed File System Lustre

The Lustre distributed file system consists of two segments: the first segment is located at MLIT JINR, and the second segment at the VBLHEP JINR.



Management Servers

2x RSC Tornado TDN711:

2x Intel Xeon Platinum 8268
96 cores @2.90GHz, 192 GB RAM
Intel Omni-path 200 Gb/sec.

Storage Servers

2x RSC Tornado TDN551 AFS:

2x Intel Xeon Gold 6248R
48 cores @3.00GHz, 376 GB RAM
Intel Omni-path 200 Gb/sec
32x Intel 30,73 TB

Storage capacity: 1,5 PB



Management Server

Supermicro:

2x Intel Xeon E5-2690 v4
56 cores @2.60GHz, 256 GB RAM
Mellanox 200 Gb/sec.

Storage Servers

3x Supermicro:

2x Intel Xeon Gold 6230R
52 cores @2.10GHz, 786 GB RAM
Mellanox 200 Gb/sec
16x Intel 15,36 TB

Storage capacity: 600 TB



<http://hlit.jinr.ru/hybrilit-lustre/>

Software and Information Environment of the Platform

All components of the Platform are integrated into a unified software and information environment, allowing users to utilize available application software packages, develop their own applications, and perform computations using various types of computing architectures (CPU and GPU). The software and information environment can be divided into three levels: system, software, and information.

Information Level

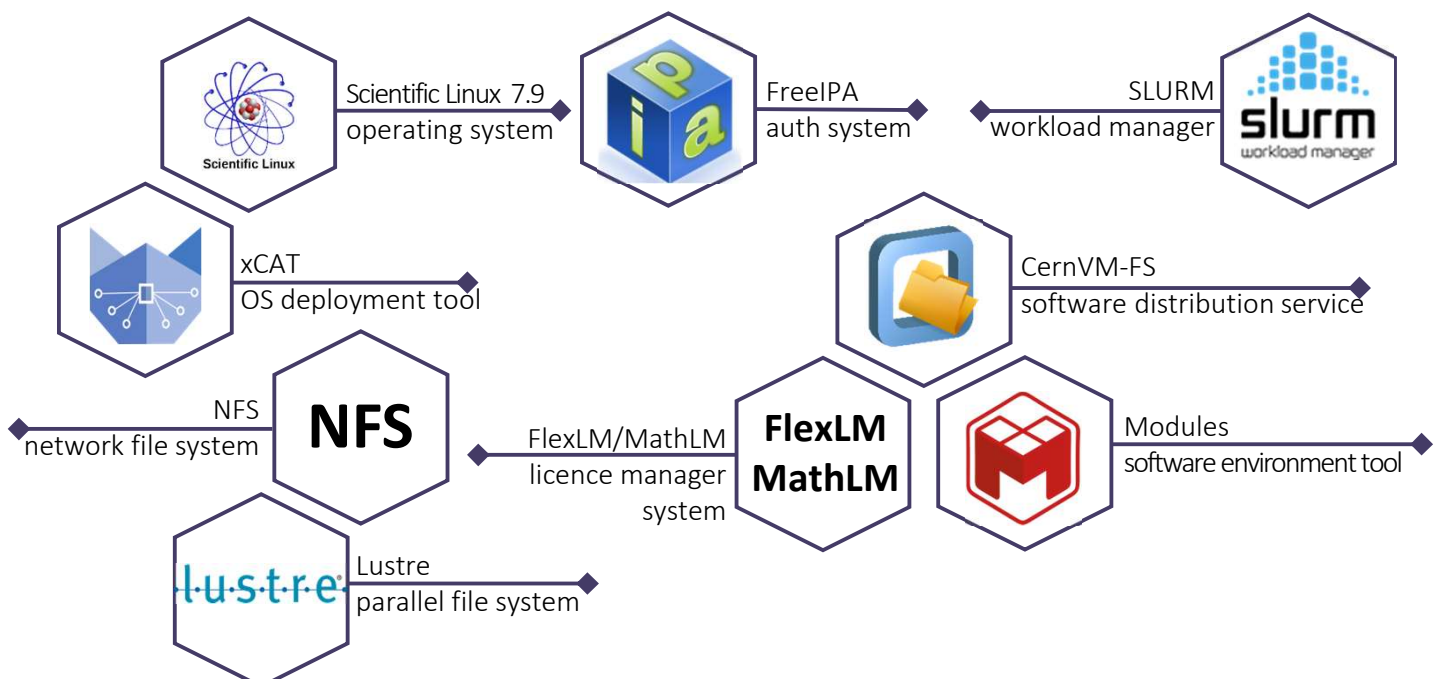
The information level includes various services that facilitate active and close interaction between users and the HybriLIT team, provide access to necessary information, and enable data exchange.

