



PLG Portal - platform for managing distributed computing resources in a federated infrastructure

Andrzej Zemła, Marek Kasztelnik

ACK Cyfronet AGH
EGI 2024 conference
30.09-4.10 2024, Lecce

Partners:





PLGrid Consortium

Established in 2007 under an agreement between the largest polish HPC centers in order to build and manage distributed computing infrastructure for science - PLGrid infrastructure

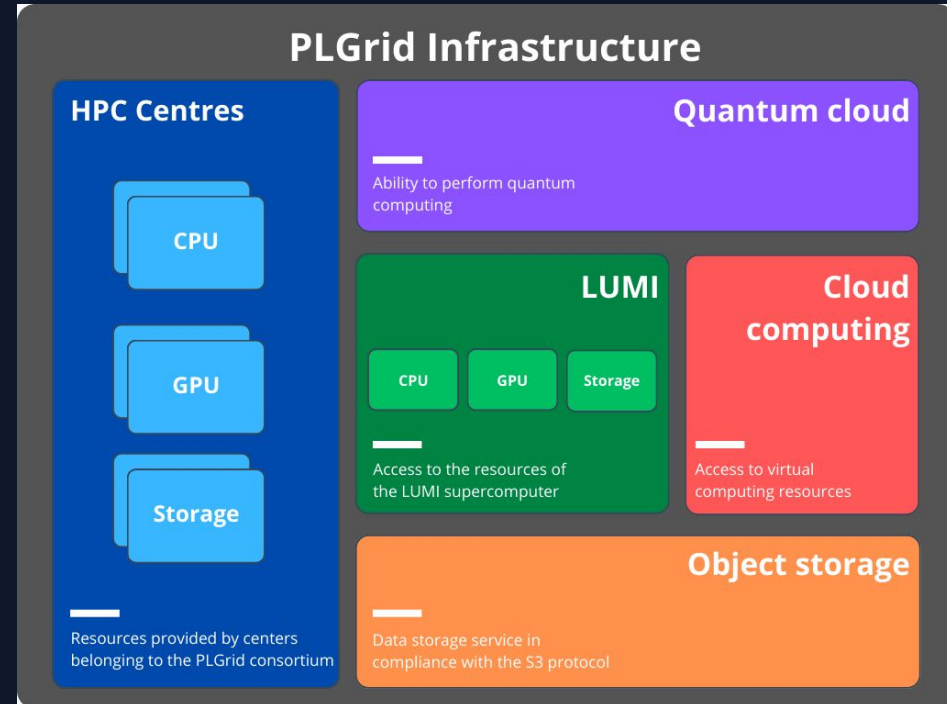




Project aims

Platform for integration of polish distributed computing resources with the infrastructures of European EuroHPC JU resources and other specialized platforms (single point of access)

- Authentication and authorization infrastructure (AAI)
- Integration platform <https://portal.plgrid.pl>
 - user management portal
 - grant (project) application process
 - resource allocation
 - integration and management of services
- Controllers





