Project Title	Hybrid Eco Responsible Optimized European Solution
Project Acronym	HEROES
Grant Agreement No.	956874
Start Date of Project	01.03.2021
Duration of Project	24 Months
Project Website	heroes-project.eu

D3.6 – Final Prototype

Work Package	**WP 3, Infrastructure as a Service**
Lead Author (Org)	**Benjamin Depardon (UCit)**
Contributing Author(s) (Org)	Davide Pastorino (Do IT Systems), Julita Corbalan (BSC), Elisabeth Ortega (HPC Now)
Reviewed by	Sablin Amon (UCit), Diego Bacchin (Do IT Systems)
Approved by	**Technical Board**
Due Date	**28.02.2023**
Date	**28.04.2023**
Version	**V1.3**

Dissemination Level

X	PU: Public
	PP: Restricted to other programme participants (including the Commission)
	RE: Restricted to a group specified by the consortium (including the Commission)
	CO: Confidential, only for members of the consortium (including the Commission)

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

Versioning and contribution history

Version	Date	Author	Notes
0.1	26.01.2023	Benjamin Depardon (UCit)	V0.1
0.2	23.01.2023	Davide Pastorino (Do IT Systems)	Added requirements
0.3	23.01.2023	Julita Corbalan (BSC)	Added news about EAR
0.4	24.03.2023	Benjamin Depardon (UCit)	Notes and todos
0.5	03.04.2023	Davide Pastorino (Do IT Systems)	Updates
0.6	06.04.2023	Benjamin Depardon (UCit)	Prototype presentation, references, videos
0.7	08.04.2023	Davide Pastorino (Do IT Systems)	Updated future works
0.8	11.04.2023	Benjamin Depardon (UCit)	Video links
0.9	20.04.2023	Benjamin Depardon (UCit)	Summary
1.0	21.04.2023	Julita Corbalan (BSC)	Added metrics on EAR
1.1	27.04.2023	Elisabeth Ortega (HPCNow!)	Updates on web UI
1.2	28.04.2023	Benjamin Depardon (UCit)	Final version
1.3	28.04.2023	Corentin Lefevre (Neovia)	Final edition for submission

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

Table of Contents

Versioning and contribution history 2
Table of Contents 3
List of Figures 3
TERMINOLOGY 4
Executive Summary 5
1 Introduction 6
2 HEROES Prototype 8
3 Requirements 13
3.1.1 HEROES Prototype Infrastructure deployment 13
3.1.2 Enrolment of existing HPC Clusters 13
4 Latest updates 14
4.1.1 HEROES Infrastructure 14
4.1.2 Migration to cloud 14
4.1.3 HELM chart for the HEROES Platform 15
4.1.4 Enhanced security and exposure of public endpoints 15
4.1.5 Integration of cloud and on premise HPC clusters 15
4.1.6 Web interface 15
4.1.7 EAR 16
5 Summary 19
6 Future work 20
References 22
List of Annexes 23

List of Figures

FIGURE 1. USER INTERACTION WITH THE HEROES PLATFORM 8
FIGURE 2. FINAL ARCHITECTURE 9
FIGURE 3 ENERGY SAVINGS EVALUATION WHEN RUNNING EAR MIN_ENERGY POLICY COMPARED WITH TURBO. CPU FREQUENCY OPTIMIZATION. 18
FIGURE 4 ENERGY SAVINGS EVALUATION WHEN RUNNING EAR MIN_ENERGY POLICY COMPARED WITH TURBO. CPU FREQUENCY OPTIMIZATION FOR GPU APPLICATIONS 18

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

TERMINOLOGY

Terminology/Acronym	Description
AI	Artificial Intelligence
API	Application Programming Interface
AWS	Amazon Web Services
CAE	Computer Aided Engineering
CPU	Central Processing Unit
CUDA	Compute Unified Device Architecture
EAR	Energy Aware Runtime software
GPU	Graphical Processing Unit
HPC	High Performance Computing
I/O	Input/Output
ML	Machine Learning
MPI	Message Passing Interface
RAM	Random Access Memory
TDP	Thermal Design Power
UI	User Interface
VM	Virtual Machine
VNC	Virtual Network Computing

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

Executive Summary

The HEROES Project is aiming at developing an innovative European software solution allowing industrial and scientific user communities to easily submit complex Simulation and ML (Machine Learning) workflows to HPC (High Performance Computing) Data Centres and Cloud Infrastructures. It will allow them to take informed decisions and select the best platform to achieve their goals on time, within budget and with the best energy efficiency.

This deliverable presents the HEROES prototype integrating all the different modules created by work packages 2, 3 and 4. We present the final version of the architecture and the prototype, the requirements needed to deploy a HEROES platform, and some updates on the various modules already presented in previous deliverables.

This document in accompanied by a set of videos in annex that present in more details the HEROES prototype.

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

1 Introduction

The purpose of this document is to present the final state of the HEROES platform prototype.