Software Level

The software level contains application software packages designed to support scientific computations, engineering calculations, data processing and analysis, as well as a number of specialized services (ML/DL/HPC ecosystem, Quantum Computing Testbed, HLIT-VDI, MOSTLIT, Medical Data Testbed, etc.).

System Level

The system level contains core software components that ensure the operation of the entire system and monitoring services that allow tracking the Platform's performance and load.



ML/DL/HPC Ecosystem

To support the development of machine learning and deep learning methods and algorithms, as well as data analysis and visualization, the ML/DL/HPC ecosystem has been deployed within the software and information environment of the HybriLIT platform.

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>



All components of the ecosystem are equipped with core libraries and frameworks for machine learning and deep learning tasks (TensorFlow, PyTorch, Keras), enabling users to develop algorithms, train neural network models, and conduct research on both central processing units (CPUs) and graphical processing units (GPUs) within the JupyterLab environment.

http://hlit.jinr.ru/ecosystem-for-ml_dl_bigdataanalysis-tasks/

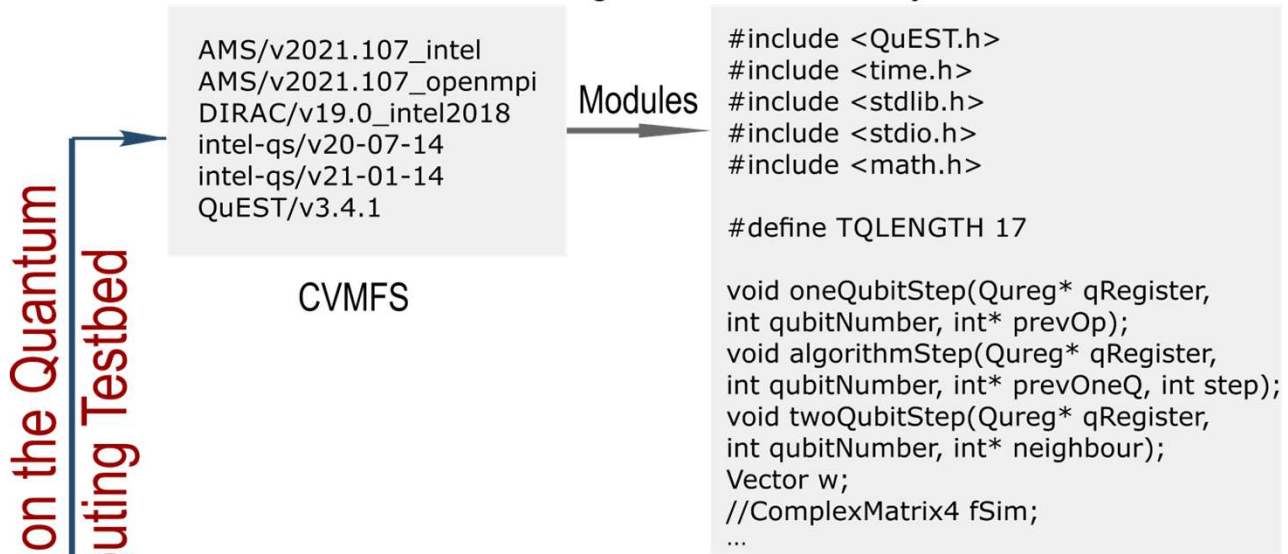
Quantum Computing Testbed

Within the resources of the ML/DL/HPC ecosystem, the Quantum Computing Testbed operates to support tasks related to the development of quantum algorithms and the use of quantum computing simulators.

