

# EuroHPC Federation Platform

A FEDERATED PLATFORM FOR HPC INFRASTRUCTURE IN  
EUROPE, BUILT WITH OPEN SOURCE SOFTWARE

FOSDEM 2025, HPC, BIG DATA & DATA SCIENCE DEVROOM

HENRIK NORTAMO, CSC - IT CENTER FOR SCIENCE, FINLAND

# Structure

- What is the EuroHPC JU & what types of resources are we federating
- Current issues and challenges for end users
- Architecture
- Main components
- Open source aspect

*Side note: The presented solution is fairly large-> I will not have time to cover it fully.*

# The EuroHPC JU


- "EuroHPC JU (The JU) is a joint initiative between the EU, European countries and private partners to develop a World Class Supercomputing Ecosystem in Europe."
- The JU owns procured compute infrastructure which is hosted and co-funded by several separate consortiums consisting of one or more countries.
- The entity hosting the infrastructure is called a Hosting Entity (HE)
- End users of the compute infrastructure are academic researchers, research institutes, public authorities, and industry



**EuroHPC**  
Joint Undertaking



MareNostrum 5 Supercomputer  
4480 x NVIDIA H100 GPUs

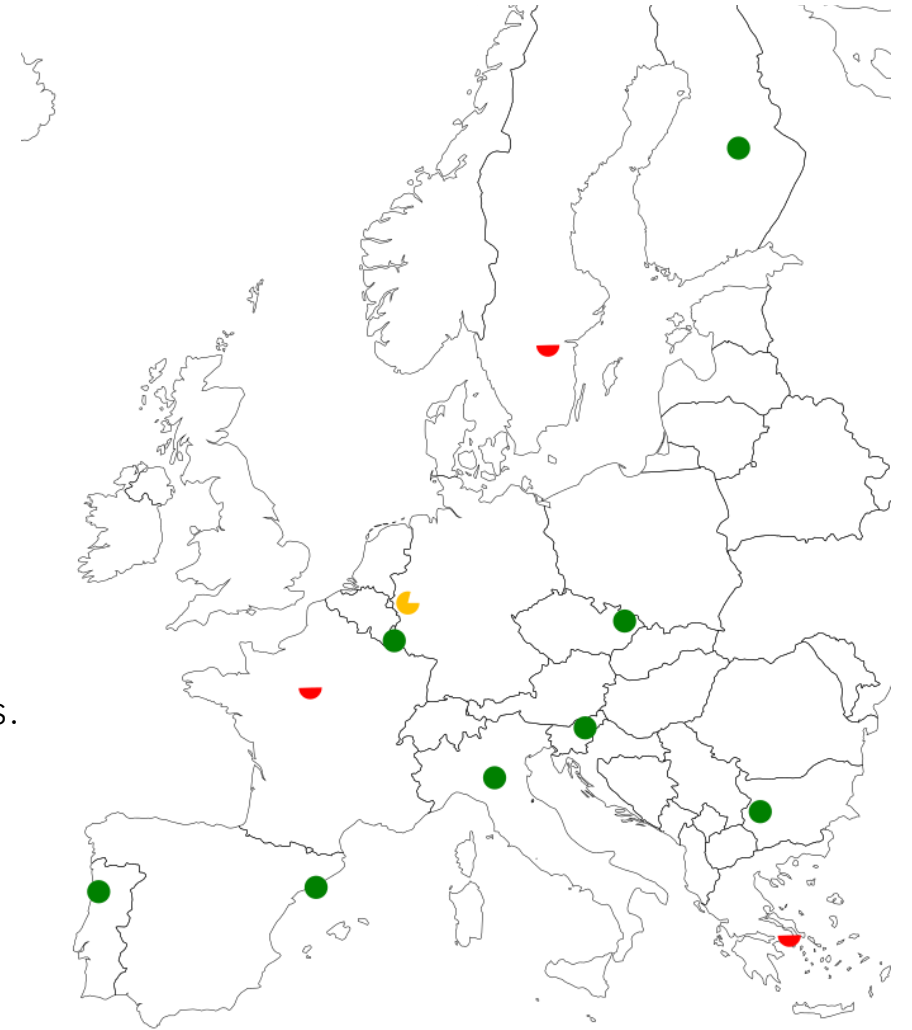


Leonardo Supercomputer  
13824 x NVIDIA A100 GPUs



# Federation targets for the platform

- Classic supercomputers (Active (9) planned (3), shown on map)
  - Massive amount of Linux servers, connected via a parallel filesystem and highspeed interconnect. Compute is done via a batch queueing system (e.g Slurm) otherwise standard Linux with local users.
  - e.g LUMI: 2978 GPU nodes with 4 x MI250x GPUs and 1x64 core AMD EPYC each + 80PB of parallel storage.
- AI Factories (7 announced )
  - Supercomputing infrastructure with AI-optimized computing capabilities.
  - First proposal selected end of 2024 -> Do not exist yet.
  - Personal guess: Lots of GPU's, compute with more features around data and 'cloud-like' features
- Quantum computers (8 in various stages)