The objective of the prototype is to implement the HEROES architecture initially presented in deliverable "D3.1 – Architecture Design" and updated in subsequent deliverables. This architecture is built around a set of modules that provide the actions required by the end-users to submit and manage their computational workflows, and administrators to run the platform:

- Identity Management
    - The capacity of the platform to provide an authentication method to the Identity Management module from the UI/API (User Interface / Application Programme Interface) module for users registered within organization of the platform.
    - The capacity to provide basic authorizations to the users of the platform corresponding to their organization and roles.
- Data Management
    - The capacity of the platform to provide a data management method, including the possibility to upload, download and move data within the platform by the authenticated and authorized users.
- Workflow and Job Management
    - The capacity to manage, submit, monitor and any other basic action by the Workflow and Job Management module depending on the authenticated user requests.
- Platform Administration
    - The capacity to manage organizations (e.g.: creation, deletion) and users within the organizations by the authenticated and privileged users through the UI/API module.
- Decision module
    - The capacity to automatically decide on which target platform a workflow should be executed.
    - Job submission parameters optimization to request only the resources needed by the jobs and optimize resource usage at the cluster level.
- Energy monitoring and optimization
    - Energy metrics gathering at the job and node level.
    - Energy optimization to adapt to the applications' performance needs while reducing power consumption.
- End-user interface

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

- Web interface to allow end-users to easily interact with the platform.

The details of these actions and associated modules have been presented in the previous deliverables, more specifically:

- D2.1 – Definition of CAE Workflow and AI Workflow
- D2.2 – Workflows Containers
- D2.3 – Workflow Interfaces
- D3.1 – Architecture Design
- D3.2 – HEROES Deployment Suite
- D3.3 – Initial Platform Mockup
- D3.4 – Best Practice Library
- D3.5 – Cost Service
- D4.1 – Updated energy aware runtime
- D4.2 – Decision module

In this document, we will present the updates on the final architecture and prototype compared to the previous deliverables in Section 2. In Section 3 we present the requirements necessary to deploy a HEROES platform. We then give the latest updates on some of the modules in Section 4, before concluding in Section 5 and providing some ideas for future works in Section 6.

Along with this document, a set of videos have been created to present in more details the different parts of the architecture and their integration. The videos should be seen in the following order:

1. HEROES Backend and APIs – https://youtu.be/_pQ4myKcrLI
2. HEROES Infrastructure – https://youtu.be/TKyQx5CnzKM
3. HEROES Workflow and Containerization – https://youtu.be/w17BeklNMKM
4. HEROES Energy Aware Runtime (EAR) – https://youtu.be/UfgTrvU6Jjg
5. HEROES Decision Module – https://youtu.be/0Z9PVvIAJjE
6. HEROES Web User Interface – https://youtu.be/aGQ4vy_-XpQ

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

2 HEROES Prototype

The HEROES prototype aims at providing end-users tools to easily submit their computational workflows on a set of available clusters. Access to the platform is done through a web portal plugged onto web services APIs, allowing any authorized user and software tool to run a workflow on selected target platforms. Both workflow execution and data management are handled by the solution in a transparent manner. The platform also embeds a decision module that helps users in the selection of the computing platform of choice (HPC centre or Public Cloud) and to submit their jobs more easily – ideally in a transparent way – while taking into account their constraints (performance, energy and costs...). Figure 1 presents the user workflow (blue arrows) and the associated actions and modules (inside the green rectangle and arrows) involved in the process.

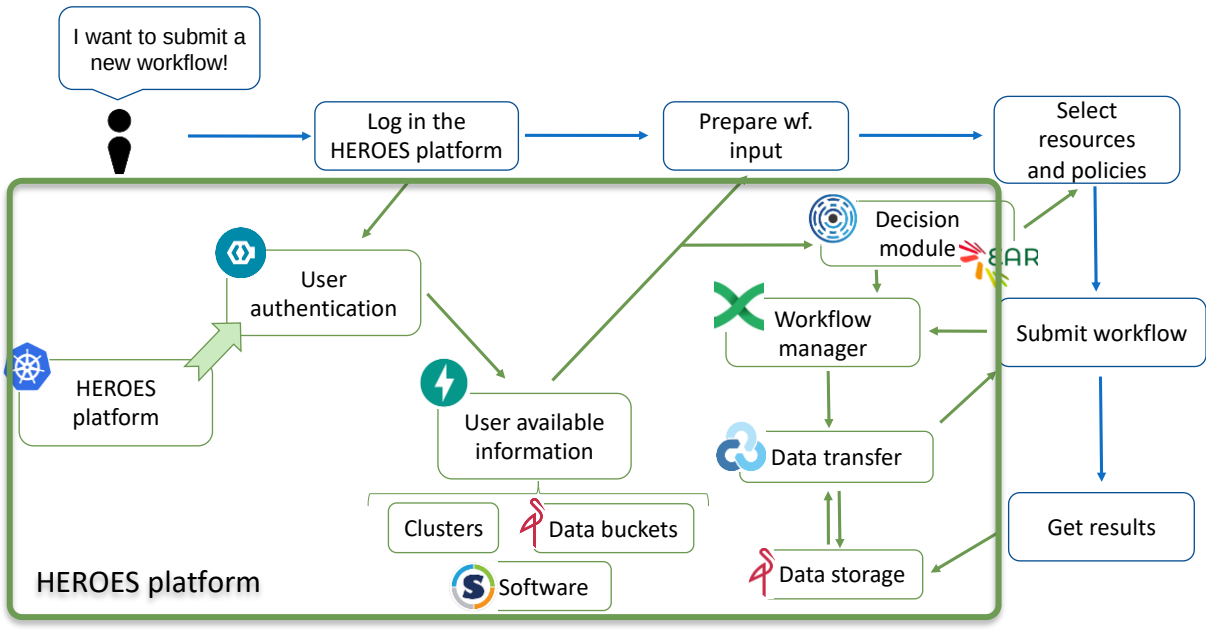

Figure 1. User interaction with the HEROES platform

The HEROES architecture is depicted in Figure 2. Its design is very close to the one presented in Deliverable D3.3 Initial Platform Mockup. The main differences are:

- The user interfaces implemented with ReactJS (see D2.3 Workflow Interfaces).
- The use of CloudSHaper for the cost module (see D3.5 Cost Service).
- The use of MinIO as a registry to store the workflow containers and definition files.
- The use of Elasticsearch [1] in the Telemetry module.
- Greater integration with RabbitMQ [14] to decouple user interactions from platform actions.

