Data Services at IICT HPC Systems

Emanouil Atanassov[1, 2 [0000-0002-7442-7096]], Aneta Karaivanova[1[0000-0002-6493-7981]], Alexander Kirilov[1[0009-0004-3338-5813]], Svetlozar Yordanov[1[0000-0001-7671-6804]]

[1] Institute of Information and Communication Technologies, Bulgarian Academy of Sciences, Sofia, Bulgaria
[2] Centre of Excellence in Informatics and Information and Communication Technologies, Bulgaria
{emanouil,anet,svetlozar}@parallel.bas.bg, alex.kirilov@acad.bg

Abstract. IICT operates a state-of-the-art computational infrastructure, developed over many years of continuous expansion. The latest acquisition has been the HEMUS supercomputer, inaugurated in November 2023. The infrastructure serves mainly diverse scientific communities with widely varying requirements for computational power and data storage and processing, which motivates a feature-rich organization of the data storage and analysis facilities. The main design goal has been to provide efficient access to the data, corresponding to the needs of the applications and workflows. Apart from that, the needs for ensuring regular backups and secure long term storage, taking into account the Open Science paradigm and FAIR principles are also considered. In this paper we outline our key technical solutions, the organization of the various data services at our datacenter as well as our near-term roadmap for deployment of new services.

Keywords: HPC, Supercomputer, Data Services, Deployment, Accounting System.

1 Introduction

Advanced data services refer to techniques and technologies used to manage, process and analyse complex and large datasets (Sterling et al., 2017). Advanced data services are needed for real-time or near-real-time data processing, predictive analytics, machine learning training and inference, 3d visualization, etc. Such services are needed for effective digital transformation in industry as well as in cutting edge science. IICT operates large and complex IT infrastructure, including two supercomputers (Atanassov et al., 2024), built on the basis of state-of-the-art hardware and software technological solutions under the Centre of Excellence in Informatics and ICT project (established under the Grant No BG05M2OP001-1.001-0003, financed by the Science and Education for Smart Growth Operational Program and co-financed by the European Union through the European Structural and Investment funds). The infrastructure has 4 levels with a harmonized design and mutually complementary functionality. At the ground level we have a laboratory for 3D digitization and microstructural analysis, with capacity for industrial computed tomography with two X-ray tubes with nano- and micro-focus; 3D laser scanning systems; dynamic analysis system, etc., as well as distributed data preparation and visualization laboratories equipped with powerful workstations. At the second level we have a data centre with a data storage and analysis system capable of storing and processing efficiently large volumes of data with a raw capacity of 6.72 petabytes. At the top level we have the supercomputer HEMUS with petascale performance, aided by the older Avitohol supercomputer. This infrastructure provides full data lifecycle support for users including data storage, archiving, backup, collaborative access, domain specific interfaces, data annotation and metadata management, etc.

Our approach to deployment of data services follows a 3-step process, where the first step is an initial deployment and internal testing and validation, the second step is reconfiguration and enhancement of capabilities where advisable or cutting off problematic features and the third step is opening the services for widespread use.

This paper presents the initial Data Service Portfolio of the Centre of Excellence in Informatics and ICT (in short, CI&ICT) (CoE in Informatics and ICT, 2025). First, we present the foundational hardware and software organization. In the following section we describe our metadata management service, which is a representative example of a higher-level data-oriented service. We also describe the accounting service, which is a necessity from both administrative point of view but also gives important technical information to improve the overall efficiency of the data centre operations. In the last section we present some conclusions and discuss our future plans and near-term roadmap for further enhancement of the data services portfolio.

2 HPC Infrastructure Usage Basic Hardware and Software Infrastructure

The data services of CI&ICT are provided using hardware resources from the data storage and analysis system, which was delivered in 2021 and later upgraded with substantial commercial software deployments. The main data resource consists of 8 servers, deployed in pairs for high availability, where each pair has access to a 140 disk SAS enclosure, thus forming a 6.72 Petabytes disk-based storage resource. 1100 TB of disk space (based on 1.68 PB raw disk space) is used for multi-purpose NFS sharing across the datacentre, using OpenZFS with RAID-Z2 VDEVs. 840 TB are deployed on 3 servers for triple data redundancy using the HPE Ezmeral Runtime Enterprise. 210 TB are deployed on 2 servers with HA failover capability enabled using the IBM Storage Scale (previously known as GPFS) and are used for faster shared storage directly available for user of the HEMUS supercomputer. The HPE Ezmeral ML Ops software stack is also available on the so-called Big Data machines – servers with 4 CPUs and 3TB RAM each, with capacity for orchestration (Kubernetes). Although the main filesystem used for HPC in HEMUS is the one provided via IBM Storage Scale, various communities of users have requested storage with specific requirements and such data spaces are served via the other storage systems, **Fig. 1.** Backup is provided generically to all users, with some specific arrangements active for certain user groups.

