EuroHPC Federation Platform

10TH EASYBUILD USER MEETING *TUE-THU 25-27 MARCH 2025 @ JÜLICH, GERMANY*

HENRIK NORTAMO, CSC – IT CENTER FOR SCIENCE, FINLAND

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

MareNostrum 5 Supercomputer
4480 x NVIDIA H100 GPUs

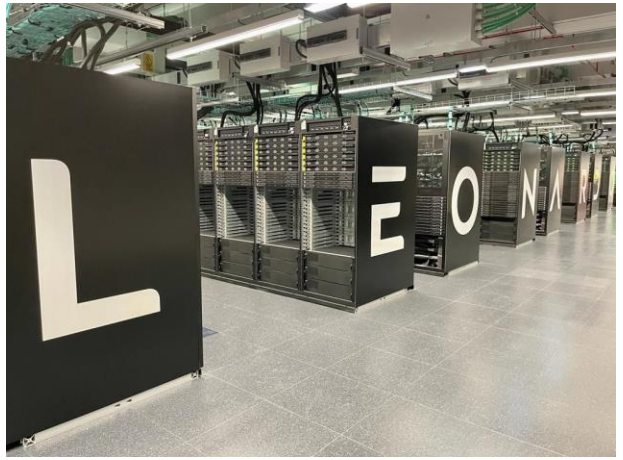

Leonardo Supercomputer
13824 x NVIDIA A100 GPUs

EFP Targets

Current supercomputers

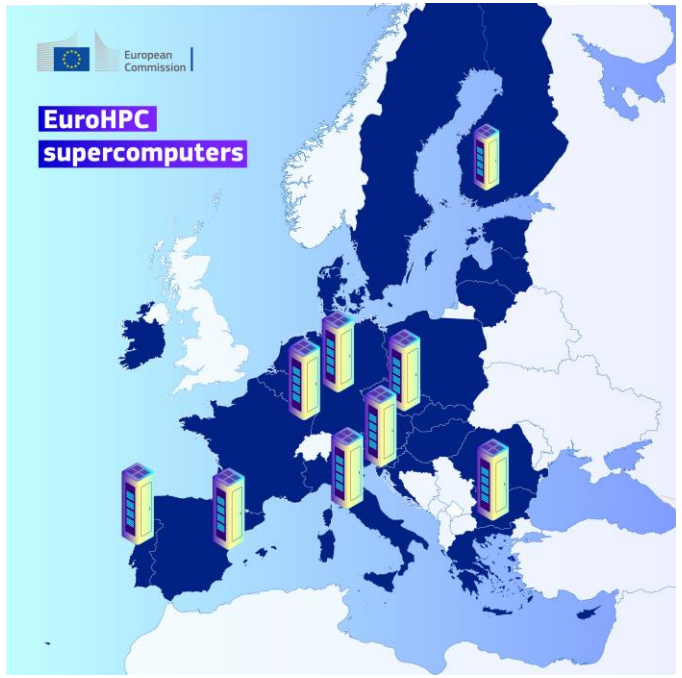

Alice Recoque, Arrhenius, DAEDALUS, Deucalion, Discoverer, JUPITER, Karolina, Leonardo, LUMI, MareNostrum 5, MeluXina, Vega.

AI Factories

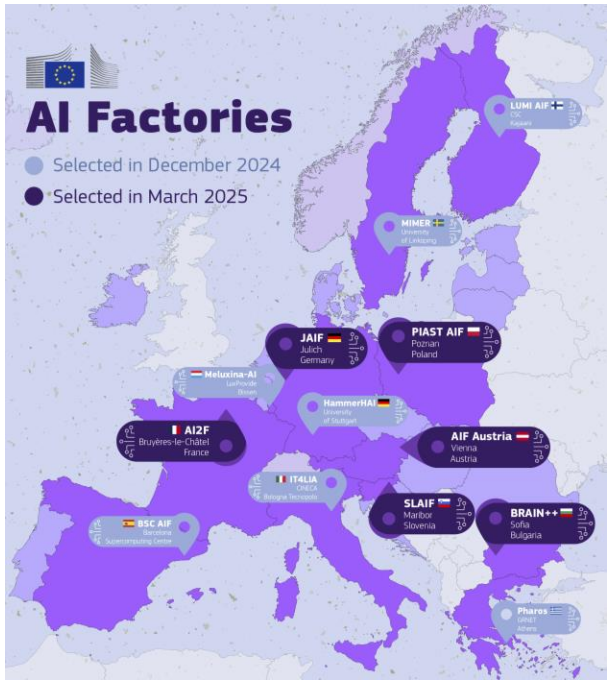

- BSC AIF, HammerHAI, IT4LIA, LUMI AIF, Meluxina-AI, MIMER, Pharos
- AI2F, AIF Austria, BRAIN++, JAIF, PIAST AIF, SLAIF
- …

Quantum Computers

Euro-Q-Exa, EuroQCS-France, EuroQCS-Italy, EuroQCS-Spain, EuroQCS-Poland, LUMI-Q,