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

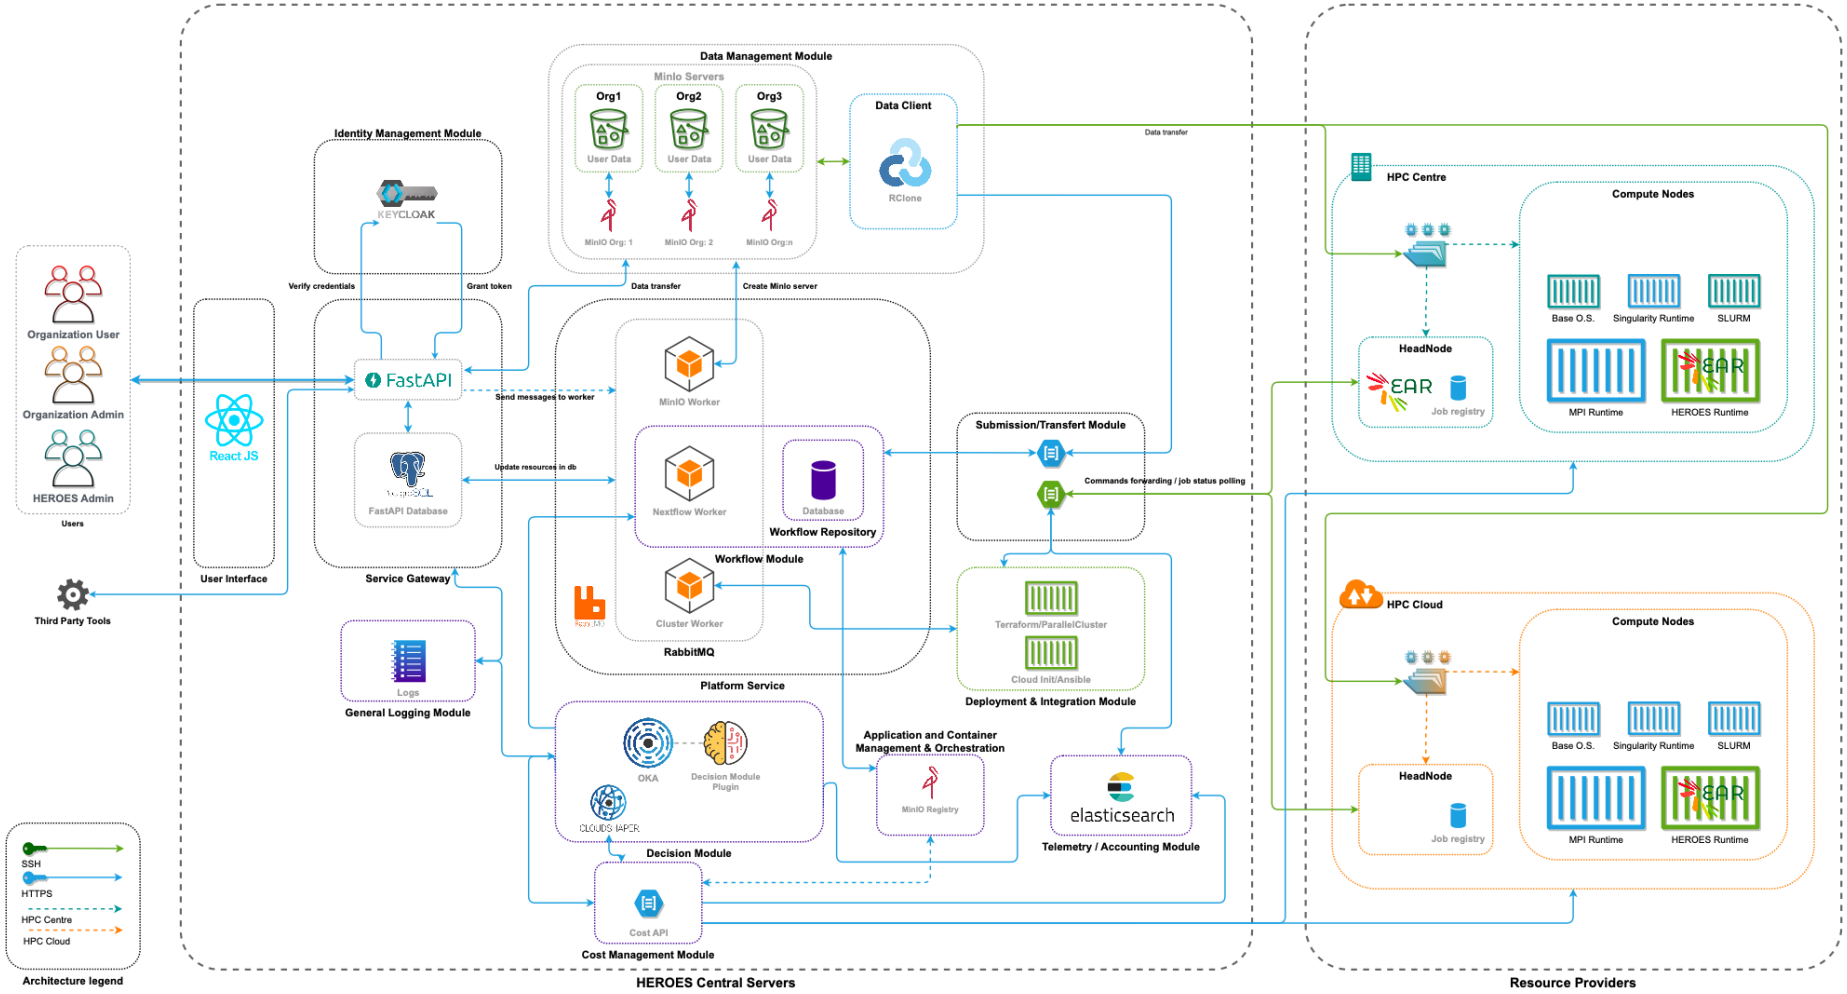

Figure 2. Final architecture

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

The APIs have been extended to finalize the prototype. Without going into the details of all of them, we present below the list of all the available APIs. The full documentation is available in the Annex HEROES API.html.

