DIRAC System as a Mediator Between Hybrid Resources and Data Intensive Domains

Vladimir Korenkov[1,3], Igor Pelevanyuk[1,3], and Andrei Tsaregorodtsev[2,3]

[1] Joint Institute for Nuclear Research, Dubna, Russia
korenkov@jinr.ru, pelevanyuk@jinr.ru
[2] CPPM, Aix-Marseille University, CNRS/IN2P3, Marseille, France
atsareg@in2p3.fr
[3] Plekhanov Russian Economics University, Moscow

Abstract. Data and computing-intensive applications in scientific research are becoming more and more common. And, since different computing solutions have different protocols and architectures, they should be chosen wisely during the design stage. In a modern world of diverse computing resources such as grids, clouds, and supercomputers the choice can be difficult. Software developed for integration of various computing and storage resources into a single infrastructure, the so-called interware, is intended to facilitate this choice. The DIRAC interware is one of these products. It proved to be an effective solution for many experiments in High Energy Physics and some other areas of science. The DIRAC interware was deployed in the Joint Institute for Nuclear Research to serve the needs of different scientific groups by providing a single interface to a variety of computing resources: grid cluster, computing cloud, supercomputer Govorun, disk, and tape storage systems. The DIRAC based solution was proposed for the Baryonic Matter at Nuclotron experiment which is in operation now as well as for the future experiment Multi-Purpose Detector on the Nuclotron-based Ion Collider fAcility. Both experiments have requirements making the use of heterogeneous computing resources necessary.

Keywords: Grid computing, Hybrid distributed computing systems, Supercomputers, DIRAC

1 Introduction

Data intensive applications became now an essential mean for getting insights of new scientific phenomena while analyzing huge data volumes collected by modern experimental setups. For example, the data recording to the tape system at CERN exceeded in total 10 Petabytes per month in 2018 for all the 4 LHC experiments. In 2021, with the start of the Run 3 phase of the LHC program, the experiments will resume data taking with considerably increased rates. The needs for the computing and storage capacity of the LHCb experiment, for instance, will increase by an order of magnitude [1].

In a more distant future, with the start of the LHC Run 4 phase, the projected data storage needs of the experiments are estimated to exceed 10 Exabyte's. These are unprecedented volumes of data to be processed by distributed computing systems, which are being adapted now to cope with the new requirements.

Other scientific domains are quickly approaching the same collected data volumes: astronomy, brain research, genomics and proteomics, material science [3]. For example, the SKA large radio astronomy experiment [2] is planned to produce about 3 Petabytes of data daily when it will come into full operation in 2023.

The needs of the LHC experiments in data processing were satisfied by the infrastructure of the World LHCb Computing Grid (WLCG). The infrastructure still delivers the majority of computing and storage resources for these experiments. It is well suited for processing the LHC data ensuring massively parallel data treatment in a High Throughput Computing (HTC) paradigm. WLCG succeeded in putting together hundreds of computing centers of different sizes but with similar properties, typically providing clusters of commodity processors under control of one of the batch systems, e.g. LSF, Torque or HTCondor. However, new data analysis algorithms necessary for the upcoming data challenges require new level of parallelism and new types of computing resources. These resources are provided, in particular, by supercomputers or High Performance Computing (HPC) centers. The number of HPC centers is increasing and there is a clear need in setting up infrastructures allowing scientific communities to access multiple HPC centers in a uniform way as it is done in the grid systems.

Another trend in massive computing consists in provisioning resources via cloud interfaces. Both private and commercial clouds are available now to scientific communities. However, the diversity of interfaces and usage policies makes it difficult to use multiple clouds for applications of a particular community. Therefore, providing uniform access to resources of various cloud providers would increase flexibility and the total amount of available computing capacity for a given scientific collaboration.

Large scientific collaborations typically include multiple participating institutes and laboratories. Some of the participants have considerable computing and storage capacity that they can share with the rest of the collaboration. With the grid systems this can be achieved by installing complex software, the so-called grid middleware, and running standard services like Computing and Storage Elements. For managers of local computing resources who are usually not experts in the grid middleware, this represents a huge complication and often results in underused resources that would otherwise be beneficial for the large collaborations. Tools for easy incorporation of such resources can considerably increase the efficiency of their usage.