For more info on the systems and timelines: <https://eurohpc-ju.europa.eu/>

# Current Issues

- End-users acquire fully separate accounts and projects/allocations through completely different processes for each system.
  - Steps like initial user identification must also be re-done for each system
- Increasing number of new user-groups which are not as familiar with classical supercomputing/computing in general or are used to a different set of tools or paradigms .
  - AI being one of the prime examples, and general industry usage another one.
- Growing heterogeneity of both compute hardware and environments
- Compute is much easier to move than data

# The federation platform

*The consortium delivering the service*

A platform federating the access to all EuroHPC systems, with the main features being:

- **Federated identity and Single-Sign-On (SSO).** Users utilize the same login and identity (e.g. granted via their home institution) to authenticate to all services and access all supercomputers.
- **Resource allocation, management and monitoring across systems.** Users can see what allocations they have on each system in a single place.
- **Direct access utilizing SSH certificates.** Short lived certificates which are obtained via a login flow with optional MFA.
- **Interactive web based usage** with e.g. remote desktop, shell sessions and Jupiter notebooks. Ability to launch batch jobs and browse files on the supercomputers.
- **Federated software catalogue** providing a pre-installed pseudo-uniform software stack on all systems
- **Advanced workflows and data transfer.** Workflow execution and data transfers across systems with smart scheduling capabilities

Work started in January this year, production with currently online EuroHPC systems integrated Q1/2026