Authentication and Authorisation Infrastructure

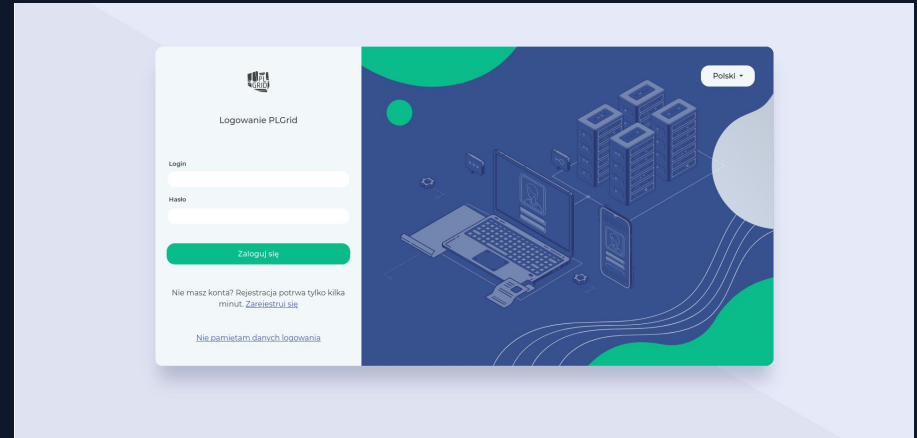


Authentication and Authorisation Infrastructure (AAI) <https://sso.plgrid.pl>

Open source solution -

<https://www.keycloak.org/>

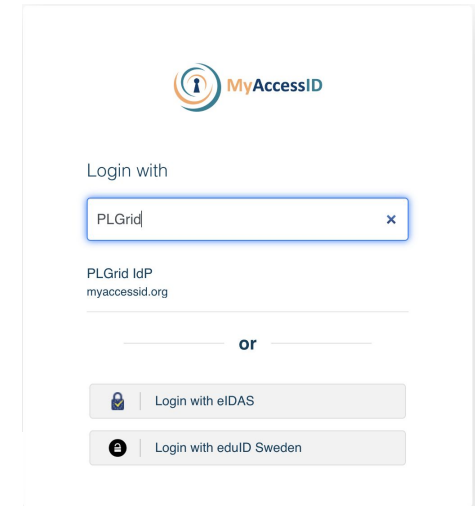
- Centralized Management
- Single-Sign On - login once to multiple applications and services
- Standard Protocols - OpenID Connect, OAuth 2.0 and SAML 2.0
- Easy integration with services



Integration with MyAccessID

Integration with **MyAccessID** is necessary to be able to create allocations on the LUMI supercomputer

- added required attributes by MyAccessID,
- integrated with the PLGrid IdP,
- ability to assign allocations to specific users by controller,
- the user have the ability to add SSH keys, which allow access to the login nodes.



Kerberos & SSH keys

Save

Public ssh Key:



User's public SSH key.



Integration platform

<https://portal.plgrid.pl>

The screenshot displays the PLGrid Portal interface. On the left is a dark sidebar menu with the following items: Profile & affiliations, Services, Research teams, Subordinates, Grants, Operator, Unit manager, and Service admin. The main content area features a grid of six grant application cards:

- PLGrid proper grant**: A predefined grant containing specific resources available in the PLGrid infrastructure. Includes an "Apply" button.
- PLGrid pilot grant**: Pilot grant to test user needs for the use of PLGrid Infrastructure. Includes an "Apply" button.
- LUMI test access**: As part of the test access, the user can use the LUMI CPU resources for 1 month (30 days). Test access provides 10,000 CPU hours that can be used. Includes an "Apply" button.
- Quantum computing**: Quantum computing grants implemented as part of PLGrid infrastructure. Available soon... Includes an "Apply" button.
- Quantum pilot grant**: Quantum computing pilot grants (D-Wave Leap). Resources granted until December 31, 2023. Includes an "Apply" button.
- Grants are currently unavailable**: A card with a blue diagonal banner indicating that grants are currently unavailable.

At the bottom left of the screenshot, the URL https://new-portal.pre.plgrid.pl/bundle_categories/1 is visible.



Profile - comprehensive user data

Providing a verified email in the domain of the scientific unit as the essential step to gain full access to the resources.

Details about user affiliation is an integral and crucial component of the user profile.

The possibility of providing one of the most popular scientific database identifiers like ORCID number or ResearchID.

PLGrid Portal / Affiliations

Personal settings

Affiliations

Expires at: 2026-05-31 Active

Akademia Górniczo-Hutnicza im. Stanisława Staszica w Krakowie
Akademickie Centrum Komputerowe
Cyfronet AGH

[New affiliation](#)

Profile & affiliations

Services

Research teams

Subordinates

Grants

Profile

Affiliations

Basic information

Extended information

Agreements

SSH keys

Personal settings management

Profile & affiliations

Choose affiliation type



Academic unit employee

I am employed and carry out research at a university in a Polish academic unit, have at least PhD degree. I can supervise students, PhDs and external collaborators.



Scientific institution employee

I have at least a PhD and carry out my research in a Polish research unit, for example PAN.



Subordinate

I am a student of Bachelor or Master studies. I have a research supervisor with at least a PhD degree, employed at a Polish academic or scientific unit.



PhD student

I am a student of a Polish doctoral school. I have a supervisor with at least a PhD degree.



HPC center employee

I am employed in one of the Polish HPC centers.



External collaborator

I am employed and carry out research in a non-Polish research unit, but I collaborate with an employee(s) of a Polish research unit (with at least PhD degree).



Commercial customer

Resource allocation for business customers.

Affiliation - user association

Flexible and intuitive affiliation type selection system.

Different permissions depending on the type of affiliation.