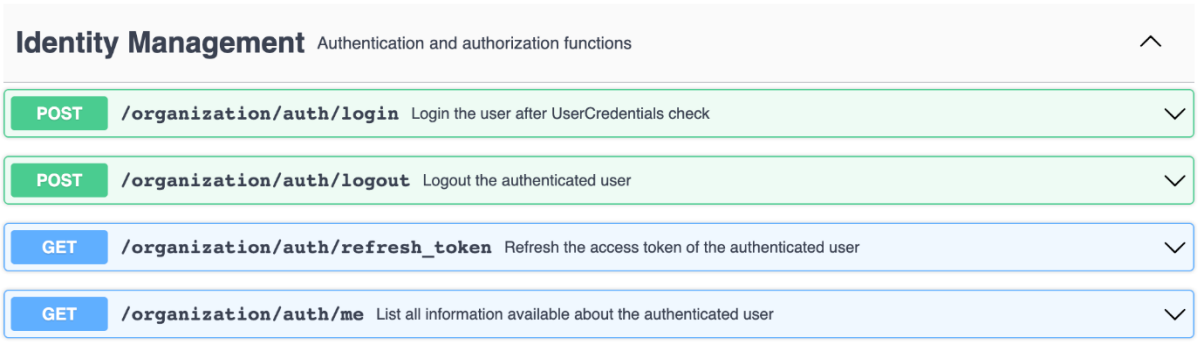

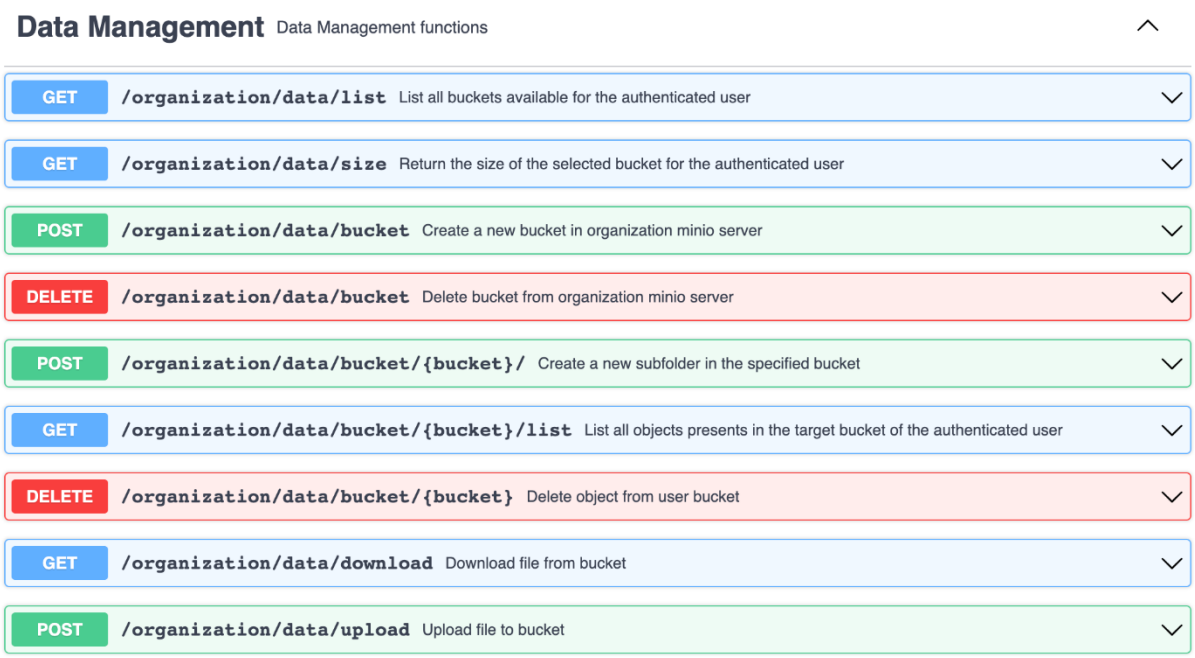

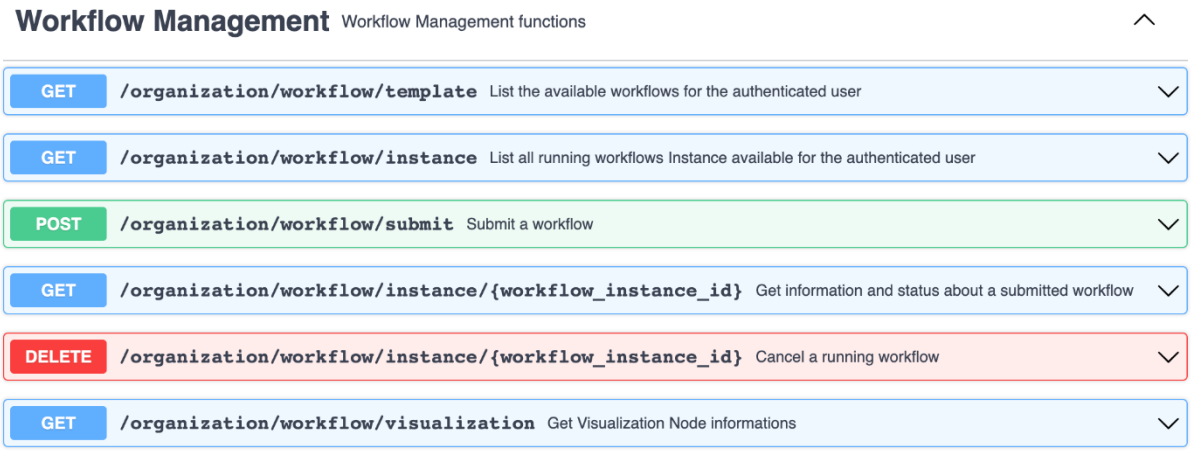

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

Organization Admin Organization Administration functions

GET	/organization/admin/users/	List users
POST	/organization/admin/users/	Create new user
GET	/organization/admin/users/{username}	Get user information
DELETE	/organization/admin/users/{username}	Delete user
PATCH	/organization/admin/users/{username}	Update user
GET	/organization/admin/groups/	List groups
POST	/organization/admin/groups/	Create group
GET	/organization/admin/groups/{groupname}	Get group information
DELETE	/organization/admin/groups/{groupname}	Delete group
PATCH	/organization/admin/groups/{groupname}	Update group
GET	/organization/admin/roles/	List roles
POST	/organization/admin/roles/	Add role to user
GET	/organization/admin/roles/{rolename}	Get role information
DELETE	/organization/admin/roles/{rolename}	Delete role
GET	/organization/admin/cluster/	List Clusters
POST	/organization/admin/cluster/	Create Cluster
GET	/organization/admin/cluster/{cluster_id}	Get cluster information
DELETE	/organization/admin/cluster/{cluster_id}	Delete cluster
POST	/organization/admin/workflow/template	Create workflow template
DELETE	/organization/admin/workflow/template/{workflow_template_id}	Delete cluster
GET	/organization/admin/workflow/instance/{workflow_template_id}	List all running workflows Instance available for the organization, can be filtered on workflow template id

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

HEROES Admin Platform Administration functions

- GET `/heroes/admin/organizations/` List Organizations
- POST `/heroes/admin/organizations/` Create a new organization
- DELETE `/heroes/admin/organizations/` Delete organization
- GET `/heroes/admin/organizations/{organization_target}` Read Organization
- GET `/heroes/admin/organizations/{organization_id}/clients/` Read Clients
- POST `/heroes/admin/organizations/{organization_id}/clients/` Create Client For Organization
- GET `/heroes/admin/organizations/{organization_id}/servers/` List Organization Servers
- POST `/heroes/admin/organizations/{organization_id}/servers/` Create a new organization server