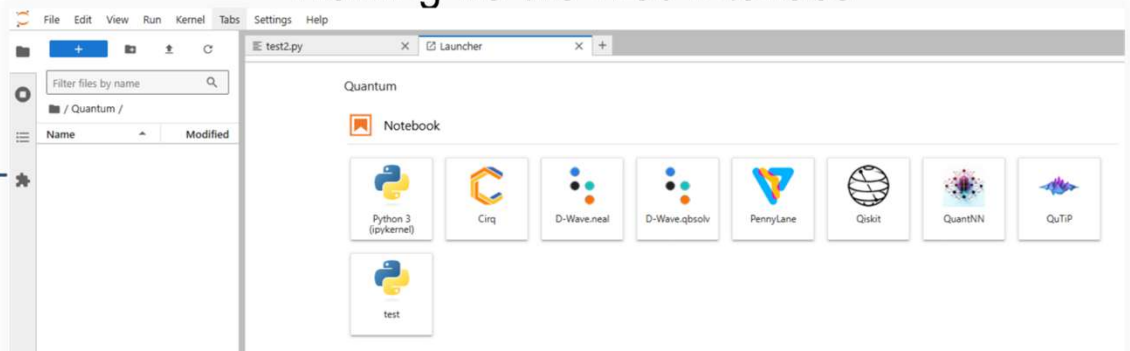
Work on the testbed is organized in two modes:

1. Using a job scheduler (via SLURM queues)
2. In interactive mode through a web browser

Working via the Batch System



Working via the Web Interface



Quantum Simulators:

- Cirq
- Qiskit
- PennyLane
- QuTiP
- DWave.neal
- DWave.qbsolv
- QunatNN

Server Specifications:

[<https://ampere05.jinr.ru>]

2x AMD EPYC 7763 (64 Cores @ 2.45 GHz),
2 TB RAM,
8x NVIDIA Tesla A100 SXM4 80 GB HBM2

<http://hlit.jinr.ru/quantum-polygon/>


MOSTLIT: Service for Detection and Analysis of Radiation-Induced Foci

The web service MOSTLIT is created to automate radiation-induced foci (RIF) analysis. It allows the user to observe the identified cell nuclei in an uploaded fluorescent image, choose the desired nuclei, and automatically get the marked foci.

Upload your image


Browse file

3931.tif




Identify Cells

All Detected Cells




After the user chooses to delete the cell numbers 8 and 11:

Selected Cells




Detect Foci


cell_00_green




cell_00_red




Colocalized_00




cell_01_green



cell_01_red



Colocalized_01



At the end, a table containing numerical characteristics such as the number of RIF per cell, foci area, and colocalizations is provided.

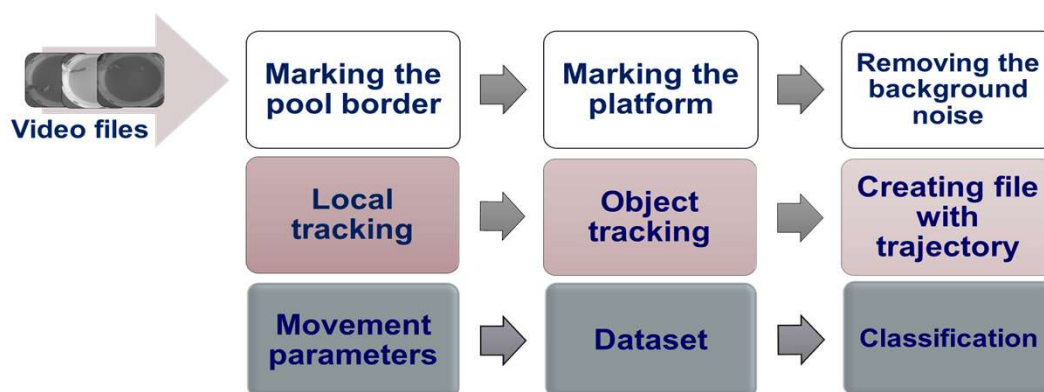
This work is a joint project between MLIT, LRB JINR, and Burnazyan FMBA center.