**NORDUnet**  
Nordic Gateway for Research & Education

**GÉANT**  
Networks • Services • People

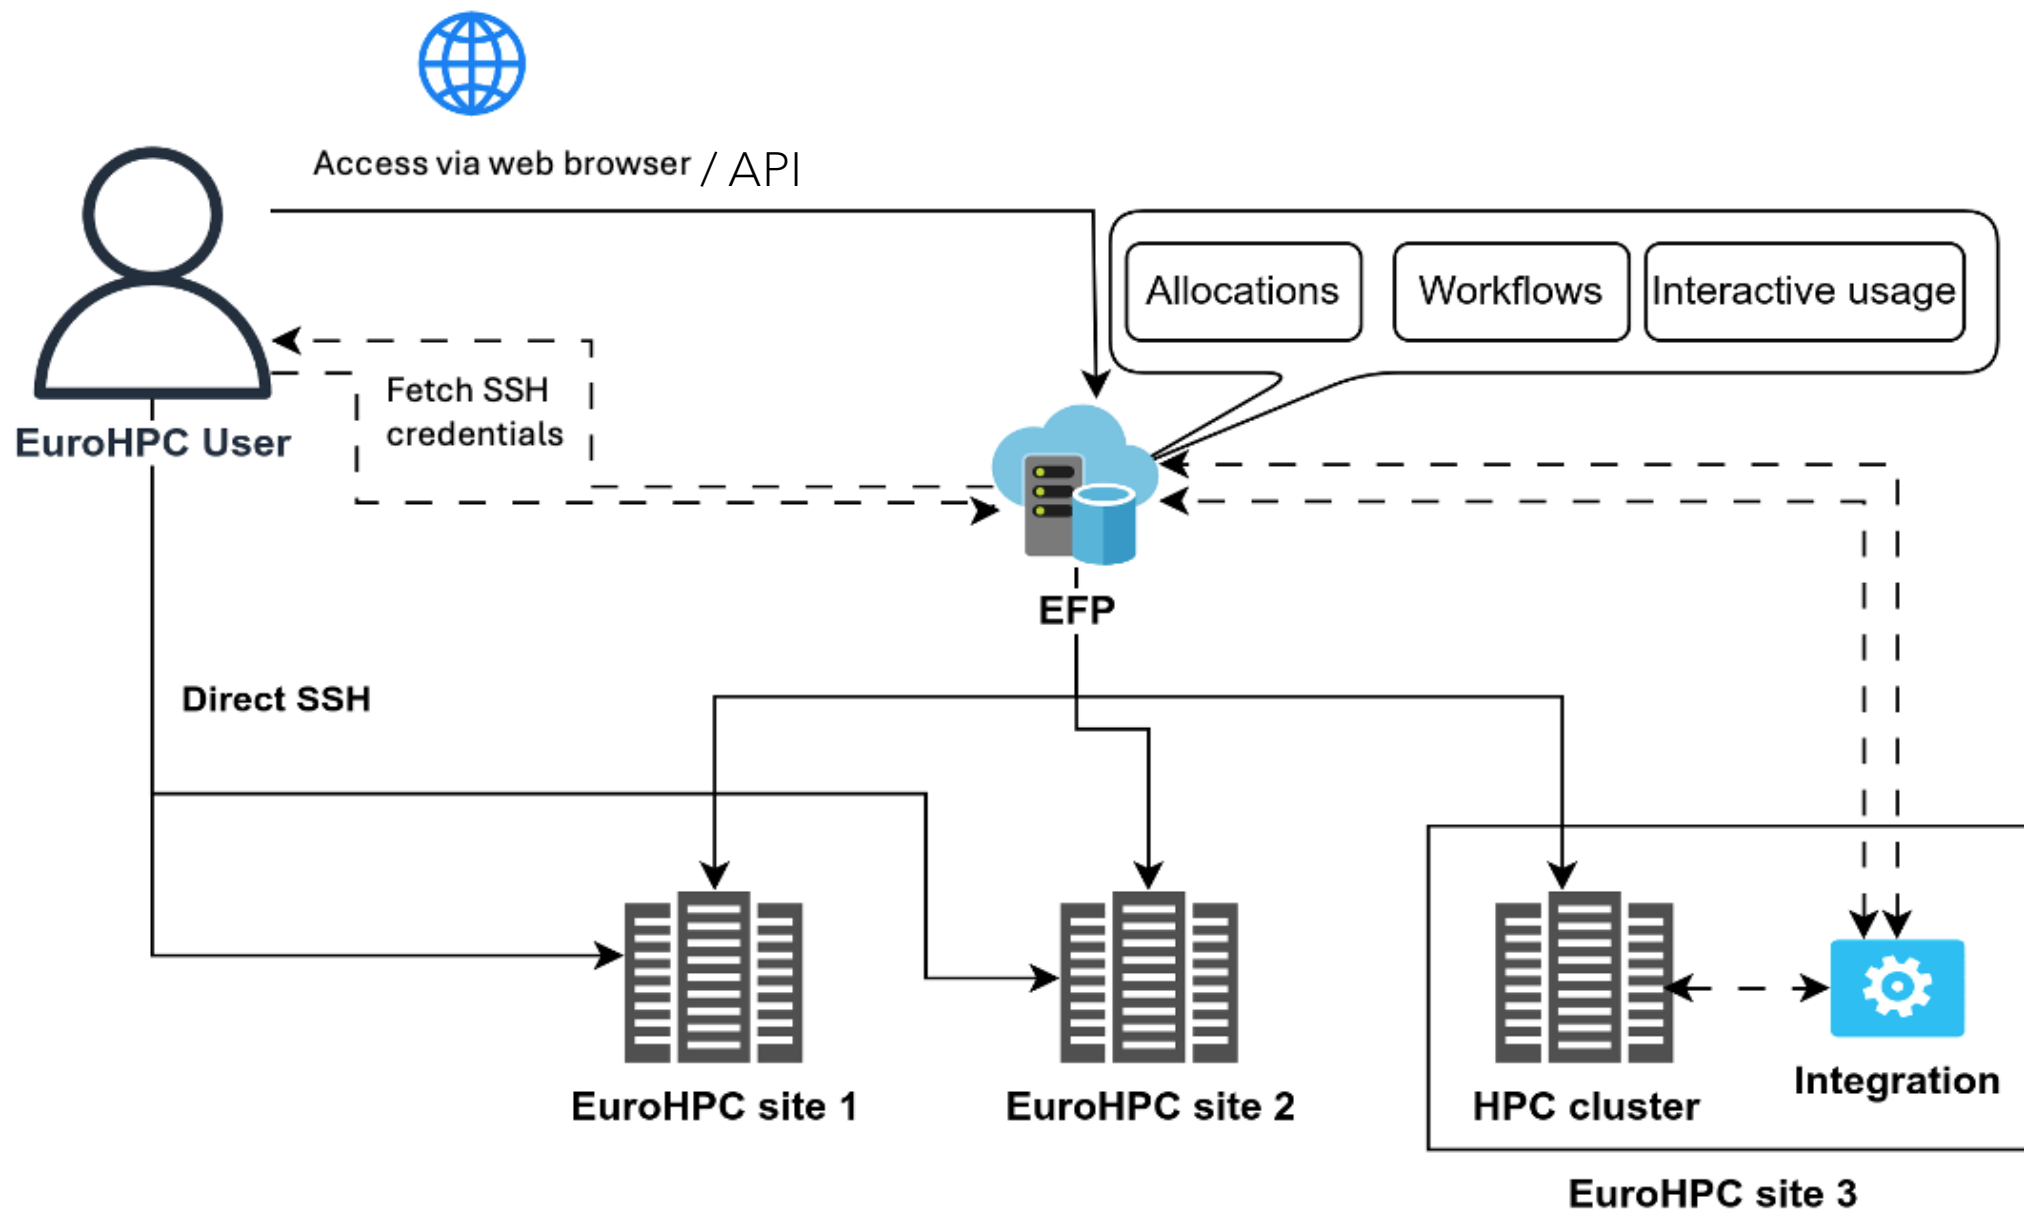
VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER




**Ghent University**



ICT Solutions for Brilliant Minds



# Architecture



Modularity, Flexibility and API centric are among the core design principles, along with being as non-intrusive as possible



# Main components powered by open source software

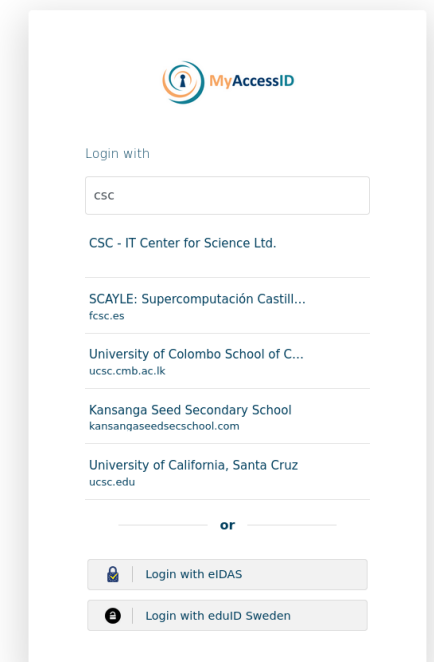
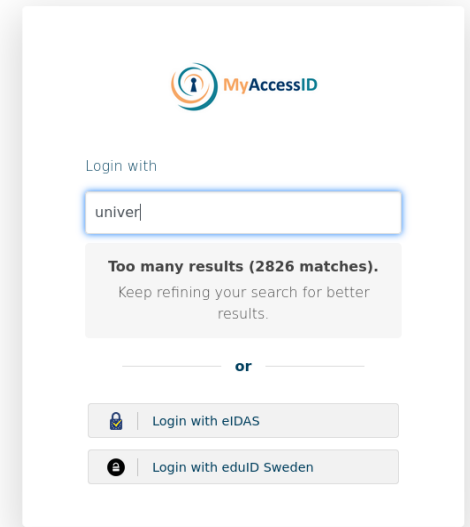
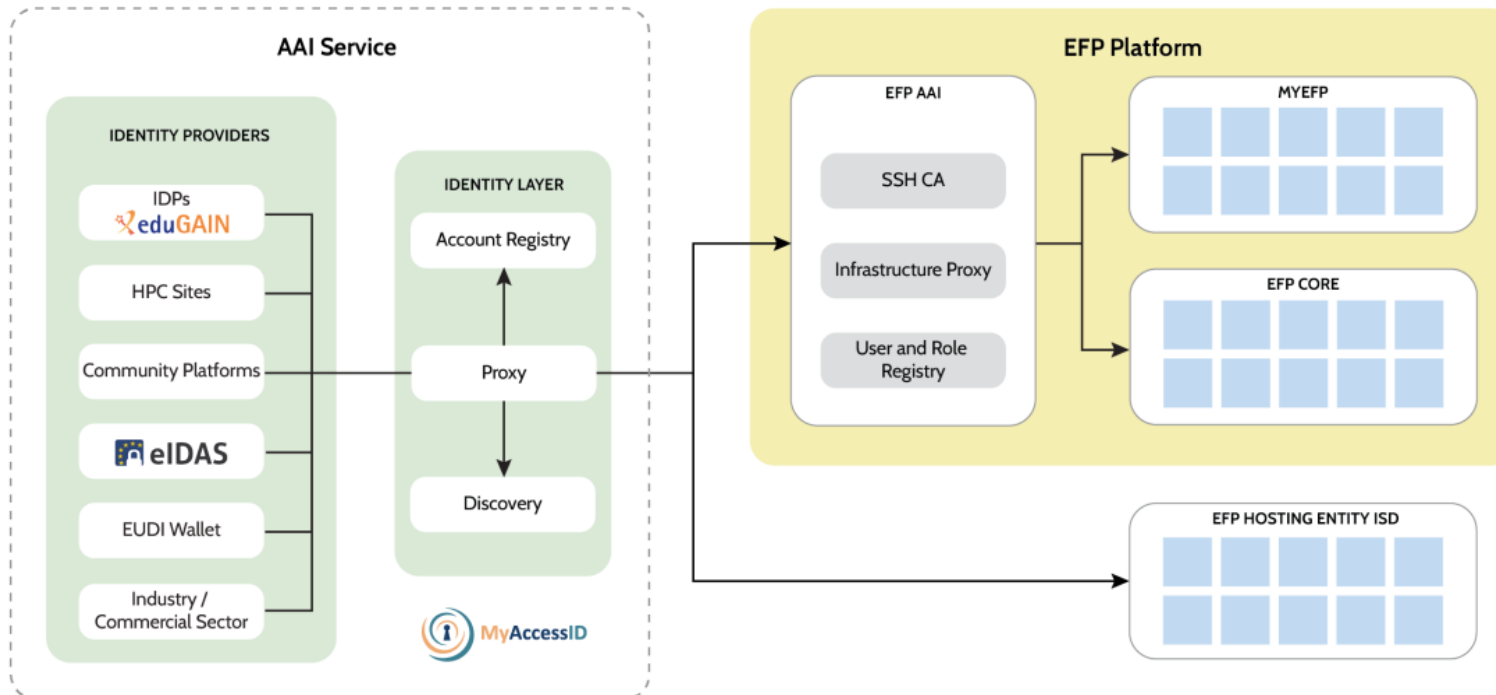
- EuroFP AAI -> Leveraging [MyAccessID](https://wiki.geant.org/display/MyAccessID), <https://wiki.geant.org/display/MyAccessID> SSH CA, Infrastructure Proxy
- EuroFP Allocations -> [Waldur](https://waldur.com/) <https://waldur.com/>
- EuroFP Interactive -> [Open OnDemand](https://openondemand.org/) <https://openondemand.org/>
- EuroFP Workflows -> [LEXIS Platform](#), [HEAppE Middleware](#)
- EuroFP Reporting (also includes monitoring) -> [Grafana](#), [Icinga](#), [OpenSearch](#)
  - Waldur also shows some information on resource consumption
- EuroFP Helpdesk -> [Zammad](https://zammad.com) <https://zammad.com>
- EuroFP Software Catalogue -> [EasyBuild](#) + [EESSI](#) via [CernVM-FS](#)
  - <https://easybuild.io/> <https://www.eessi.io/> <https://cernvm.cern.ch/fs/>