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

CSC ICT Solutions for Brilliant Minds

GÉANT Networks • Services • People

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

TARTU ÜLIKOOL 1632 UNIVERSITAS TARTUENSIS

GHENT UNIVERSITY

NORDUnet Nordic Gateway for Research & Education

Architecture

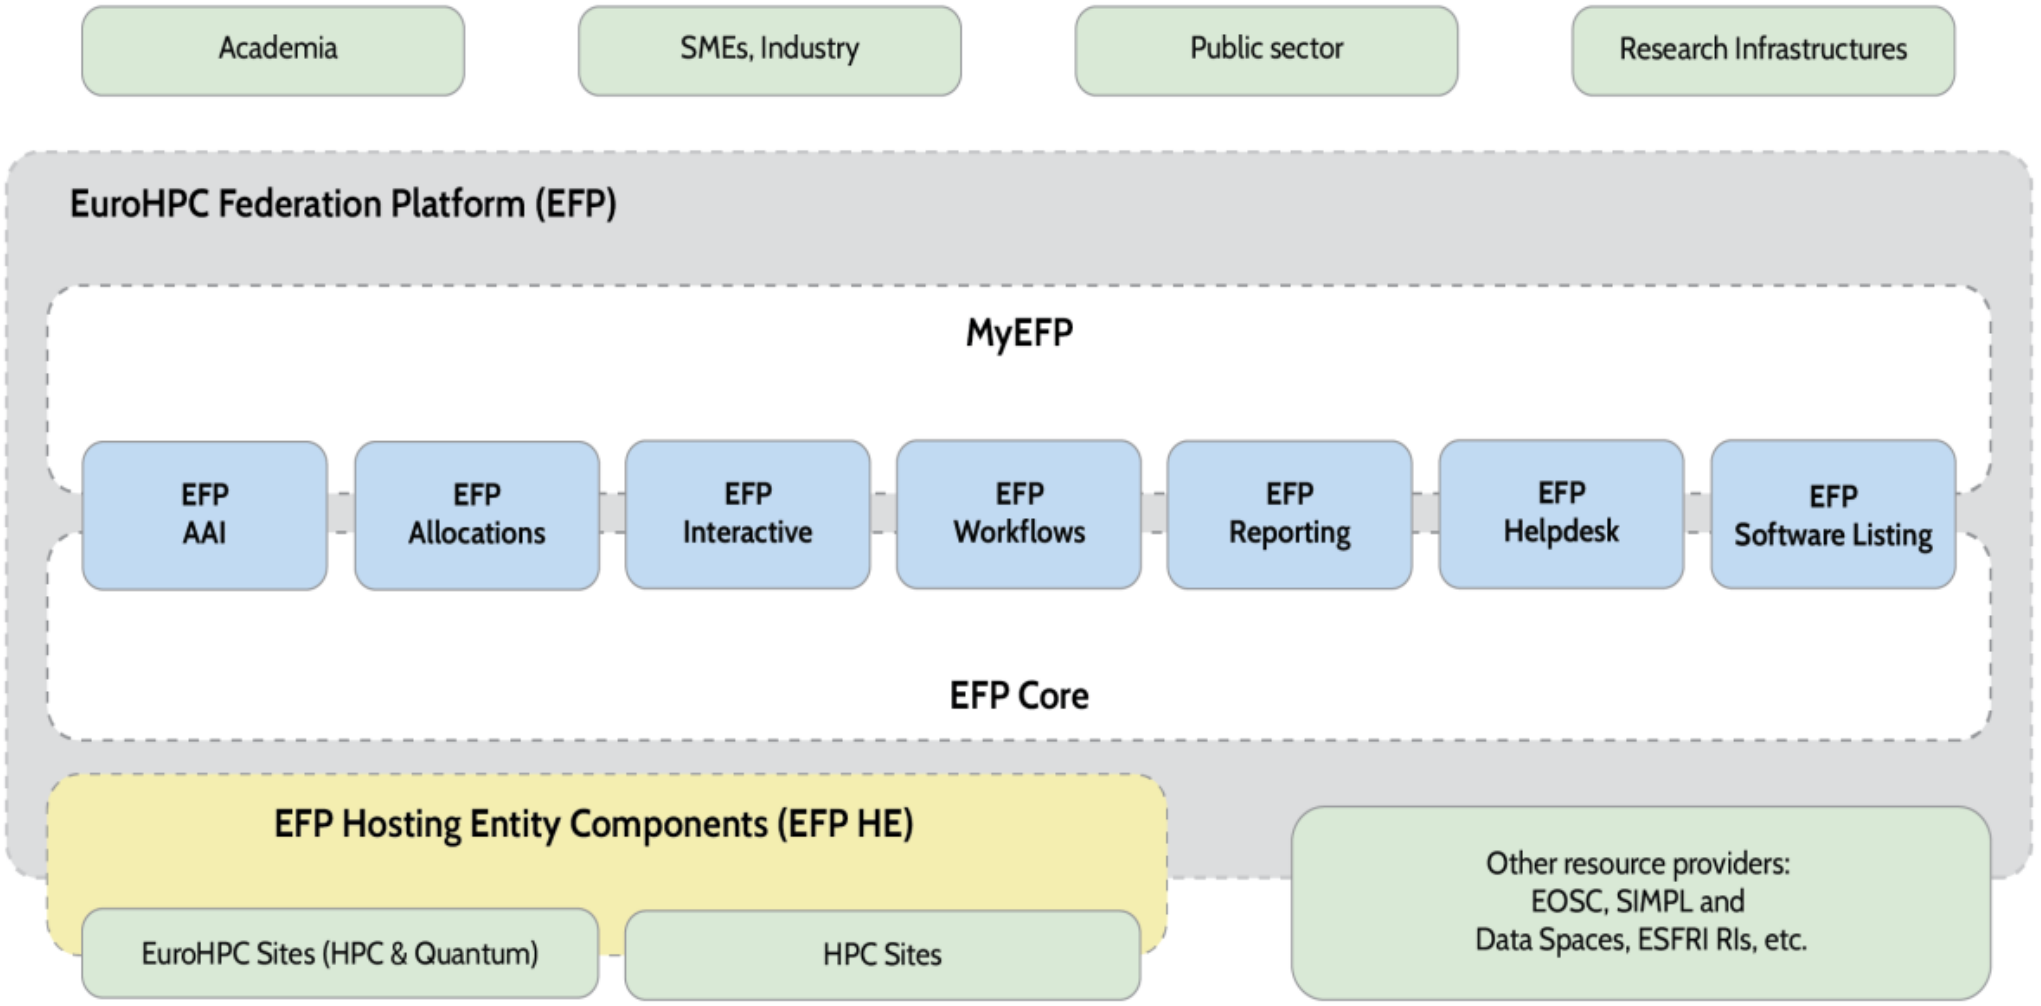

Modularity, Flexibility and API centric are among the core design principles, along with being as non-intrusive as possible

Main components

- The major components of the platform are all based on **open-source technologies** which are already in production use on several systems and have active communities.