<https://mostlit.jinr.ru/>



Web Service for Analyzing Trajectories of Small Laboratory Animals in the Morris Water Maze Behavioral Test

The Morris Water Maze behavioral test is aimed at studying spatial memory and learning ability in small laboratory animals. Experiments generate large amounts of video data. Developing convenient tools significantly reduces research time, provides more accurate and comprehensive quantitative data, and captures the trajectory of the experimental animal, which is highly informative for this test.


The web service allows users to extract movement trajectories, obtain necessary analytical movement data, and create a database for subsequent trajectory classification tasks. Users can also verify the constructed trajectories.



Trajectory



name	score
Total time traveled (sec)	61.0
Total distance (cm)	1635.39
FPS	30
Speed (cm/s)	71.49
First frame	26
Last frame	1856



- ✓ An algorithm for trajectory construction has been developed and tested
- ✓ Obtained trajectories have been annotated
- ✓ A dataset for classification has been created

<http://bio-dashboards.jinr.ru/morris>

Medical Data Testbed

Within the resources of the ML/DL/HPC ecosystem of the HybriLIT heterogeneous platform, a Medical Data Testbed operates for working with computed tomography (CT), magnetic resonance imaging (MRI), and functional MRI (fMRI) data in DICOM and NIfTI formats, within the JupyterLab environment. The testbed allows reading and storing data using the installed libraries.



Server Specifications

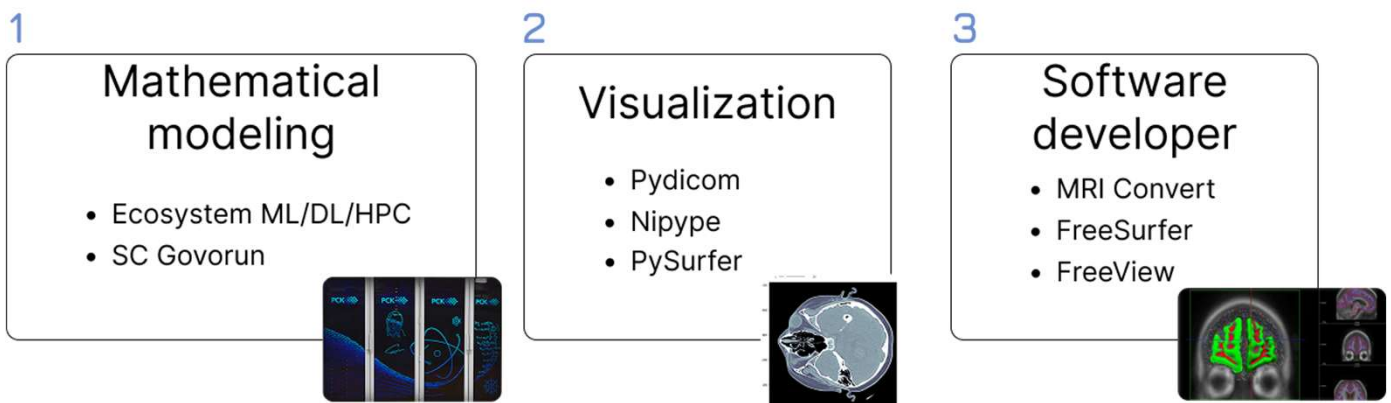
2x Intel Xeon Gold 6126, 48 core @ 2.60 GHz,
256 GB RAM

Using the testbed

Access to the testbed is provided from the JINR network. To gain access to the testbed resources, please send an email to the administrators:

[<aanikina@jinr.ru>](mailto:aanikina@jinr.ru), [<zuevmax@jinr.ru>](mailto:zuevmax@jinr.ru).

For working with the installed Python libraries in the JupyterLab interface, the `ct-mri` environment is available as a separate “button.”