Verification of affiliation based on the RAD-on (system of information on science and higher education).

Access for the world of science and business.



Research teams

- easy management of your team
- possibility to invite members via email and by sending an invitation link
- archiving teams from completed projects

PLGrid Portal

Research team

Basic Information

Members

Invitation link

Invitations by email

Actions

PLGrid Portal / Research teams / plggstructure

plggstructure

Basic information

Team ID	plggstructure
Team name	Structure and Properties of Nanoparticles in Medicine

[Edit team basic information](#)

PLGrid Portal / Research teams / plggstructure / Invitations by email

plggstructure

Invitations by email


Emails *

[Send an invitation](#)

You haven't sent any invitations yet.



Grant lifecycle



The user obtains all the necessary permissions to apply for a grant.

Start

The user creates a draft, fills out the application, specifies which resources and how many are needed. The user submits the application for review.

Start



Draft



The operator reviews the application. If necessary, it is sent back to the user for corrections. If everything is fine, the application is forwarded to the resource administrator.

Start

1

Draft

2

Verification

—————

—————

The resource administrator reviews the application, may propose resource negotiations, and send it back to the user. Once approved, the grant and its resources become active.

The grant lasts up to 12 months. The user can request resource renegotiation, subject to admin approval.



The grant ends,
and the user has
30 days to settle it.



The operator reviews and approves the settlement, closing the grant.





Resource allocation

Customization of the resource allocation process

Integration with different types of resources, especially from other infrastructures, must ensure consistency and interoperability.

We cannot assume that the procedural and technical requirements will be the same.


Main points of customization:

**application
process/form**

resource metrics

granting policy

Glossary



The class of grants with common set of policies, mostly related to infrastructure (ex. PLGrid, LUMI)

The exact HPC system, tier Cloud, etc... (ex. Athena, BEM, LUMI, Object Storage)








Specific type of computing or storage elements (ex. CPU, GPU, capacity, VM)

Requested resource parameters

Grant type

Access to specific types of resources may have different application policies:

- application process
 - continues recruitment
 - regular calls
 - fast/test access (pilots)
- time frame
- review process (if required)
- settlements/reports
- evaluation survey

 Basic information
 Bundles
 Bundle categories
 Calls
 Grant form extensions
 Review types
 Grant evaluation



Grant type

Basic template created on the basis of plgrid grants

Splitted into sections (basic informations, resources, grant details, etc.)

Can be extended with various types of fields, ex:

- rich text
- string
- select
- rating

Grant details

Grant start date * Calendar icon Duration in months * Dropdown arrow

The application must be reviewed and evaluated before the grant is allocated. The above date will be considered as the actual start date of the grant. In case of prolonged negotiations, the grant will be launched as soon as the application is accepted.

Field of Science * Dropdown arrow

Research purpose * Rich text editor icons: Bold, Italic, Underline, Link, H1, H2, H3, H4, H5, H6, Bulleted list, Numbered list, Indent, Outdent, Undo, Redo

Substantive description of the research, public abstract - information about planned works under the grant and what the user wants to achieve. In this field you can add graphic files, links, pdf files. (min. 250, max. 1000 characters)

Computational readiness * Rich text editor icons: Bold, Italic, Underline, Link, H1, H2, H3, H4, H5, H6, Bulleted list, Numbered list, Indent, Outdent, Undo, Redo

Justify the requested computational and storage resources. Provide basis on which the requested resources were estimated. Describe the computational techniques and platforms that you will use, including application code names (together with exact version), programming languages, libraries, and other software used. Describe also parallelization and scalability aspects, including parallel programming system used (e.g., MPI, PGAS, "embarrassingly parallel"). If possible, provide data for your application parallel performance, speedup and scalability.

Attachments

Drag and drop to upload or [click here](#) to browse
10 MB file size maximum. Allowed file type: pdf.

Previous publications of the user on the subject being the basis of the grant.

[Save changes](#)



Site has a relationship to the type of grant, provides various resources and has dedicated administrators. They can:

- edit basic information like name, logo,
- accept the resources under grants,
- add public API authentication key (JWT),
- define resources:
 - postfix,
 - operator acceptance criteria.



Full site name in Polish *

Full site name in English *

Resources:

ACK Cyfronet Object Storage
DC-Nawojki

Key *

NOTE: All white chars will be converted to - (dash) and capital letters to lowercase due to normalization.