Architecture

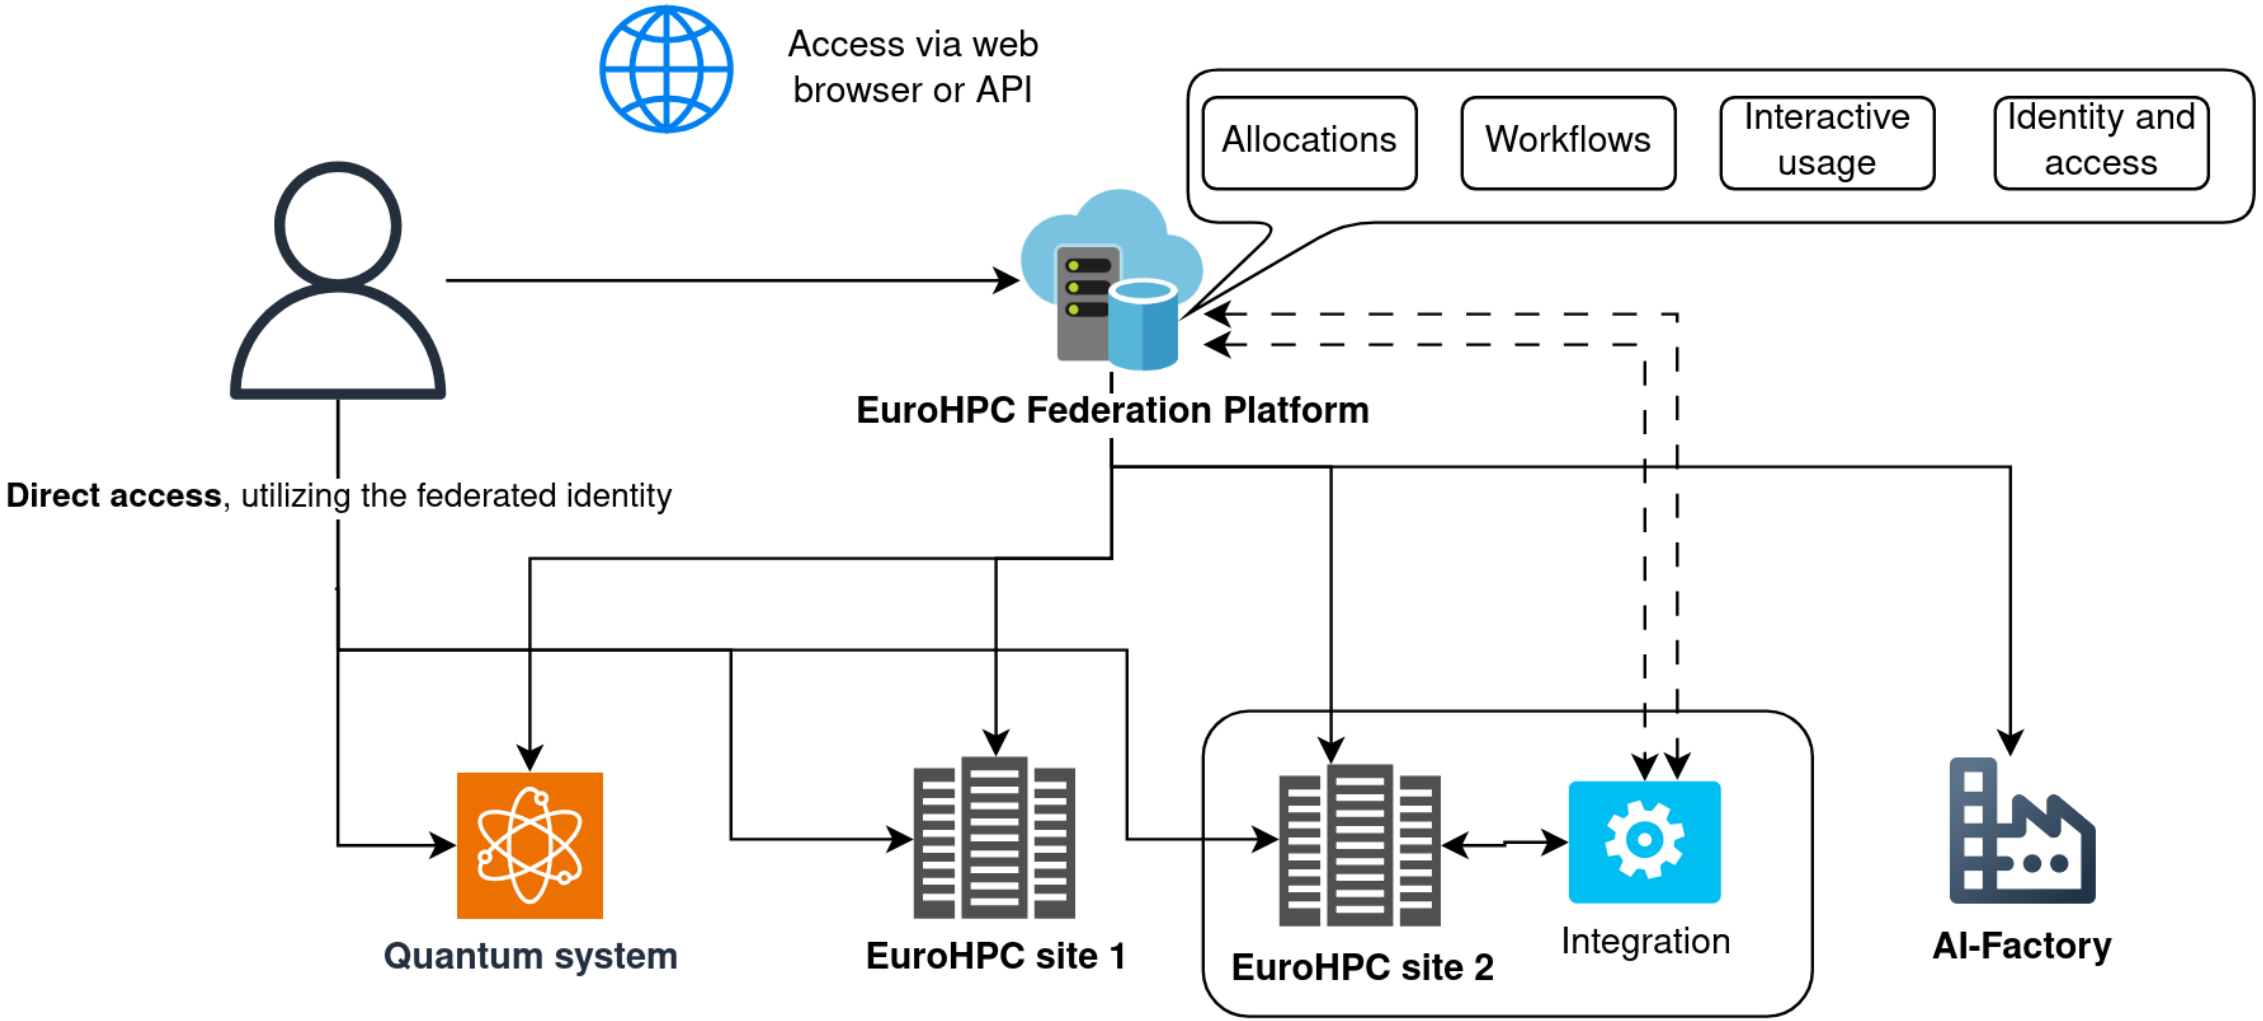

Main components powered by open source software

- EuroFP AAI -> Leveraging MyAccessID,https://wiki.geant.org/display/MyAccessID SSH CA, Infrastructure Proxy
- EuroFP Allocations -> Waldur https://waldur.com/
- EuroFP Interactive -> Open OnDemand 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
- 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 Certificate authority for direct access to the systems