<http://hlit.jinr.ru/polygon-ct-mri/>

HLIT-VDI Service

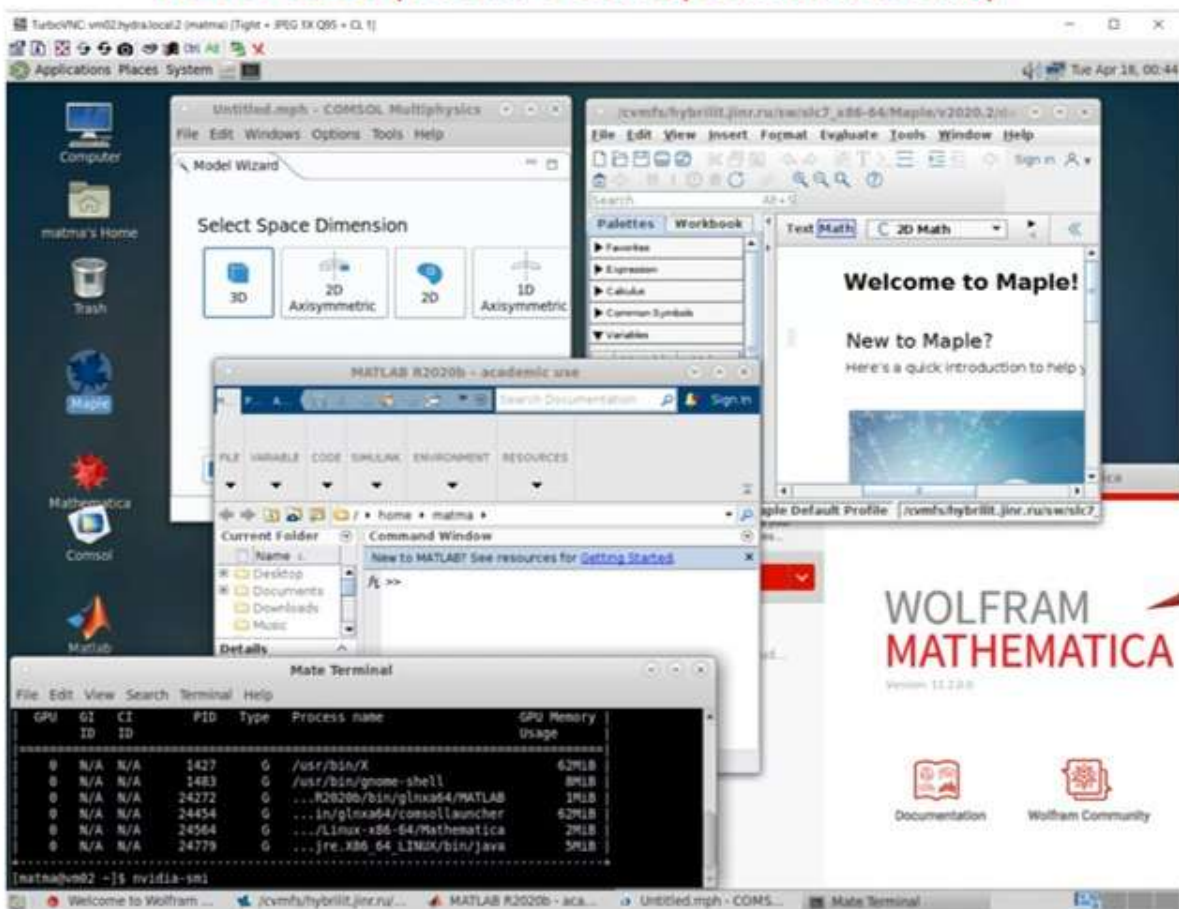
The HLIT-VDI service is designed for working with application software packages that have a graphical user interface in a remote desktop mode using the TurboVNC Viewer client on virtual machines. The service is hosted on a dedicated server with an Nvidia Tesla M60 GPU. Four virtual machines are available.

Virtual Machine Specifications

- 24 GB RAM
- 4x Intel Xeon E5-2697A v4
- Nvidia Tesla M60, 8 GB
- File Systems: ZFS, Lustre
- OS Centos 7.9

The HLIT-VDI service allows users to utilize complex mathematical and physics software packages in their research, such as Comsol, Matlab, Mathematica, Maple, Geant4, and ROOT.