The DIRAC Interware project is providing a framework for building distributed computing systems using resources of all different types mentioned above and putting minimal requirements on the software and services that should be operated by the resources providers. Developed originally for the LHCb experiment at LHC, CERN, the DIRAC interware was generalized to be applicable for a wide range of applications. It can be used to build independent distributed computing infrastructures as well as to provide services for existing projects. DIRAC is used by a number of High Energy Physics and Astrophysics experiments but it is also providing services for a number of general-purpose grid infrastructures, for example, national grids in France

[4] and Great Britain [5]. The EGI Workload Manager is the DIRAC service provided as part of the European Grid Infrastructure service catalog. It is one of the services of the European Open Science Cloud (EOSC) project inaugurated in the end of 2018 [6]. The EGI Workload Manager provides access to grid and cloud resources of the EGI infrastructure for over 500 registered users.

In this paper we describe the DIRAC based infrastructure deployed at the Joint Institute for Nuclear Research, Dubna, putting together a number of local computing clusters as well as connecting cloud resources from JINR member institutions.

2 DIRAC Interware

The DIRAC Interware project provides a development framework and a large number of ready-to-use components to build distributed computing systems of arbitrary complexity. DIRAC services ensure integration of computing and storage resources of different types and provide all the necessary tools for managing user tasks and data in distributed environments [7]. Managing both workloads and data within the same framework increases the efficiency of data processing systems of large user communities while minimizing the effort for maintenance and operation of the complete infrastructure.

The DIRAC software is constantly evolving to follow changes in the technology and interfaces of available computing and storage resources. As a result, most of existing HTC, HPC and cloud resources can be interconnected with the DIRAC Interware. In order to meet the needs of large scientific communities, the computing systems should fulfill several requirements. In particular, it should be easy to describe, execute and monitor complex workflows in a secure way respecting predefined policies of usage of common resources.

2.1 Massive Operations

Usual workflows of large scientific collaborations consist and creation and execution of large numbers of similar computational and data management tasks. DIRAC is providing support for massive operations with its Transformation System. The system allows definition of Transformations – recipes to create certain operations triggered by the availability of data with required properties. Operations can be of any type: submission of jobs to computing resources, data replication or removal, etc. Each Transformation consumes some data and derives ("transforms") new data, which, in turn, can form input for another Transformation. Therefore, Transformations can be chained creating data driven workflows of any complexity. Data production pipelines of large scientific communities based on DIRAC are using heavily the Transformation System defining many hundreds of different Transformations. Therefore, each large project developed its own system to manage large workflows each consisting of many Transformations. There was a clear need to simplify the task of managing complex workflows for the new communities. In order to do that a new system was introduced in DIRAC – Production System. The new system is based on the experience of several community specific workflow management systems and provides a uniform way to create a set of Transformations interconnected via their input/output data filters. It helps production managers to monitor the execution of so created workflows, evaluate the overall progress of the workflow advancement and validate the results with an automated verification of all the elementary tasks.

2.2 Multi-community Services

In most of the currently existing multi-community grid infrastructures the security of all operations is based on the X509 PKI infrastructure. In this solution, each user has to, first, obtain a security certificate from one of Certification Authorities (CA) recognized by the infrastructure. The certificate should be then registered in a service holding a registry of all the users of a given Virtual Organization (VO). The user registry keeps the identity information together with associated rights of a given user. In order to access grid resources, users are generating proxy certificates which can be delegated to grid remote services in order to perform operation on the user's behalf.

The X509 standard based security is well supported in academia institutions but is not well suited for other researchers, for example, working in universities. On the other hand, there are well-established industry standards developed mostly for the web applications that allow identification of users as well as delegation of user rights to remote application servers. Therefore, grid projects started migration to the new security infrastructure based on the OAuth2/OIDC technology. With this technology, user's registration is done by local identity providers, for example, a university LDAP index. On the grid level a Single-Sign-On (SSO) solution is provided by federation of multiple identity providers to ensure mutual recognition of user security tokens. In particular, the EGI infrastructure has come up with the Check-In SSO service as a federated user identity provider.