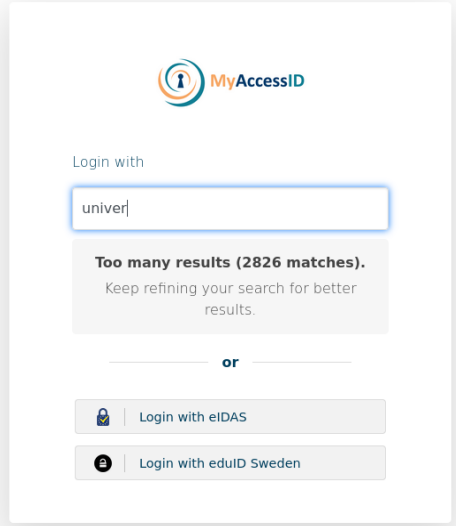

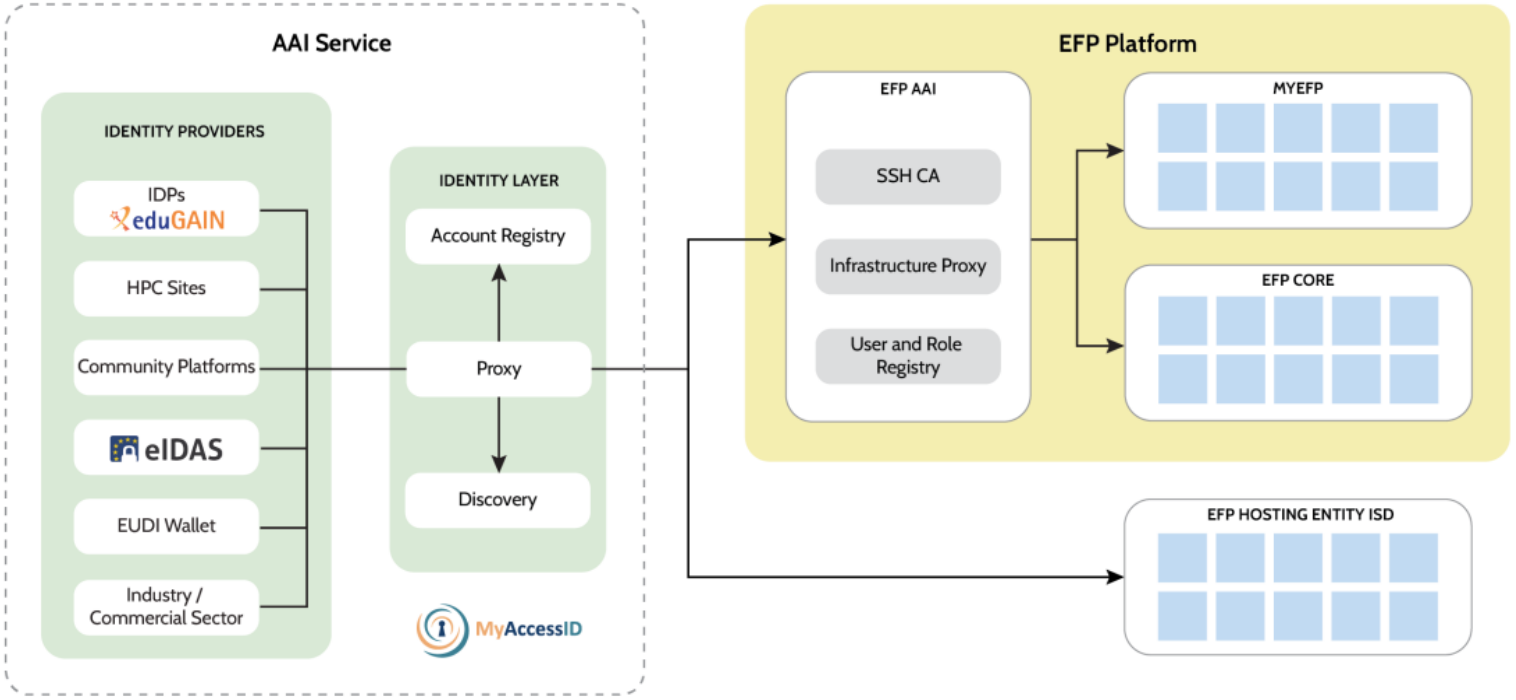

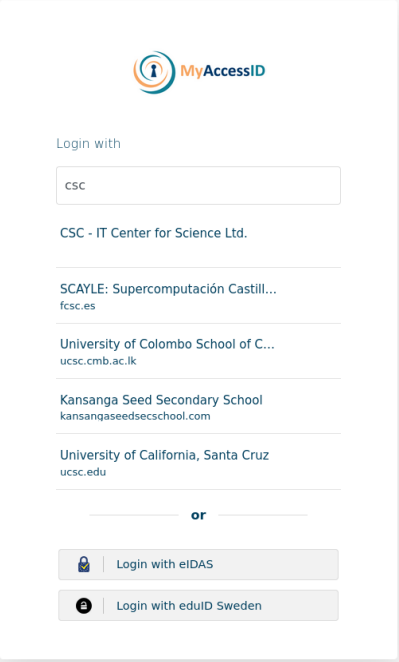

EFP Allocations -> WALDUR

- 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.