Most of the platform consists of a bunch of open source components, which we then glue together and work out the necessary integrations to the federated systems.

For all the highlighted stuff you can just go test it out. Components are also used separately in production on multiple systems.

# EFP AAI

- EFP Utilizes MyAccessID for identity federation, but does not provide it
  - Simplified -> an AAI Proxy + Discovery service + Account registry
- MyAccessID can be integrated as an OpenID Connect Provider or SAML Identity Provider
- Identity federation can also be utilized directly without going through the platform
- SSH CA for direct access to the systems



# EFP Allocations ->



- Provides a unified allocation and project membership management capability
- Users and resource management portal provides project-centric dashboards.
- EuroFP Allocations allows project members to see and manage their allocations, request new ones either directly or via published calls for access.
  - Resource allocations for EuroHPC systems will be done via a separate EuroHPC Peer-Review Platform.

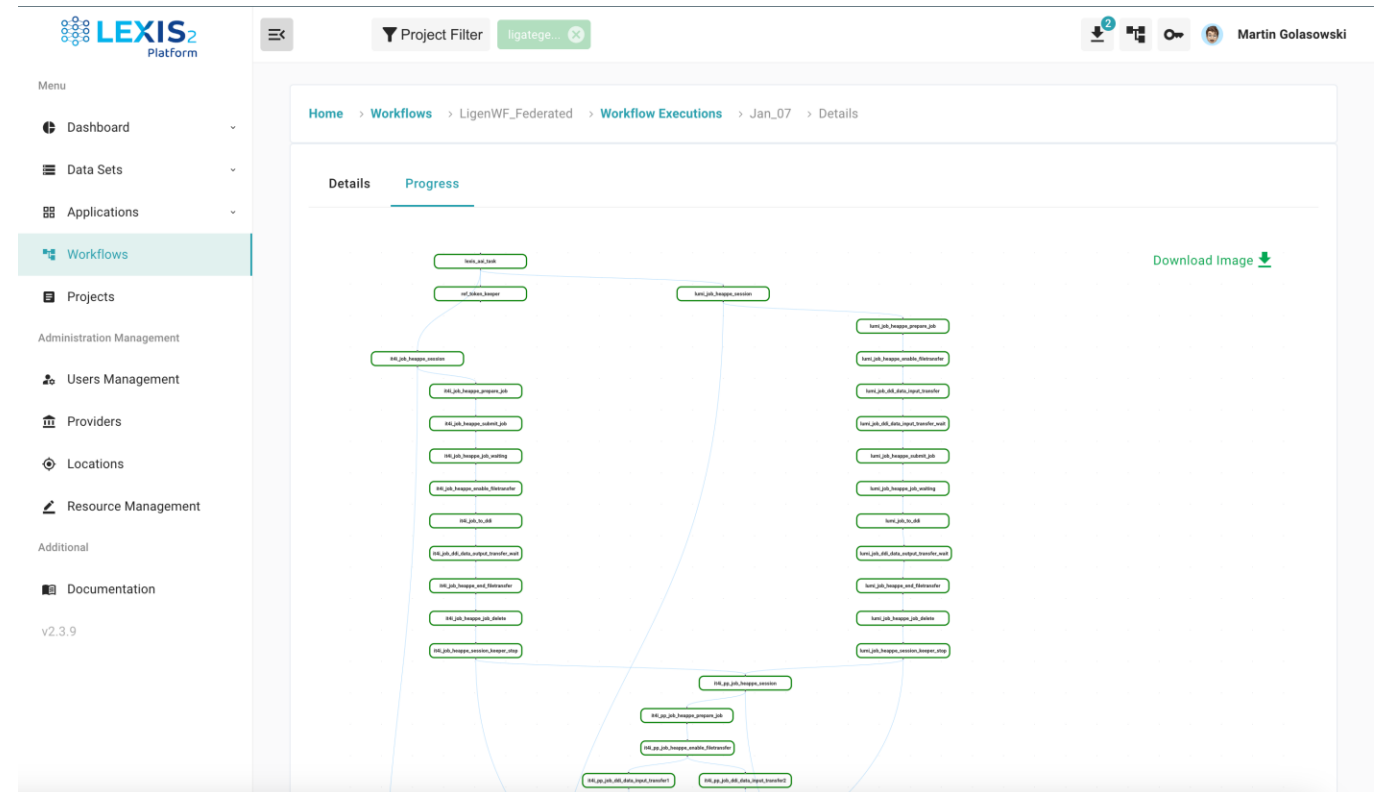
The screenshot displays the EuroHPC portal interface. On the left is a navigation sidebar with options: Add resource, Organizations, Projects, Resources (with sub-options for All resources and HPC), Reporting, Calls, Service catalog, Support, and Administration. The main content area shows the breadcrumb path: Organizations / The ...king / Projects / 6-20...ntry / HPC / Allocation for 6-200 - Finish Carpentry. Below this, there's a card for the allocation with a status of 'OK' and an offering name 'MARENOSTRUM 5'. To the right, there are progress indicators for CPU, GPU, QPU, and Storage allocations, all at 0/100, along with a Refresh button and an Actions dropdown. Below the allocation card is a 'Getting started' section with tabs for Usage and Resource metadata. The 'Getting started' content includes a welcome message, a list of prerequisites (SSH Client, HPC Account, Open Your SSH Client, Connect to the HPC System, Authenticate), basic commands for navigating directories, listing files, and copying files to/from the HPC system, and instructions for exiting the HPC system.

# EFP Workflows ->



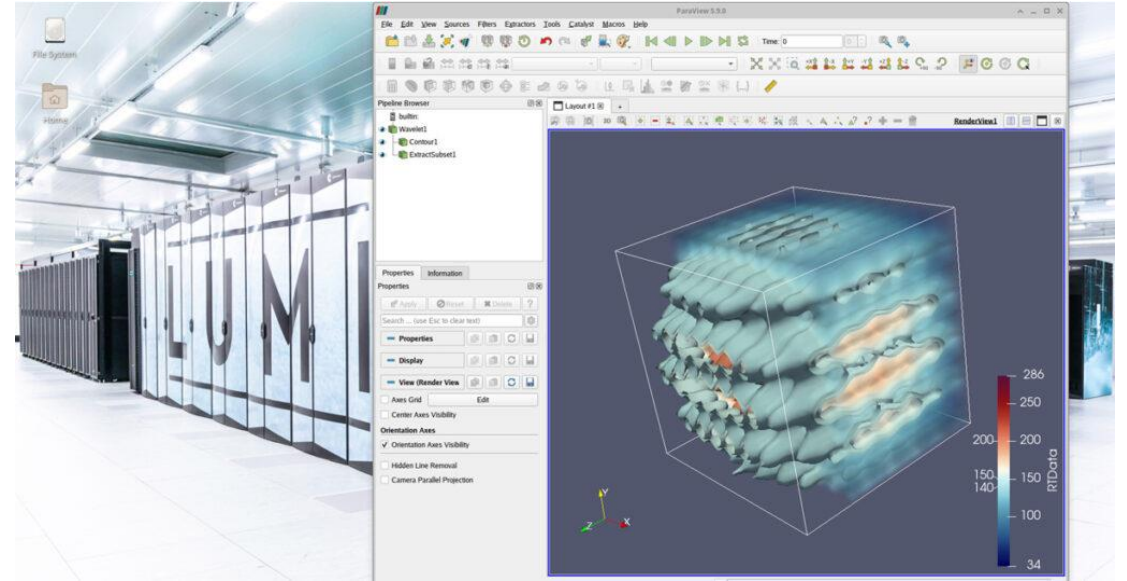
- Easy access to visually managed workflows and distributed data management across the federated resources.
  - Also allows API usage from various languages like Python or R.
- Smart scheduling policies across federated resources in workflow executions based on metrics from HPC sites.
- Multi system workflows
- Data staging from external sources
- Multiple backends e.g HPC batch scheduler and kubernetes