The DIRAC user management subsystem was recently updated in order to support this technology. Users can be identified and registered in DIRAC based on their Check-In SSO tokens which contain also additional user metadata, e.g. membership in VOs, user roles and rights. This metadata are used to define user membership in the DIRAC groups, which define user rights within the DIRAC framework. This allows managing community policies, such as resources access rights and usage priorities that will be applied by DIRAC to the user payloads. The DIRAC implementation of the new security framework is generic and can be easily configured to work with other SSO systems.

2.3 DIRAC Software Evolution

The intensity of usage of the DIRAC services is increasing and the software must evolve to cope with the new requirements. This process is mostly driven by the needs of the LHCb experiment, which remains the main consumer and developer of the DIRAC software. As was mentioned above, the order of magnitude increase in the data acquisition rate of LHCb in 2021 dictates revision of the technologies used in its data processing solutions.

Several new technologies were introduced recently into the DIRAC software stack. The use of Message Queue (MQ) services allows passing messages between distributed DIRAC components in an asynchronous way with the possibility of message buffering in case of system congestions. The STOMP message passing protocol is used and all the MQ service supporting this protocol can be used, e.g. ActiveMQ, RabbitMQ and others. The MQ mechanism for the DIRAC component communications is considered to be complementary to the base Service Oriented Architecture (SOA) employed by DIRAC. This solution increases the overall system scalability and resilience.

The DIRAC services states are kept in relational databases using MySQL servers. The MySQL databases have shown very stable operation over the years of usage. However, the increased amount of data to be stored in databases limits the efficiency of queries and new solutions are necessary. The so-called NoSQL databases have excellent scalability properties and can help in increasing the efficiency of the DIRAC components. The ElasticSearch NoSQL (ES) database solution was applied in several DIRAC subsystems. In particular, the Monitoring System, which is used to monitor the current consumption of the computing resources, was migrated to the use of the ES based solution. This information is essential in implementation of the priority policies based on the history of the resources consumption to ensure fair sharing of the common community resources.

This and other additions and improvements in the DIRAC software aim at the overall increase of the system efficiency and scalability to meet requirements of multiple scientific communities relying on DIRAC services for their computing projects.

3 JINR DIRAC Installation

The Joint Institute for Nuclear Research is an international intergovernmental organization, a world-famous scientific center that is a unique example of the integration of fundamental theoretical and experimental research. It consists of seven laboratories: Laboratory of High Energy Physics, Laboratory of Nuclear Problems, Laboratory of Theoretical Physics, Laboratory of Neutron Physics, Laboratory of Nuclear Reactions, Laboratory of Information Technologies, Laboratory of Radiation Biology. Each laboratory being comparable with a large institute in the scale and scope of investigations performed.

JINR has powerful high-productive computing environment that is integrated into the world computer network through high-speed communication channels. The basis of the computer infrastructure of the Institute is the Multifunctional Information Computer Complex (MICC). It consists of several large components: grid cluster, computing cloud, supercomputer Govorun. Each component has its features, advantages, and disadvantages. Different access procedures, different configuration and connection with different storage systems do not allow simple usage of all of them together for one set of tasks.

3.1 Computing Resources

Grid cluster. The JINR grid infrastructure is represented by the Tier1 center for the CMS experiment at the LHC and the Tier2 center.

After the recent upgrade, the data processing system at the JINR CMS Tier1 consists of 415 64-bit nodes: 2 x CPU, 6–16 cores/CPU that form altogether 9200 cores for batch processing [8]. The Torque 4.2.10/Maui 3.3.2 software (custom build) is used as a resource manager and a task scheduler. The computing resources of the Tier2 center consist of 4,128 cores. The Tier2 center at JINR provides data processing for all four experiments at the LHC (Alice, ATLAS, CMS, LHCb) and apart from that supports many virtual organizations (VO) that are not members of the LHC (BES, BIOMED, COMPASS, MPD, NOvA, STAR, ILC).