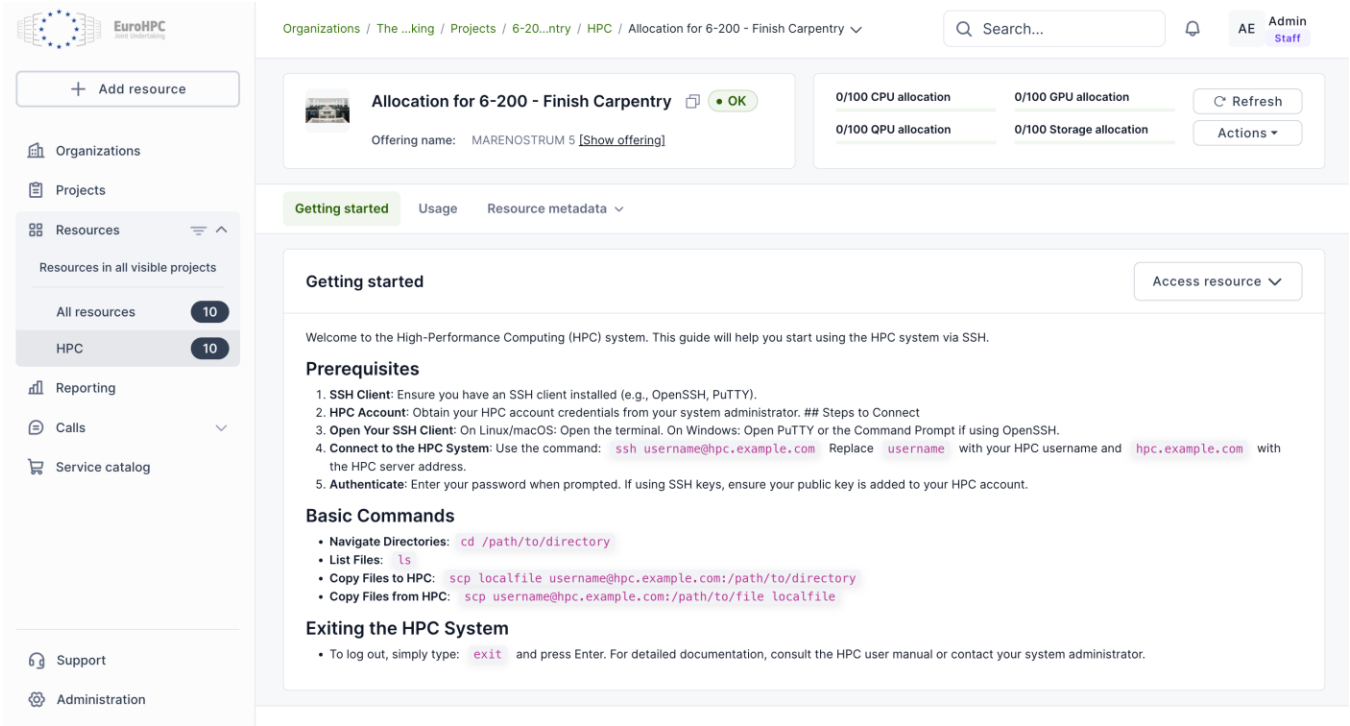

EFP Workflows ->

- 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

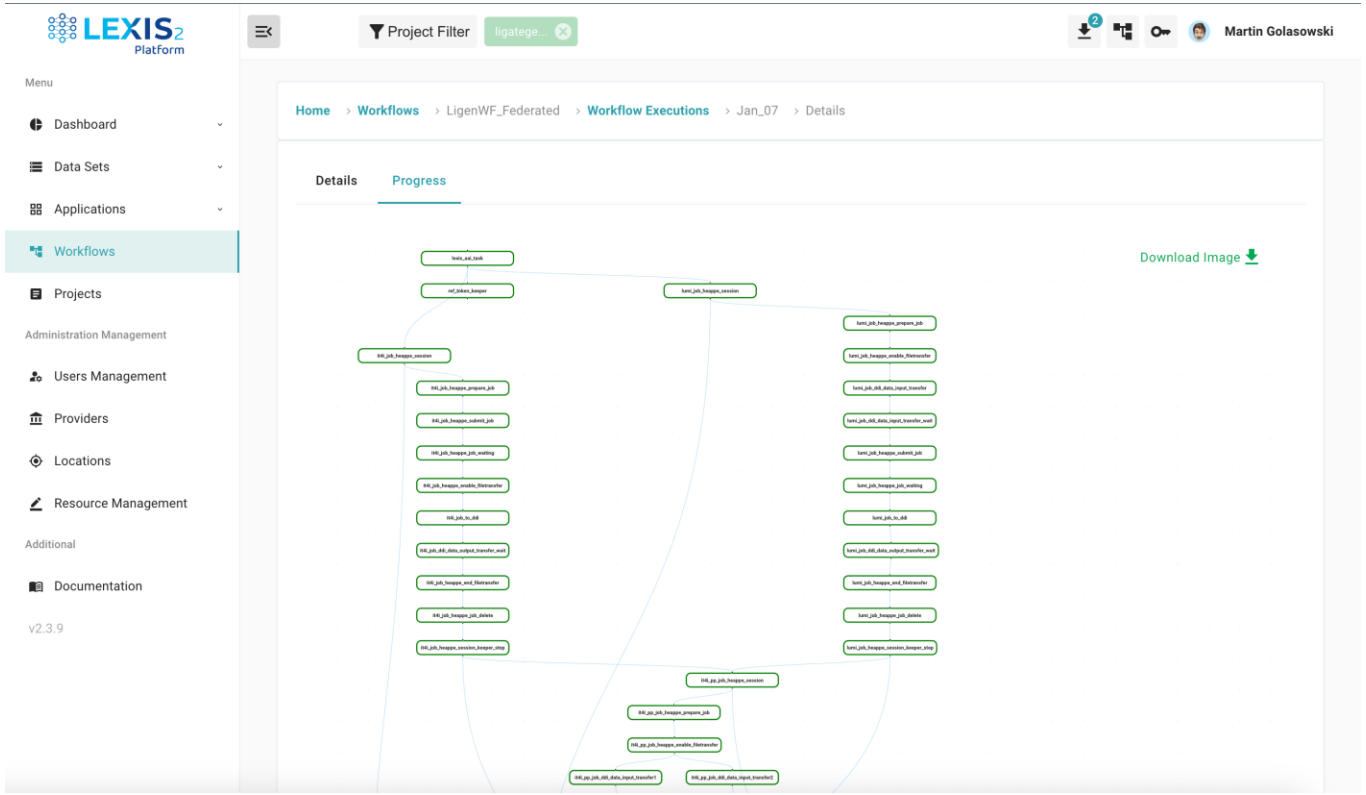

- HEAppE Middleware: REST API for restricted access to HPC infrastructure

EFP Interactive -> OPEN OnDemand

- 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

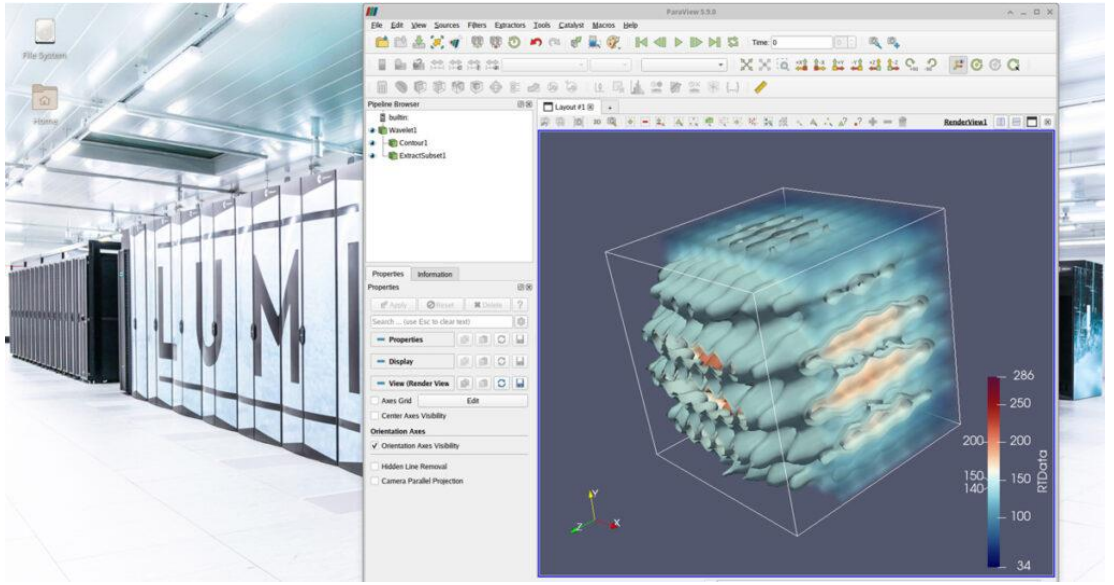

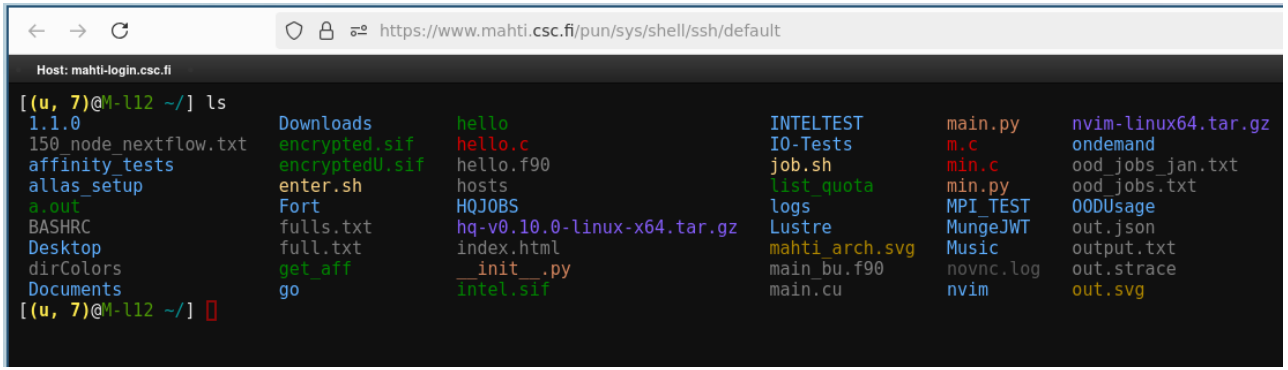

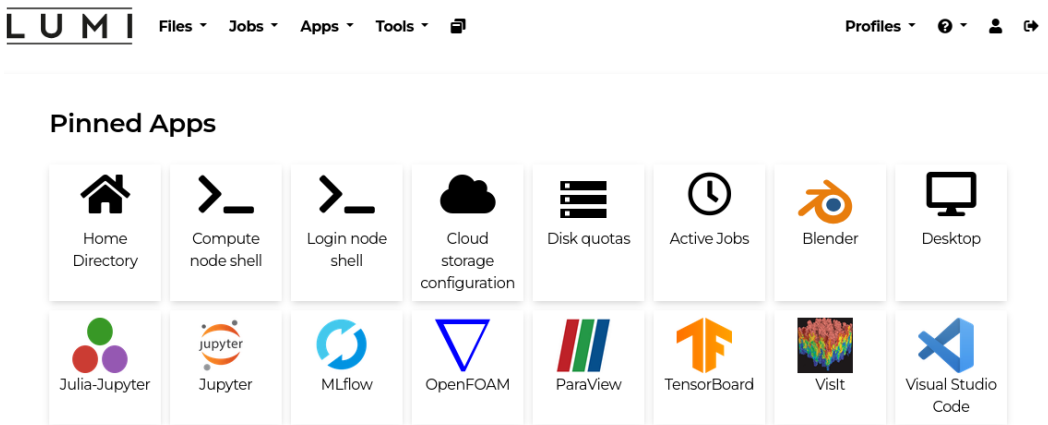

EFP Software Catalogue