- DELETE `/heroes/admin/organizations/{organization_id}/servers/` Delete organization

Decision Module

- POST `/organization/decision/decide` Decide
- POST `/organization/decision/ingest_ear_files` Ingest Ear Files
- GET `/organization/decision/policies` Get Policies

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

3 Requirements

The purpose of this section is to describe the requirements for the deployment of the HEROES solution components for the prototype.

This paragraph will outline the requirements for:

- The deployment of the HEROES prototype infrastructure.
- The enrolment of an existing HPC Cluster within the HEROES prototype infrastructure.

3.1.1 HEROES Prototype Infrastructure deployment

To deploy the current revision of the HEROES prototype, the basic requirements are:

1. An account on AWS.
2. A PC with Kubespray [3] code downloaded locally.
3. The Terraform [4] CLI and the HELM [9] binary, are both installed on a VM or on a computer with administrative access to the cloud account.
4. The HEROES HELM chart, downloaded from the main HEROES repository.
5. Internet access from the cloud private network.

Alternatively, if a Kubernetes cluster is already available, the requirements are:

1. The HELM binary, on a computer with administrative access to the Kubernetes cluster.
2. The HEROES HELM chart is downloaded from the main HEROES repository.
3. Internet access from the Kubernetes cluster.

The results of running the Helm charts will be:

- the Kubernetes [5] cluster and related VMs.
- all HEROES infrastructural microservices active and ready.

At the end of the deployment, the architecture described in the previous paragraphs will be available in the cloud platform account. The process will provide output details about the deployed VMs, and the bastion host used to access them.

3.1.2 Enrolment of existing HPC Clusters

The process to enrol an existing HPC Cluster has been designed to have the least possible requirements. However, it will still require the following:

- SLURM [6] as a workload scheduler.
- a user account on the HPC cluster that can submit and run batch jobs.
- the possibility to connect to the HPC cluster with SSH and SFTP, using key-based authentication.