Grid cluster is an example of a High-Throughput Computing paradigm. It means that the primary task of this cluster is to run thousands of independent processes at the same time. Independent means that once a process has started and until it finishes, the process does not rely on any input that is being produced at the same moment by other processes.

Jobs may be sent to the grid using CREAM Computing Element – service installed in JINR specifically for grid jobs. Computing element works as an interface to the local batch farm. Its primary task is to authenticate the owner of the job and redirect it to the right queue. For the users, it is required to have X509 certificate and be a member of Virtual Organization supported by the Computing Element.

Cloud infrastructure. The JINR Cloud [9] is based on an open-source platform for managing heterogeneous distributed data center infrastructures – OpenNebula 5.4. The JINR cloud resources were increased up to 1564 CPU cores and 8.1 TB of RAM in total. Cloud infrastructure is used primarily for two purposes: to create personal virtual machines and to create virtual machines to serve as worker nodes for jobs. We are going to focus on the second purpose.

The biggest advantage of cloud resources as computing capacity is their flexibility. In case of grid or batch resources, several jobs working on one worker node share between them: operating system, CPU cores, RAM, HDD/SSD storage, disk Input/Output capabilities, and network bandwidth. If a job needs more disk space or RAM it is not straightforward to submit the job to the grid without the help of administrators, who in most cases have to create a dedicated queue for this particular kind of jobs. In the case of clouds, it is much easier to provide a specific resource that the job requires. It may be a virtual machine with a large disk, specific operating system, required number of CPU cores, RAM capacity and network.

When a job destinated to the cloud enters the system the corresponding virtual machine is created by DIRAC using the OpenNebula API. During the contextualization process, the DIRAC Pilot is installed in the VM and configured to receive jobs for this cloud resource. Once the job is finished, the pilot attempts to get the next job. If there are no more jobs for the cloud, the pilot will request the VM shutdown. The pilot in the cloud environment is not limited by the time and may work for weeks. These features make cloud resources perfect for specific tasks with unusual requirements.

Govorun supercomputer. The Supercomputer Govorun was put into production in March 2018[10]. It is a heterogeneous platform built on several processors' technologies: GPU part and two CPU parts. GPU part unites 5 servers DGX-1. Each server consists of 8 NVIDIA Tesla V100 processors. The CPU part is a high dense liquid-cooled system. Two types of processors are used inside: Intel Xeon Phi 7290(21 servers) and Intel Xeon Gold 6154(40 servers). The total performance of all the three parts is 1 PFlops for operations with single precision and 0.5 PFlops for double precision. SLURM 14.11.6 is used as the local workload manager. Three partitions were created to subdivide tasks in the supercomputer: gpu, cpu, phi.

The supercomputer is used for tasks, which require massive parallel computations. For example: to solve problems of lattice quantum chromodynamics for studying the properties of hadronic matter with high energy density and baryon charge and in presence of strong electromagnetic fields, mathematical modeling of the antiproton-proton and antiproton-nucleus collisions with the use of different generators. It is also used for simulation of collision dynamics of relativistic heavy ions for the future MPD experiment on the NICA collider.

Right now, the supercomputer utilizes its own authentication and authorization system. Every user of the supercomputer should be registered and allowed to send jobs. Sometimes a part of the supercomputer is free from parallel tasks and may be used as a standard batch system. Special user was created for DIRAC. All jobs sent to the Govorun are executed with this user identity. This frees actual users from additional registration procedures.

3.2 Storage Resources

EOS storage on disks. EOS [11] is a multi-protocol disk-only storage system developed at CERN since 2010 to store physics analysis data physics experiments (including the LHC experiments). Having a highly-scalable hierarchical namespace, and with the data access possible by the XROOT protocol, it was initially used for physics data storage. Today, EOS provides storage for both physics and user use cases. For the user authentication, EOS supports Kerberos (for local access) and X.509 certificates for grid access. To ease experiment workflow integration, SRM as well as GridFTP access is provided. EOS supports the XROOT third-party copy mechanism from/to other XROOT enabled storage services.