HLIT-VDI (Virtual Desktop Infrastructure)



<http://hlit.jinr.ru/hlit-vdi/>

How to Become a Platform User



For JINR Employees

Steps:

1. Go to the website: hlit.jinr.ru → *For Users* → *Registration*.
2. Select the registration type: *JINR Employees*.
3. Fill out the registration form.
4. Submit the form and wait for confirmation.
5. After approval, you will receive the following via the provided email:
 - Login credentials
 - Instructions for accessing the Platform

For Employees of External Organizations

For external users, the procedure includes additional approval.

Steps:

1. Go to the website: hlit.jinr.ru → *For Users* → *Registration*.
2. Select the option: *User from an External Organization*.
3. Fill out the form.
4. After submission, the form will go through moderation.
5. Upon approval, you will receive:
 - Login credentials
 - Instructions for accessing the Platform

For Access to the “Govorun” Supercomputer

It is additionally necessary to fill out the *Extended Task Description* form (available on the website in the *Registration* section).

http://hlit.jinr.ru/for_users/registration/

How to Connect to the Platform

For Windows Users

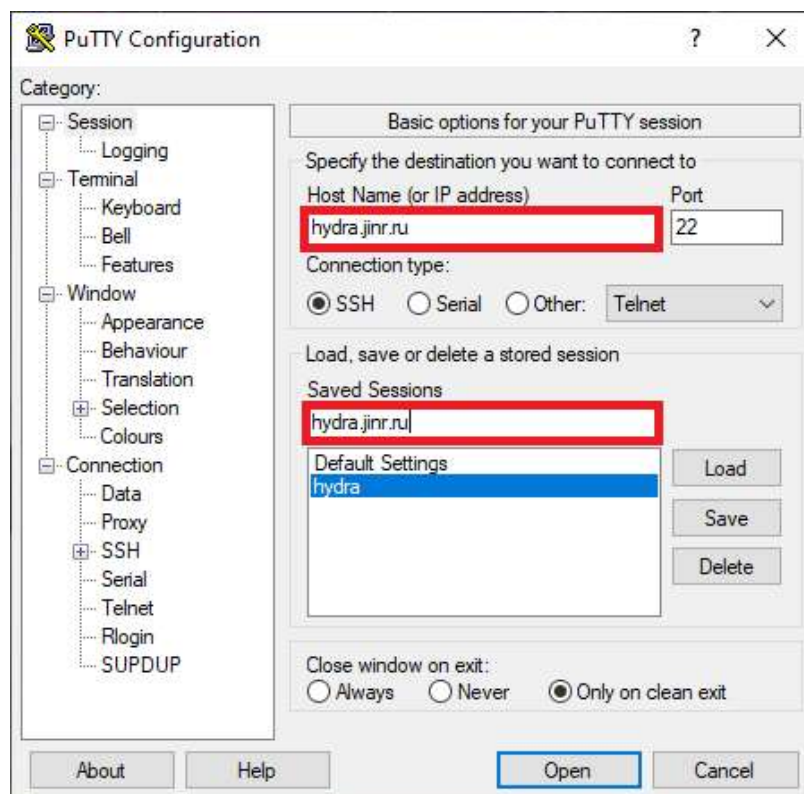
To connect to the Platform from Windows, you need to use a special program — an SSH client, for example, PuTTY.

The download link is available on the website:

http://hlit.jinr.ru/guide_hybrilit/#_2

PuTTY Settings for Platform Access:

- In the *Host Name (or IP address)* field, enter the server address: *hydra.jinr.ru*.
- In the *Saved Sessions* field, enter a name for the connection: *hydra.jinr.ru*.
- To use the remote graphical X11 interface, go to *Connection* → *SSH* → *X11* and check *Enable X11 forwarding*.
- On the *Connection* → *SSH* → *Tunnels* tab, check *Local ports accept connections from other hosts*.
- Return to the *Session* tab and click *Save* to store all settings.
- Click *Open* to connect to the HybriLIT Platform and enter the login and password you received during registration.