Multiple datasets have been identified as containing important results of the various work packages of the Center of Excellence in Informatics and ICT project and they are maintained in data repositories for the purpose of storing intellectual property of the project and enhancing open science interactions. In the next section we explain more about how our data discovery service contributes in this direction.

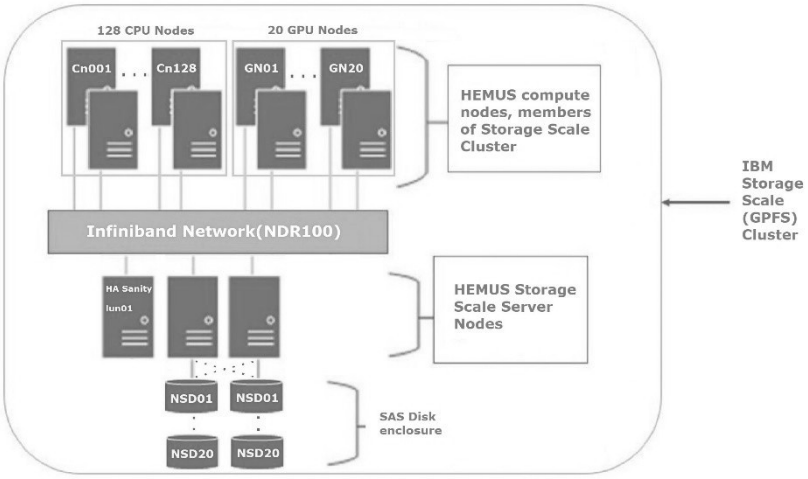

Fig. 1. Schematic view of the IBM Storage Scale deployment at IICT.

3 Data Discovery Service

The data discovery service has been developed and deployed with the purpose of enabling searching of data using the available metadata information. The metadata is provided by the owners of the actual datasets. Metadata is mapped onto standardized facets and is collected from data repositories and provides users with the possibility for flexible search and browsing.

It is possible to search for keywords, partial phrases, creator, organization, publisher, time of publishing, versions, tags, research areas and communities etc. The results are presented in a user-friendly form. In order to enter metadata to the service, the users need to be registered. Searching of a particular dataset is performed using easy to use command-line Python scripts or a simple web accessible form. The search task can be either different types of free-text search or so-called faceted search, concerning tags stored in the metadata accompanying the data. The users may refine their searches inside the received results.

Typically a data discovery service consumes a lot of hardware resources, especially when many users use it simultaneously with sophisticated search queries. Therefore, the service is running on a virtual machine and we maintain the option to upgrade the hardware configuration depending on the usage load. The Data Discovery Service is based on the so-called CKAN, providing the functionality of quick keyword search combined with data tags like the EUDAT B2FIND tags.

Following the organization in CKAN, the authorization is the primary way to control who can see, create and update datasets. Datasets that are marked as "public" are visible to everyone. Private datasets could be seen for example within the frame of the different scientific communities, within the separate team members, or others. The main methods of finding and accessing the data are through an easy to use web-based interface, or using the command-line, but an API (CKAN's Action API) is also available.

The core usage of the Data Discovery Service is in that applications can publish metadata to it. The metadata may be searched by users from the project or by other users, depending on the permissions.

The Internet protocol which allows data records provisioning is OAI-PMH (Open Archives Initiative Protocol for Metadata Harvesting) which is a low-barrier mechanism for repository interoperability. Repositories that expose structured metadata via OAI-PMH can be harvested by Service Providers through OAI-PMH service requests.

The research communities can select which datasets and with what permissions will be exposed through the Data Discovery Service. The communities themselves decide which metadata are made available and how their metadata elements are mapped. More specifically, the Data Discovery Service can be integrated with existing repositories so that basic metadata about the stored datasets is synchronized automatically, based on the trusted relationship between the two services. Rich metadata information about the datasets can be gathered, depending on their usage needs of the respective communities. In general, only keywords, tags and small documents (e.g., input files) are stored in the Data Discovery Service, while datasets or individual large files are presented only as links (URIs).