Postfix

Operator acceptance criteria



Resource

Metrics
Metric '1'
Metric '2'
Metric 'n'

Different types of resources can have different metrics and/or units

- CPU - kilo-hours
- GPU - hours
- QPU - seconds!!
- storage capacity- GB (static pool) or TBh
- cloud resources (# VM, VM spec, network spec... the whole bunch of metrics)

Additional metrics:

- memory (per slot or working node)
- specific cpu/gpu type in heterogeneous machines
- max walltime


And any combinations, names whatever we can think of.

The platform has to ensure flexibility!

The screenshot shows the configuration page for ACK Cyfronet Cloud. At the top, there is a logo for 'CYFRONET' and the text 'ACK Cyfronet Cloud'. Below this, it says 'Cloud Cyfronet'. A dark blue button labeled 'CLOUD-3.0 - Szkic' is visible. The main configuration area contains several input fields with labels and units:

- Liczba publicznych adresów IP: 10
- Liczba rdzeni: 4
- Maksymalna pamięć RAM: 4 GB
- Maksymalna pamięć dyskowa wolumenów: 4 GB
- Maksymalna liczba wirtualnych maszyn: 10

At the bottom right, there are two buttons: 'Anuluj' and 'Stwórz alokację'.



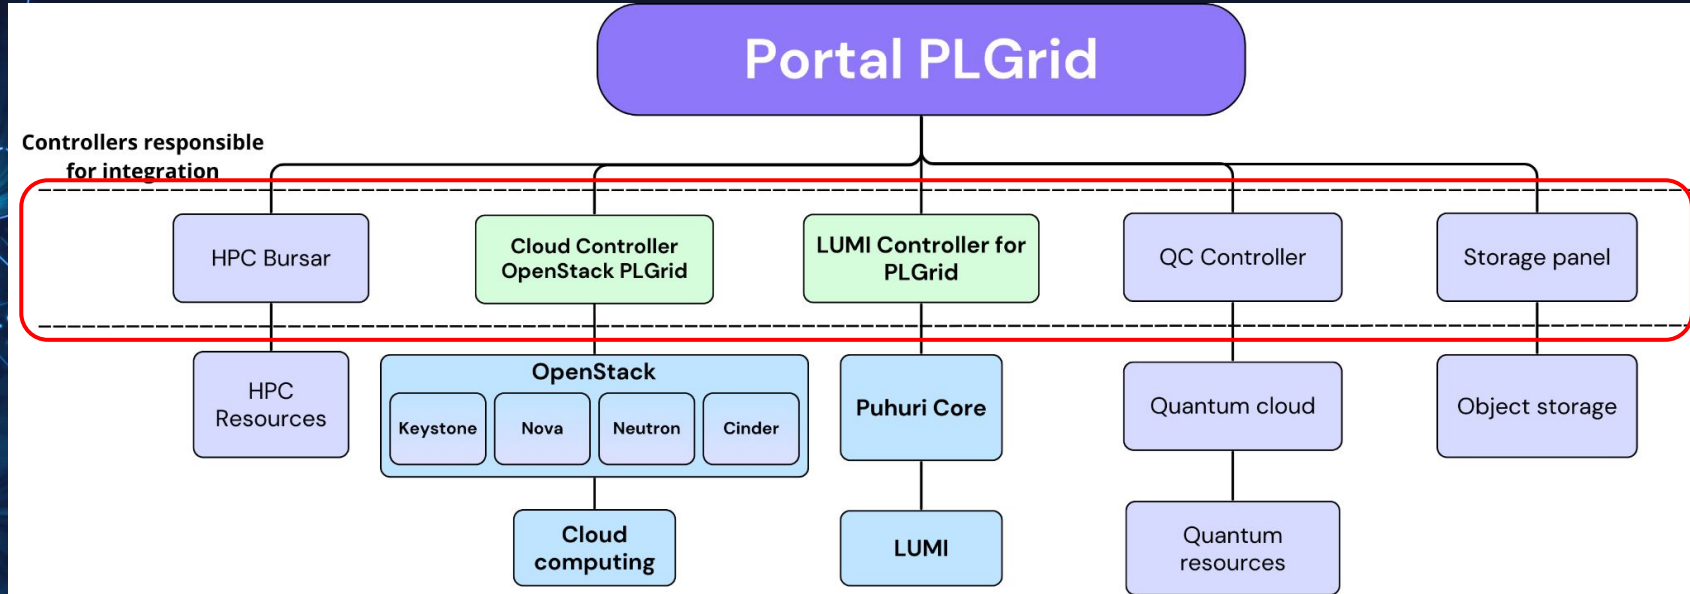
- resource metrics scheme is fully customizable
- metric attributes:
 - name
 - value
 - types:
 - int (with min and max bounds)
 - string
 - select box
 - unit value
- It allows to define virtually any type of resources