The EOS was successfully integrated into the MICC structure. The NICA experiments already use EOS for data storage. At the moment there are ~200TB of "raw" BM@N data and ~84GB of simulated MPD data stored in the EOS instance. EOS is visible as a local file system on the MICC worker nodes. It allows users authorized by the Kerberos5 protocol to read and write data. A dedicated service was installed to allow usage of X509 certificates with VOMS extensions.

dCache disk and tape storage. The core part of the dCache has been proven to efficiently combine heterogeneous disk storage systems of the order of several hundreds TBs and present its data repository as a single filesystem tree. It takes care of data, failing hardware and makes sure, if configured, that at least a minimum number of copies of each dataset resides within the system to ensure high data availability in case of disk server maintenance or failure. Furthermore, dCache supports a large set of standard access protocols to the data repository and its namespace. It supports DCAP, SRM, GridFTP, and xRootD [12].

dCache at JINR consists of two parts: disk storage and tape storage. The disk part operations are similar to EOS. The tape works through the dedicated disk buffer servers. When data are uploaded to the dCache tape part, they are first uploaded to the disk buffer. If the disk buffer is occupied above a certain threshold (which is 80% in our case), all the data is moved from disk to tape and removed from the disk buffer. While data stay in the buffer, access to them is similar to access to the dCache disk data. But once the data are moved to tape and removed from the disk, access to them may require time. The time required to select the right tape and transfer data from tape to the disk depends on the tape library task queue. Generally, the time varies from 20 seconds up to several minutes.

Tape library should be used only for archive storage and preferably for big files. Otherwise, it may bring unnecessary load on the tape library. It is much easier to write many small files to the tape than to read it back.

Ceph storage. Software-defined storage (SDS) based on the Ceph technology is one of the key components of the JINR cloud infrastructure. It runs in production mode since the end of 2017. It delivers object, block and file storage in one unified system. Currently, the total amount of raw disk space in that SDS is about 1 PB. Due to triple replication, effective disk space available for users is about 330 TB. Users of Ceph can attach part of the storage to a computer using the FUSE disk mounting mechanism. After that, it is possible to read and write data to the remote storage as if it is connected directly to the computer.

The Ceph storage was integrated into DIRAC installation for tests. Since Ceph does not allow authentication by X509 certificates with VOMS extensions, a dedicated virtual machine was configured to host DIRAC Storage Element – a standard service which works as a proxy to a file system. It checks certificates with VOMS extensions before allowing writing and reading to a dedicated directory. Right now, Ceph storage does not allow massive transfers since it relies on one server with Ceph attached by FUSE. The test demonstrated that the maximum speed of transfer is not exceeding 100 MB/s which is a consequence of 1Gb network connection. The way to increase the performance of this storage is an improvement of the network speed up to 10Gb/s and a possible creation of additional DIRAC Storage Elements which can share the load between themselves.

Performance test of EOS and disk dCache. In the case of massive data processing, it is crucial to know the limitation of different components. The limitations may depend on the usage of resources. In many use-cases it is crucial to transfer some amount of data first, so we decided to test storage elements. The synthetic test was proposed: run many jobs on one computing resource, make them start download of all the same data at the same moment, measure how much time it takes to get the file.

Every test job had to go through the following steps:

1. Start execution on the worker node.
2. Check the transfer start time.
3. Wait until the transfer time moment.

4. Start the transfer.

5. When the transfer is done, report the information about the duration of the transfer.

6. Remove the downloaded file.

Two storage systems were chosen for the tests: EOS and dCache since only they are accessible for read and write on all the computing resources right now. We chose the test file size to be 3 GB. The amount of test jobs in one test campaign depends on the number of free CPU cores in our infrastructure. We initiate 200 jobs during one test campaign. Not all of them could start at the same time, which means that during the test less than 200 jobs may download data. This is taken into consideration when we calculate total transfer speed.