For example, the Digital Cultural Heritage (DCH) community operates with virtually any type of data: from simple arrays containing numbers that represent chemical or physical characteristics of a given artifact, plain texts, to high-resolution 2D and 3D images, video clips, audio records etc. In addition, many artifacts have unique characteristics, therefore they cannot be grouped together effectively and summarizing queries in PivotTable style cannot be performed on them. Other research communities have fewer types of data to process and the data formats are well documented and standardized. Hence, various DCH metadata are needed to make the datasets usable both inside and outside the data centre. That is why the format of the metadata for each application should be thoroughly and carefully defined, taking into account the needs of all potential end users and applications developers. It is advisable that the uploading of datasets to the repository or to any application-specific database be coupled with the uploading of the corresponding metadata to the Data Discovery Service. The end-users then will be able to search in the Data Discovery Service and obtain either the needed metadata immediately or a link to the datasets where they can obtain the parts of the data they are interested in. Thus there are two flows in this usage scenario:

- dataset metadata flow to the Data Discovery Service (DDS);
- end-user utilization of Data Discovery Service to obtain the metadata directly or the URI for the datasets.

The end-user utilization in this case is less complicated, because a user can log in to the Data Discovery service, be authenticated and authorized within the research community and use the detailed form to perform searching of datasets. Another option would be to use parts of the application for that, through the API. The further processing, based on the received URI, can be done by the application or using other tools.

For the scientific communities dealing with climate modelling and weather forecasting the data formats are well established and in many cases similar software tools are used, although in different ways and in combination with some tools or scripts developed by the researchers themselves. Because of this level of standardization some of the input files may be added to the Data Discovery Service directly together with keywords and descriptions of the types of computations being performed – e.g., validation, simulation, forecasting. The links to the software codes from the code repository service can also be provided. In this way an interested researcher should be able to either repeat the computations or use the same combination of codes on data from their own region, validating the computational approaches and evaluating their applicability for other scenarios. For each computation there should be clear description in the metadata about what kind of assumptions it is based on, in order to facilitate comparisons and increase transparency.

For researchers from the domain of Life Sciences it is important to take into consideration the complex workflows usually utilized. Many types of computations produce large output datasets from relatively small input files. In that case the input files shall be stored together with URI for the result dataset, as well as metadata about the computation itself, e.g., how it was performed, on which machine, etc. The popular format DICOM has some inherent searching capabilities, so it is possible that most of the searching can be done within a DICOM-specific portal, while our Data Discovery Service will play only an auxiliary role.

Examples of some specific datasets from the domain of pure and applied mathematics that are also maintained at IICT:

- The first 100,000 elements of a 1000-dimensional Sobol sequence for simulation/modelling;
- Generators of the stabilizers and description of the orbits under the action of the affine group;
- Set of codes over Zq for different q given by their verification matrices.

4 IICT Accounting System Services

A comprehensive accounting system was developed in order to accurately track and report resource usage across the multiple supercomputers and HPC clusters within the IICT infrastructure. It consolidates logging data from the various cluster environments and respective batch systems and provides automated reporting as well as a user-friendly web interface for browsing and visualizing the data. The data sources comprise of the supercomputers Avitohol (Avitohol supercomputer, 2025) and HEMUS

(HEMUS, 2025), GPU-based virtual laboratory and the Petabyte data storage and analysis system. Apart from the data from the batch system, the Linux wtmp files are also used to extract login session data. (Slurm manager, 2025)

The data processing pipeline is written mostly in python and performs collection, cleaning, storage and reporting. Reports are necessary for various managerial tasks, and that is why they include not only user-level data, but also projects and research infrastructures that are associated with the particular user activity. **Fig. 2** shows a schematic view of the architecture of the accounting system.

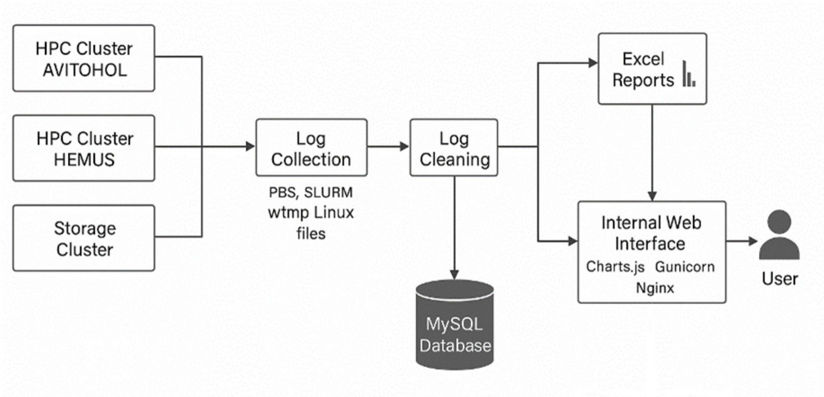

Fig. 2. Scheme of resource utilization information gathering by the accounting system.

A web interface has been developed for online monitoring of the accounting data. It has a backend using Flask (Python), uses Gunicorn and Nginx for serving scalability and performance, and Bootstrap and Charts.js for dynamic rendering of charts and producing visual summaries. Users can browse historical usage data, filter and generate charts on demand. Example chart is shown in **Fig. 3**.