Based on the metrics schema, an appropriate API is generated

```
[
  {
    name: String,
    grantName: String,
    resource: String,
    site: String,
    pilot: Boolean,
    start: Date,
    end: Date,
    lastChange: Time,
    status: Enum(accepted, blocked),
    parameterValues: [
      {
        key: String,
        value: String/Number
      },
      ...
    ]
  },
  ...
]
```



Integration scheme





Why controllers layer approach?



Security

An external app **run periodically based on data from the API** is more secure due to less access to it from the outside world.

Moreover, the web application does not need to have information enabling authentication to external applications with the administrator role.





Consistency

The use of controllers independent of the main web application allows you to **maintain order** in the code and **consistency in the operation** of the application.

Thanks to this, the PLGrid Portal does not have to contain unnecessary and non-generic functionalities needed for integration with other e-infrastructures.





Maintenance

Maintaining a small controller independent of the main web application is very easy.

The local configuration of such a controller is independent of the PLGrid Portal and requires the installation of only those dependencies that are used in the controller.





Flexibility

Controllers provide a **high level of flexibility** in mapping information returned by the PLGrid Portal API.

This makes it **easier to transform** the received information to meet the requirements of external e-infrastructures.





Possibility to choose the best-suited technology

In the case of some controllers, dependencies available only for specific technologies are used for integration purposes.

One example is a CCOP controller, designed to integrate the PLGrid Portal with cloud computing resources. For this purpose, it uses, among others, the OpenStack SDK3 library, the only official version of which was created for Python.

Creating separate controllers allows you to select technology and tools in such a way that their implementation is as simple as possible.



Team



**Marek
Kasztelnik**



**Konrad
Perłowski**



**Bartłomiej
Bańda**



**Mieszko
Cholewa**



**Arkadiusz
Zięba**




**Radosław
Duma**



**Jakub
Perlak**



**Magdalena
Paluchowska**



**Łukasz
Kitowski**



**Jacek
Chudzik**

Partners:





Thank you for your attention



Outline

PLGrid is a nationwide computing infrastructure designed to support scientific research and experimental development across a wide range of scientific and economic fields. PLGrid provides access to supercomputers, quantum computers, specialized accelerators for artificial intelligence, cloud computing, disk storage, optimized computing software and assistance from experts from the entire Poland. The Polish PLGrid infrastructure is managed by the PLGrid Consortium, established in January 2007, which includes the following computing centers: Cyfronet, ICM, PSNC, CI TASK, WCSS, NCBJ.

In order to make it easier for users to use the available distributed resources, it was necessary to create a centralized platform that includes many applications, tools, and solutions, with the PLGrid Portal as its main component. The platform has been developed since the beginning of the PLGrid Consortium, going through successive new versions. Thanks to the experience gained over 10 years, it has become a mature and flexible solution. This allows us to easily adapt to changing requirements, which makes us able to effectively respond to new challenges in terms of both user and operational convenience in federated infrastructure.

As the main application in PLGrid Infrastructure, the Portal PLGrid consists of many elements from creating an account to requesting distributed resources through PLGrid grants. All the necessary functionalities for the User like creating and managing an account, affiliations, subordinates, teams, services (access to resources), applications, and ssh keys. However, it is the process of requesting resources that is specific. The user fills out a grant application, which is negotiated with resource administrators - and after the grant is completed, it must be settled. From our perspective, the most important thing was to realize such Portal that would be flexible. We have a variety of types: accounts, teams, services, and grants. In an easy way, new types of resources can be defined, that will have other limitations. For example, a specific account type cannot create a specific grant type or a given type of grant has other fields that are required to be provided by the user. On the operational side, we have many roles, where the main like Resource/Service Administrator have dedicated web views, the ability to replicate data to LDAP, and API access which allows you to synchronize all data among various HPC centers and different types of resources (computing, storage, object storage, cloud, etc).