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

- a shared storage space for the user account that can be accessed from all compute nodes (e.g., HOME directory).
- cluster-wide support for Apptainer/Singularity [7] containers.

The enrolment process, at the current status of development, assumes that all cluster-side requirements are already in place at the time of enrolment.

On HEROES, the definition of the new cluster and its related parameters need to be manually entered into the platform configuration.

4 Latest updates

4.1.1 HEROES Infrastructure

During the second year of development the HEROES infrastructure team focused on migrating the HEROES platform from the Do IT Systems Laboratory to the cloud, using AWS as the cloud platform of choice.

Additionally, the tool we selected to provide backend monitoring and logging facilities for the HEROES platform, OpenSearch [2], ended up being more difficult to integrate than expected.

Therefore, the team proceeded to do a new cycle of tool selection; in the end the following tools were chosen to support the new deployment approach:

- KubeSpray [3], a Kubernetes provisioning tool based on Ansible [8].
- HELM [9], a tool to manage and deploy applications on a Kubernetes cluster.
- Prometheus [10], a monitoring tool that provides out-of-the-box monitoring for Kubernetes.
- Grafana [11], a dashboard web application that can integrate and visualize multiple sources.
- Loki [12], a logging component for the Grafana/Prometheus stack.

The other tools used for the laboratory (Rook [13], RabbitMQ [14], Keycloak [15], Argo CD [16], Gitlab [17]) were successfully integrated with the new infrastructure, providing the previously available interfaces to the HEROES development team as well as some new features.

4.1.2 Migration to cloud

The Do IT Systems laboratory environment was manually installed on bare metal servers. The migration to cloud required a new approach, and we wanted to be as close to “cloud native” as possible. To achieve the desired result on AWS we developed a new deployment template based on Terraform and Ansible, using KubeSpray as a basic templated and expanding on it based on the project’s needs.

Additional work was also performed to ensure enhanced data reliability and replication between the nodes of the Kubernetes cluster.

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

This resulted in a “production ready” Kubernetes infrastructure on which the HEROES platform could be reliably hosted.

4.1.3 HELM chart for the HEROES Platform

Once the new infrastructure was ready, a HELM chart was developed to perform the full deployment of all components and applications of the HEROES Platform, including basic configuration and integration tasks.

This HELM chart enables a consistent and reliable deployment of the HEROES Platform on a Kubernetes cluster.

4.1.4 Enhanced security and exposure of public endpoints

A Web Application Firewall was added to the HEROES Service Gateway to enhance the overall platform security.

Two public endpoints were exposed on the internet:

- https://portal.heroes-project.eu for the WebUI (Web User Interface).
- https://api.heroes-project.eu for API access.

Both these endpoints require a valid account on the HEROES platform as well as an authenticated session.

4.1.5 Integration of cloud and on premise HPC clusters

Two small-scale HPC clusters were integrated in the Final Prototype:

- a cluster hosted on the HEROES account on AWS.
- a bare-metal cluster of 3 nodes hosted by HPCNow! in their own data centre.

Users can submit workflows on both clusters from the HEROES WebUI.

4.1.6 Web interface

User interaction was adopted by the HEROES platform via the creation of a web interface. The initial idea was to use Open OnDemand as a portal to the HEROES platform because it is already known by the HPC community, as it was presented in the first reporting period. However, Open OnDemand technology cannot be containerized and deployed alongside the other pieces of the HEROES platform. To solve this issue, HEROES team developed a web interface from scratch using ReactJS and including the basic functionalities at user level to submit a workflow:

- Log in form.
- Data visualization and management.
- Workflow lists.
- Workflow creation and execution.
- Visualization of results in a remote VNC server.

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

4.1.7 EAR

In the last public version of EAR included in the BSC GitLab repository [1] there is full support for HEROES runtime. Some improvements since deliverable D4.1 have been done in the topics of EAR compilation and deployment, HEROES plugin and dynamic EAR library initialization.

EAR compilation and deployment has been improved to be fully automatic without any modification in Makefiles so it can be added in ansible scripts used in HEROES. The EAR-lite feature can be just enabled by adding `FEAT_EAR_LITE=1` at compile time. We decided to introduce this feature only on demand to avoid unnecessary runtime checks done by the library in scenarios where EAR is installed and then it's not needed.

The HEROES plugin has been updated and extended. It has been updated to use the final names for environment variables. New names are:

- `HEROES_TEMPLATE_WORKFLOW_ID`: the identifier of the workflow template used to instantiate (submit) the workflow.
- `HEROES_INSTANCE_WORKFLOW_ID`: the identifier of the instantiated (submitted) workflow.
- `HEROES_USER_ID`: the identifier of the end-user who has submitted the workflow.
- `HEROES_ORGANIZATION_NAME`: the name of the organization the end-user belongs to.
- `HEROES_ORGANIZATION_ID`: the identifier of organization the end-user belongs to.

Moreover, we have extended the information reported by the plugin with the elapsed time on the different phases the EAR library has detected. This information enriches the performance metrics reported by EAR because it combines multiple metrics to decide if an application is in a computational phase, I/O phase etc. This information is used by the decision module as an important hint for resource selection.