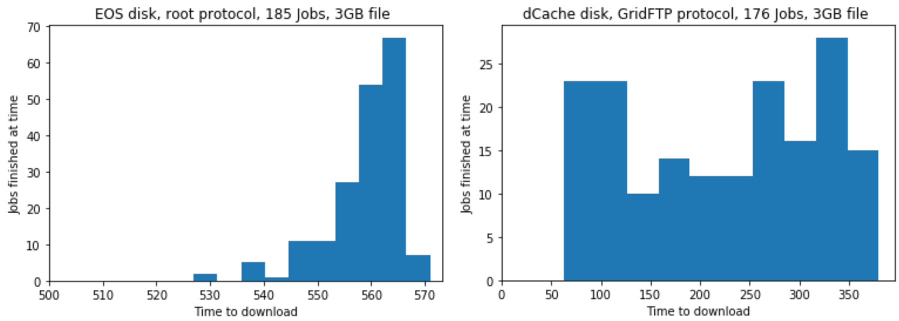

Fig. 4. Number of transfers finished at the time

Several test campaigns were performed to evaluate variance between the tests, but all of them showed similar results after all. Two representative examples were chosen to demonstrate the rates (see Fig. 1). To calculate transfer speed the following formula was used:

$$Transfer\ speed = \frac{Total\ data\ transferred}{Longest\ transfer\ duration}.$$

This formula allows calculation of the worst transfer speed of individual file during the test campaign. For EOS calculated transfer speed was 990 GB/s on 200 jobs and for dCache it was 1390 GB/s in 176 jobs. It should be mentioned that all the tests were performed on a working infrastructure, so some minor interference may be caused by other activities. On the other hand, demonstrated plots represent real transfers performed under normal conditions.

The numbers described above demonstrate that the computing and storage infrastructure at JINR is quite extensive and diverse. Nowadays, different components are used directly for different tasks. So, the workflows are bound to dedicated resources and switching between them would not be an easy task at least. Sometimes different components could be separated by a slower network, different authentication systems and different protocols. This problem becomes visible when one of the resources is overloaded while others are underloaded. In the case of good interoperability between the components, it would be possible to easily switch between them.

Of course, the resources are not fully available for all the tasks. They have to provide pledges for different tasks and experiments, but still, they could be underloaded. And since there are tasks that should not be necessarily bound to particular resources, it would be beneficial to have a mechanism to use them in some uniform way by scientific groups.

So, to improve the usage efficiency of all the resources, to provide a uniform way to store, access and process data, the DIRAC system was installed and evaluated.

4 JINR DIRAC Installation

The DIRAC installation in JINR consists of 4 virtual machines. Three of them placed on a dedicated server to avoid network and disk I/O interference with other virtual machines. The operating system on these virtual machines is CentOS 7. It appeared that some of the LCG software related to grid job submission is not compatible with CentOS. To cope with that, we created a new virtual machine with Scientific Linux 6 installed there. Flexibility of the DIRAC modular architecture allowed us to do so. The characteristics of the virtual machines hosting DIRAC services are presented in Table 1.

Table 7. Virtual machines hosting DIRAC services

	dirac-services	dirac-conf	dirac-web	dirac-sl6
OS	CentOS	CentOS	CentOS	Scientific Linux
Version	7.5	7.5	7.5	6.10
Cores	8	4	4	2
RAM	16 GB	8 GB	8 GB	2 GB

4.1 Use Cases in JINR

Up to now, we foresee two big possible use cases: Monte-Carlo generation for Multi Purpose Detector (MPD) at NICA and data reconstruction for Baryonic Matter at Nuclotron(BM@N).

Raw data was received by the BM@N detector and uploaded to the EOS storage. There are two data taking runs available now: run 6 and run 7. The data sizes are respectively: 16 TB and 196 TB. All data consists of files, for run 6 it is roughly 800 files and for run 7 it is around 2200 files. The main difficulty with these files is the fact that their sizes are very different: from several MBs up to 800 GBs per one file. This makes data processing a tough task especially on resources without small amount of local storage or bad network connection. The data could be processed using the Govorun Supercomputer, but the EOS is currently not connected to the storage. And the data may require full reprocessing one day, if the reconstruction algorithms will be changed.