- HEAppE Middleware: REST API for restricted access to HPC infrastructure




# EFP Interactive -> nDemand

- Easy to use web interface for new and industrial users to be in interactive mode with HPC
- Applications like Jupyter notebooks, remote desktops and shell
- Fine-grained job level management and file management across the federated resources.
- **NOTE:** EFP will be developing features for having Open OnDemand running further from the system



## Pinned Apps



```

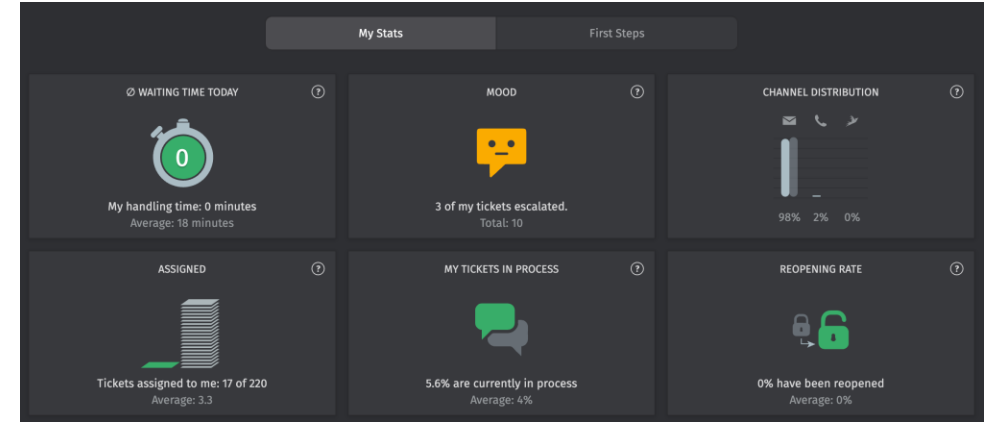
https://www.mahti.csc.fi/pun/sys/shell/ssh/default
Host: mahti-login.csc.fi
[(u, 7)@M-l12 ~/] ls
1.1.0
150 node_nextflow.txt
affinity_tests
allas_setup
a.out
BASHRC
Desktop
dirColors
Documents
Downloads
encrypted.sif
encryptedU.sif
enter.sh
Fort
fulls.txt
full.txt
get_aff
go
hello
hello.c
hello.f90
hosts
HQJOBS
hq-v0.10.0-linux-x64.tar.gz
index.html
__init__.py
intel.sif
INTELTEST
I0-Tests
job.sh
list_quota
logs
Lustre
mahti_arch.svg
main_bu.f90
main.cu
main.py
m.c
min.c
min.py
MPI_TEST
MungeJWT
Music
novnc.log
nvim
nvim-linux64.tar.gz
ondemand
ood_jobs_jan.txt
ood_jobs.txt
OODUsage
out.json
output.txt
out.strace
out.svg