The last and main improvement has been in the extension of the metrics estimation when some of them are not available. In the last months we have extended from a basic detection and configuration of the available features to the estimation of metrics based on the ones available. In more detail:

- We have extended the library to estimate the memory bandwidth using last level cache events (which typically doesn't require any type of privilege to be measured).
- In the case of the average CPU frequency, in case it is not available, we estimate it is the same than the selected CPU frequency (which makes sense in most of the cases).
- The memory frequency, current and average, is a complicated metric and we don't have any strategy to measure the current one (or even the default one). We consider for future versions to use the same strategy than for the CPU TDP. In that case, we decided to create and progressively extend a static list with known CPU models and use an environment variable for new CPUs.

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

- In the case of GPU metrics, in particular NVIDIA GPUs, it's not required to have privileges to be able to get GPU metrics. In a full EAR installation, the EARD is in charge of accumulating GPU metrics to be able to provide average metrics over a period and not instantaneous metrics (as reported by NVIDIA API). What we have included in EAR-lite is the computation of aggregated metrics done by the library directly.

Additionally, to all these estimations, the EAR disables those capabilities, such as energy optimization, when the required hardware knobs are not available. In particular, we switch the optimization policy to monitoring in case the CPU frequency management API is not available. We disable the memory frequency optimization when the memory management API is not available, and we disable GPU optimization policy if GPU management policy is not available.

- When EAR is installed in bare metal systems, we have estimated an average energy reduction of 10%. We have executed some applications to evaluate the energy reduction that can be expected when executing individual applications and/or workloads. Figure 3shows the energy saving, power saving and GFlops/W gain when running some applications with EAR and min_energy policy compared with legacy TURBO. These experiments have been done in a cluster with 2 x Intel® Xeon® Platinum 8360Y @2.4GHz 36c per node. Results are the average of three runs. Applications executed are POP (384 processes)[19], HPCG (64 processes)[22], DUMSES (512 processes)[20], and AFiD(384 processes)[23]. These applications are mostly memory bound, reporting a clear benefit by reducing the CPU frequency. Table 1 shows the CPI (cycles per instructions) and GB/s (Giga Bytes per second) corresponding with the applications evaluated. These metrics (CPI and GB/s) are two of the main inputs used by EAR to characterize the application phases. The last column shows the average CPU frequency selected when executed with min_energy optimization policy.

Table 1 Basic metrics for application characterization

Application	CPI (Cycles per instructions	GB/sec	Avg. CPU frequency
POP	0,91	222	2,23 GHz
HPCG	3,01	301	1,94 GHz
AFiD	0,81	188	2,25 GHz
DUMSES	0,99	195	2,19 GHz

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

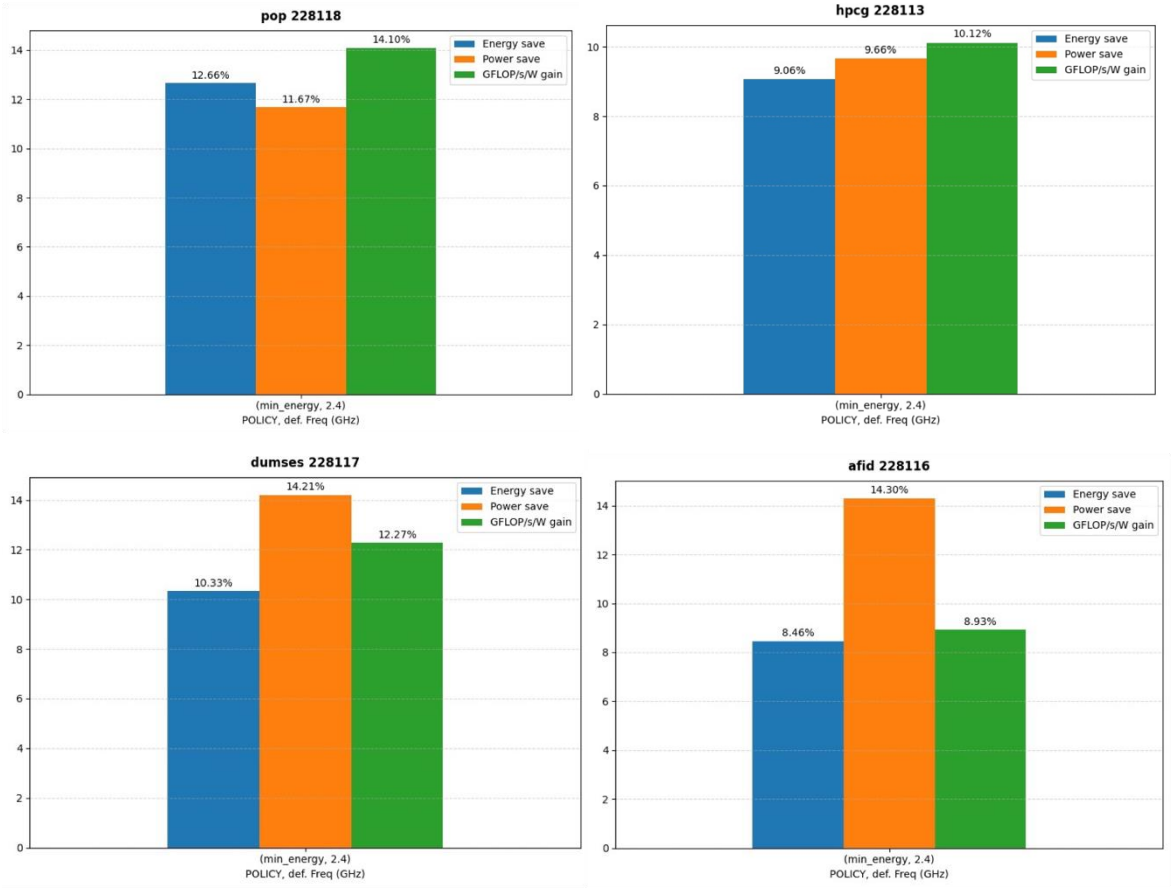

Figure 3 Energy savings evaluation when running EAR min_energy policy compared with TURBO. CPU frequency optimization.