Although the accounting service is mainly important for aggregating and analyzing HPC usage data, we also monitor the distribution and usage of data space in the various filesystems and match that with the datacentre access forms, where users are supposed to specify their needs. Frequently we observe that the data space allocated has been exhausted as the problem being solved is better understood and the applications in use are fully optimized. Our regular analysis of HPC and data usage patterns is the starting point for planning and carrying out of our upgrade plans.

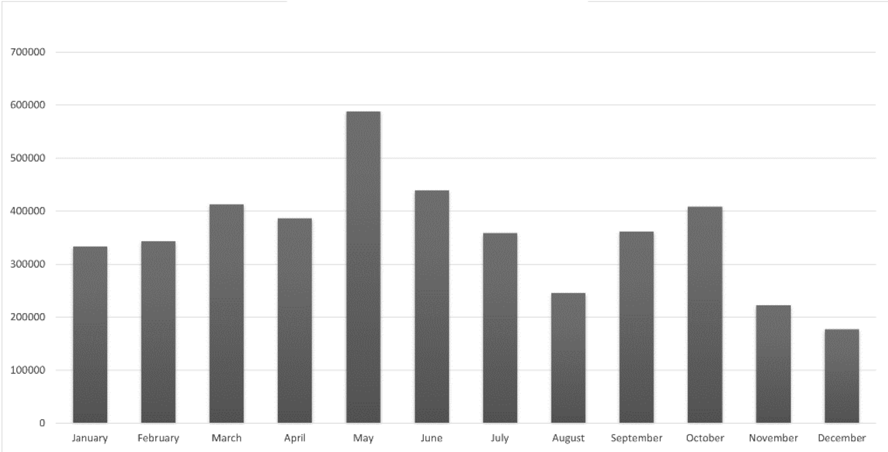

Fig. 3. Scheme of resource utilization information gathering by the accounting system.

5 Conclusions and Future Plans

In this work we showed briefly some of the services in the Data Service Portfolio of IICT and the processes involved. It is important to mention that significant part of our Data storage system capacity is used for storing various collections from the field of Digital Cultural Heritage, and the user feedback has been instrumental in the process of defining the overall architecture of the services.

As the HPC technologies and the High Performance Data Analysis (HPDA) and especially Artificial Intelligence (AI) progress rapidly with frequent disruptive developments, it is important to perform regular technology watch and plan and deploy new and improved services in order to maintain state-of-the-art environment available for advanced research. Our current focus is on making better use of our available server equipment with large amounts of RAM in order to offer temporary data space with exceptionally high throughput and low latencies for demanding AI/Machine Learning training workflows. As we are exploring the various options for collaboration with industry and making HPC technology more widely available to SMEs, for example withing the framework of the EuroCC2 project (EuroCC ACCESS, 2025) and the work of the Bulgarian National Competence Centre in HPC (NCC Bulgaria, 2025), we are also working on better implementation of encrypted data spaces with secured access protocols. Although we do not foresee important hardware acquisitions at this time, we have a comprehensive plan for deploying commercial and free/open source software for updating our data service portfolio.

Acknowledgements.

This work of authors has been partially supported by the EuroCC2 National Competence Center Bulgaria; project No 101101903, which has received funding from the

European Union through the European High Performance Computing Join Undertaking (JU) and the Ministry of Education and Science of the Republic of Bulgaria. The work of Emanouil Atanassov was partially supported by the Centre of Excellence in Informatics and ICT established under the Grant No BG05M2OP001-1.001-0003, financed by the Science and Education for Smart Growth Operational Program and co-financed by the European Union through the European Structural and Investment funds.

References

Atanassov, E., Karaivanova, A., & Gurov, T. (2024). HPC Ecosystem and Competences in Bulgaria. *Digital Presentation and Preservation of Cultural and Scientific Heritage*, *14*, 301–310. https://doi.org/10.55630/dipp.2024.14.30

Avitohol supercomputer. (2025, May). https://top500.org/system/178609

EuroCC ACCESS. (2025, May). https://www.eurocc-access.eu/

CoE in Informatics and ICT. (2025, May). http://ict.acad.bg

HEMUS. (2025, May). https://www.top500.org/system/180208

NCC Bulgaria. (2025, May). https://eurocc-bulgaria.bg/

Slurm manager. (2025, May). https://slurm.schedmd.com

Sterling, T., Anderson, M., & Brodowicz, M. (2017). *High Performance Computing: Modern Systems and Practices*. Elsevier. https://doi.org/10.1016/C2013-0-09704-6.

Received: April 15, 2025
Reviewed: May 25, 2025
Finally Accepted: June 05, 2025