```

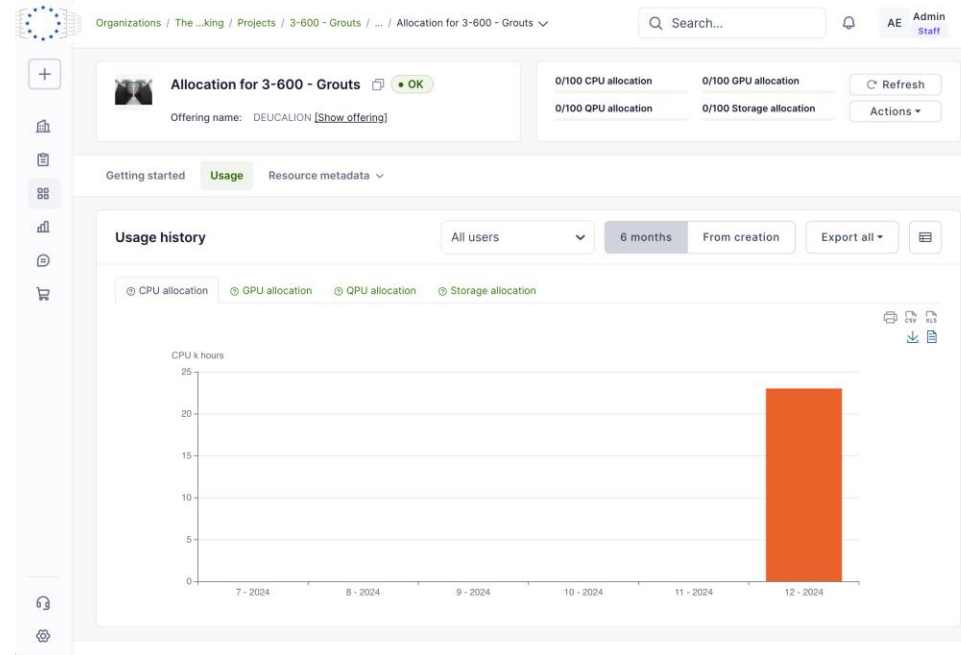
# EFP Helpdesk



It's a helpdesk/ticketing system




# EFP Reporting

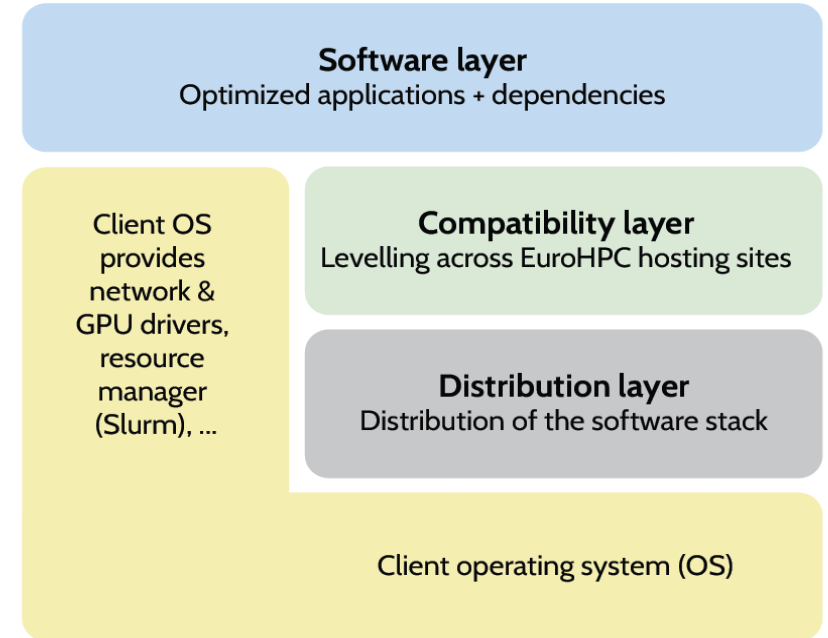


- Centralized reporting solution
- Resource consumption monitoring across different centers

# EFP Software Catalogue ->



- Federated Software Catalogue based on EESSI
  - EESSI -> precompiled architecture optimized binary distributions of common scientific software
- EFP Software Listing provides list of the available software
- Ideally accessible on nodes via native installation of CernVM-FS, but alternative approaches are possible
  - CernVM-FS -> POSIX read-only file system as a FUSE module, Files and directories are hosted on a webserver (local caching recommended)
- The EuroFP Federated Software Catalogue component will manage the heterogeneity of the different system architectures
  - So ideally, basic users would not need to care about the underlying GPU or CPU architecture.



# The open source aspect

This kind of project would probably not be possible without all the open source components:

- Starting from scratch for all the functionality and understanding the user needs is a massive undertaking
- Federation always requires a degree of trust, it's much easier to make the case for open solutions which are already used elsewhere. Also, the trust in the promised features -> If everything here was closed, you would be thinking I'm selling vaporware.
- End users would rather learn/integrate into fewer frameworks -> benefit of using what's already out there
- The ability to modify the components is the key enabler of the integration and creating a cohesive platform.
- Hopefully a mutually beneficial exchange between us and the community.



# **Give and take** ( where we hope to not just mooch off the community)

- While a lot of the work in the project is integration work, certain functionalities will be developed and hopefully contributed to the upstream projects for others to use.
  - Fortunately, the project can continuously open source produced work, as opposed to doing so in bulk at the end of the project.
- For a lot of projects actual integration is the hard part. The variety of the integrated systems will hopefully improve how easily the software projects can be integrated and the availability of examples, instructions and templates.
- A large userbase to use the various projects
  - Bug testing, feedback on features and direction
  - For certain projects, proof of impact.

# More information

- [https://eurohpc-ju.europa.eu/index\\_en](https://eurohpc-ju.europa.eu/index_en)
- [https://eurohpc-ju.europa.eu/paving-way-eurohpc-federation-platform-2024-12-19\\_en](https://eurohpc-ju.europa.eu/paving-way-eurohpc-federation-platform-2024-12-19_en)
- Come talk with me

**Questions ?**