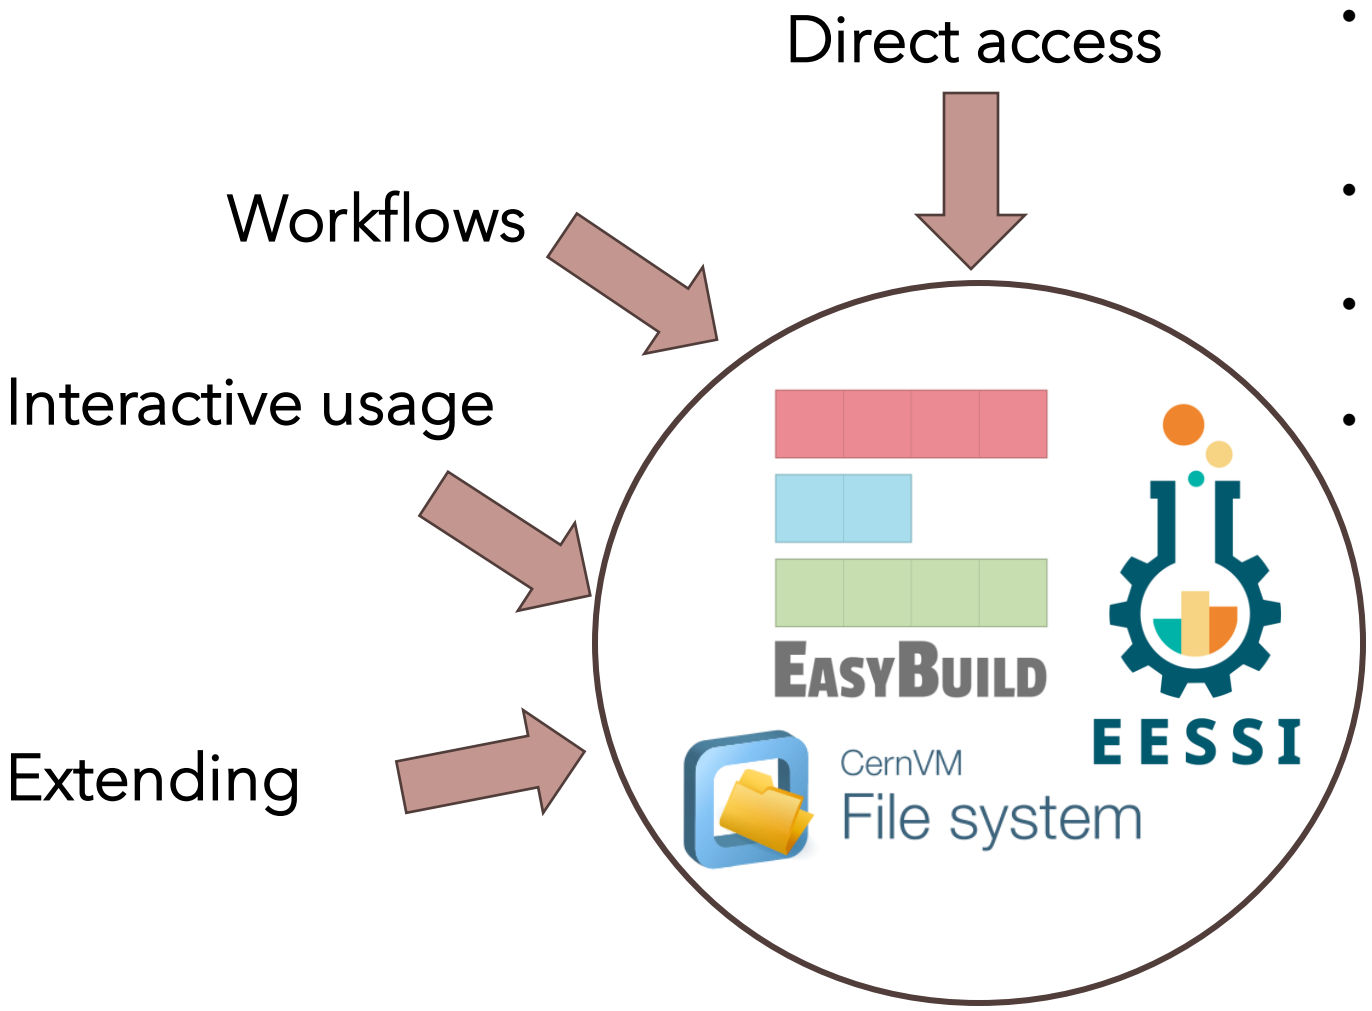

- Federated Software Catalogue based on EESSI
- EFP Software Listing provides list of the available software,
- Ideally accessible on nodes via native installation of CernVM-FS,
  but alternative approaches are possible
- Not meant to replace site provided software stacks
- Scalability of support and user self-service

- Save time for both users and support

EFP Software Catalogue

EESSI Wishlist (most of these are being worked on)
- Robust GPU support for both AMD and NVIDIA
- Possible support for accelerated visualization via e.g. EGL
- Slingshot support

Thinking about:
- Best way to leverage EasyBuild and EESSI for other environments that will be available in future systems, e.g IAAS, Kubernetes, etc.

More information

- https://eurohpc-ju.europa.eu/index_en
- https://eurohpc-ju.europa.eu/paving-way-eurohpc-federation-platform-2024-12-19_en
- Talk with me

Questions ?