For Linux Users

- Open the Terminal.
- Enter the command: `ssh USERNAME@hydra.jinr.ru` (where USERNAME is your login)..
- Enter your password.

http://hlit.jinr.ru/guide_hybrilit/#_2

Working with the SLURM Job Scheduler

1. Log in to the HybriLIT platform.
2. Load the required software packages.
For computations on the “Govorun” supercomputer,
load the module **GVR/v1.0-1**

Basic module commands:

module avail – list available software packages;
module add <module_name> – load a software package;
module list – show loaded modules;
module unload <module_name> – unload a software package;
module clear – unload all modules.

3. Prepare a SLURM job script with the required parameters.

Main parameters:

#SBATCH -p <partition_name> – queue/partition to use;
#SBATCH -n 10 – number of requested CPU cores;
#SBATCH -t 60 – requested runtime in minutes;
#SBATCH --gres=gpu:2 – number of requested GPUs;
#SBATCH --mem=1024 – requested RAM in megabytes;
#SBATCH -N 2 – number of requested nodes.

4. Submit the job to the queue with the command:

sbatch <name of the script>


Basic SLURM commands:

sinfo – show the state of queues and nodes;
sinfo -Nle – show detailed state and characteristics of queues and nodes;
squeue – show status of running jobs;
sbatch <name of the script> – submit a job to the queue;
scancel <job_id> – remove a job from the queue.

5. After the job completes, a file named **slurm-<job_id>.out** appears in the working directory. It contains the standard output of the job (program log and error messages if the job ended unsuccessfully).

http://hlit.jinr.ru/guide_hybrilit/#_4

User Support



The **HybriLIT website** provides a description of the Platform's computing component, methodological materials, and instructions for working with various programming technologies.

http://hlit.jinr.ru/guide_new-2/

The **Indico system** is used for organizing conferences, seminars, and meetings.

<https://indico.jinr.ru/>

The **Project Management HybriLIT User Support** service is used for communication between users and the HybriLIT support team and is intended for resolving issues that arise while working on the Platform.

<https://pm.jinr.ru/projects/hybrilit-user-support>


The **HybriLIT User Support Telegram channel** is used to keep users informed.

<https://t.me/hybrilit>

The **GitLab service** is intended for collaborative software development.


<https://git.jinr.ru/>

HybriLIT Platform in Numbers



The results of research conducted using the supercomputer resources since 2018 have been presented in **509** publications.

In 2024–2025, an additional **110** articles were published.

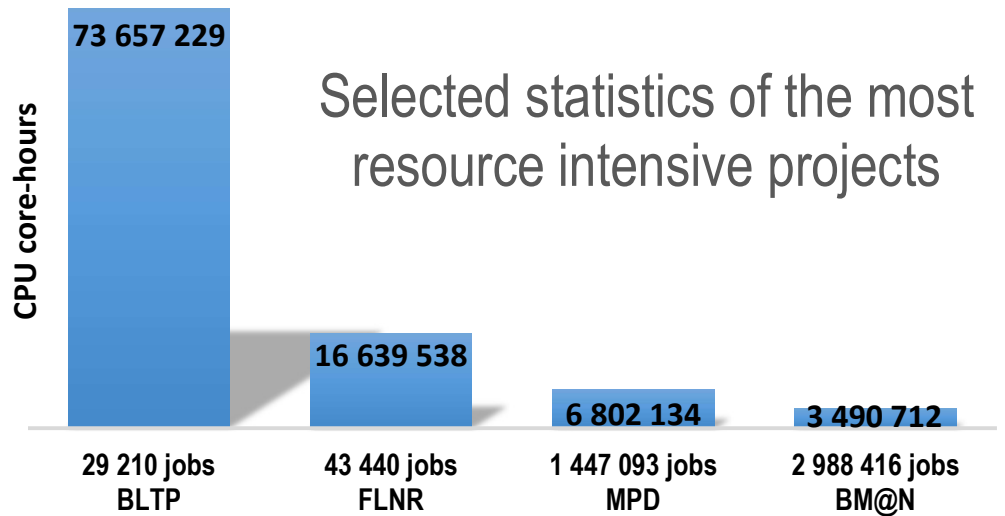


Total jobs

13,94M


Total core hours

111,87M



Selected statistics of the most resource intensive projects

GPU usage at the Heterogeneous HybriLIT platform





MLIT JINR, Dubna
2025