So far, the best would be to process big files in the cloud, other files in the grid infrastructure and sometimes, when the supercomputer has free job slots, do some processing there. But without some central Workload Management system and Data Management system this is a difficult task. The data could be placed not only in EOS but also in dCache. This would allow data delivery to the worker nodes using grid protocols like SRM or xRootD. Once the X509 certificates start working for the EOS storage, it will also be included in the infrastructure and be accessible from everywhere.

The second use case is Monte-Carlo generation for the MPD experiment. Monte-Carlo generation could be performed almost on all the components of MICC at JINR. It is a CPU intensive task less demanding in terms of disk size and input/output rates. The file size could be tuned to be in a particular range for the convenience of the future use. The use of a central distributed computing system may not be critical right now, but it will definitely be useful later, when the real data arrive. It would allow for design and testing of the production workflows, and allow different organizations to participate in the experiment.

5 Conclusion

Joint Institute for Nuclear Research is a large organization with several big computing and storage subsystems. Most of the time they are used by particular scientific groups and there is no simple way to reorganize the load throughout the whole computing center. But with the new big tasks and with the improvement in technologies it became easier to integrate computing resources and use them as a single meta-computer. This leads to improvements in terms of efficiency of usage of the computing infrastructures.

The DIRAC Interware is a good example of a product for building distributed computing systems. It covers most of the needs in workload and storage management. Putting DIRAC services into operation at JINR allowed organization of data processing not in terms of tasks, but in terms of workflows. It also provides tools for removing barriers between the heterogeneous computing and storage resources.

DIRAC services were installed at JINR in order to integrate resources used by big experiments like MPD and BM@N. The initial tests and measurements demonstrated the possibility to use it for data reconstruction and Monte-Carlo generation on all the resources: JINR grid cluster, Computing Cloud and Govorun supercomputer.

References

1. Bozzi, C. and Roiser, S.: The LHCb software and computing upgrade for Run 3: opportunities and challenges, 2017 J. Phys.: Conf. Ser. **898** 112002; doi: doi:10.1088/1742-6596/898/10/112002
2. SKA telescope. https://www.skatelescope.org/software-and-computing/, last accessed 2019/08/19
3. Kalinichenko, L. et al.: Data access challenges for data intensive research in Russia, Informatics and Applications **10** (1), 2–22 (2016); doi: 10.14357/19922264160101
4. France Grilles. http://www.france-grilles.fr, last accessed 2019/08/19
5. Britton, D. et al.: GridPP: the UK grid for particle physics, Phil. Trans. R. Soc. A **367**, 2447–2457 (2009).
6. European Open Science Cloud. https://www.eosc-portal.eu, last accessed 2019/08/19
7. Gergel, V., Korenkov, V., Pelevanyuk, I., Sapunov, M., Tsaregorodtsev, A., and Zrelov, P.: Hybrid Distributed Computing Service Based on the DIRAC Interware, Communications in Computer and Information Science **706**, 105–118 (2017). doi: https://doi.org/10.1007/978-3-319-57135-5_8
8. Baginyan, A. et al.: The CMS Tier1 at JINR: five years of operations, Proceedings of VIII International Conference “Distributed Computing and Grid-technologies in Science and Education” **2267**, 1–10 (2018).
9. Baranov, A. et al: New features of the JINR cloud, Proceedings of VIII International Conference “Distributed Computing and Grid-technologies in Science and Education” **2267**, 257–261 (2018).
10. Adam, Gh. et al.: IT-ecosystem of the HybriLIT heterogeneous platform for high-performance computing and training of IT-specialists, Proceedings of VIII International Conference “Distributed Computing and Grid-technologies in Science and Education” **2267**, 638–644 (2018).
11. Peters, A.J. et al.: EOS as the present and future solution for data storage at CERN 2015, J. Phys.: Conf. Ser. **664** 042042, doi: doi:10.1088/1742-6596/664/4/042042
12. dCache, the Overview, https://www.dcache.org/manuals/dcache-whitepaper-light.pdf, last accessed 2019/08/19