Figure 4 shows same metrics when executing BT-CUDA kernel[21]. In that case, the source of the energy saving is the reduction of the CPU frequency when the CPU is idle waiting the GPU computation.

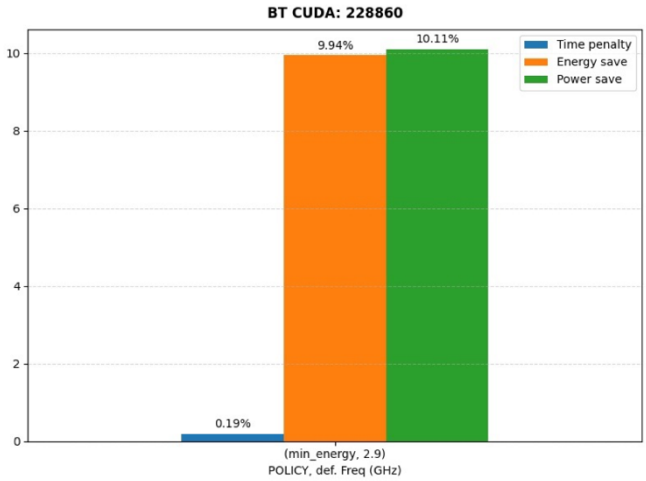

Figure 4 Energy savings evaluation when running EAR min_energy policy compared with TURBO. CPU frequency optimization for GPU applications

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

5 Summary

In this document we presented the final prototype of the HEROES Platform. It aggregates the work conducted during the 2 years of the project. The prototype is the integration of all the modules composing the HEROES platform, that had initially been designed in "D3.1 Architecture Design".

The prototype allows the following actions:

- For an end-user:
  - Connect to the platform.
  - Manage their data (input/output) in a dedicated object store (both per user and organization).
  - Submit and manage workflows, while being guided in the submission process to select the best target platform considering the user's constraints on performance, costs and energy.
- For an Organization administrator:
  - Manage users/groups/roles in the Organization.
  - Manage available clusters.
  - Manage available workflow templates.
- For the HEROES administrator:
  - Manage Organizations.
  - Administrate the HEROES infrastructure.
  - Monitor, and create reporting on the use of the platform.
- For a resource provider
  - Optimize energy consumption with EAR.
  - Provide resources to additional users/organization, through the creation of HEROES Marketplaces.

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

6 Future work

The HEROES project is just the first step towards our view of a "marketplace" for HPC infrastructures. We list below our views on the future work that could be carried on to carry-on the developments of the HEROES platform.

Concerning the general HEROES architecture:

- Take additional security constraints to connect to remote clusters, such as: 1 account per final user on the remote cluster to keep their identity from end-to-end, use of MFA devices to authenticate.
- Improve platform scalability and maintainability.
- Add support for other cloud platforms to the HEROES deployment tools.
- Provide additional support for auto-scaling the HEROES infrastructure resources to match workload.

In the case of OKA and the Decision Module, we envision the following improvements:

- The improvement of the models used in Predict-IT to predict the behaviour of jobs: increase accuracy, simplify and try to automate configuration and training processes as it still requires a lot of human expertise.
- Design algorithms capable of predicting multiple features of a job at the same time (currently one at a time).
- Have pre-trained predictors on well-known applications: allow for reuse even without a long history of existing jobs, use of transfer learning techniques to adapt existing models.
- Take into account additional metrics and constraints in the Decision Module: data localization, specific hardware requirements (e.g., GPU, architecture…).

In the case of EAR, our future plans target 4 directions:

- The enhancement of the CPU power model to improve power predictions when we cannot measure.
- The analysis of requirements when executing new workflows with potentially new requirements. For example, we have seen the requirement of supporting a list of applications to be loaded with EAR. EAR is loaded by default with a list of programming models: MPI. MKL, OpenMP, MKL, CUDA and when python is detected. For the other cases, we were offering one environment variable. However, when running a workflow, it can be not enough to support the specification of only one application name. This extension has been already included (and will be available in the next release), but new ones can appear.

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

- To introduce the concept of workflow budget or workflow optimization in cases where the characteristics of the workflow show benefits in prioritizing some jobs in front of others of the same workflow.
- The enhancement of the learning phase including not only computational kernels but also I/O, GPU etc. This information can be used for application classification and energy optimization but also to be used by the decision module to compare between different clusters.

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

References

All web pages have been visited on 06/04/2023.

[1] ElasticSearch: https://www.elastic.co/

[2] OpenSearch: https://opensearch.org/

[3] Kubespray: https://kubespray.io/

[4] Terraform: https://www.terraform.io/

[5] Kubernetes: https://kubernetes.io/

[6] SLURM: https://slurm.schedmd.com/

[7] Apptainer: https://apptainer.org/

[8] Ansible: https://www.ansible.com/

[9] HELM: https://helm.sh/

[10] Prometheus: https://prometheus.io/

[11] Grafana: https://grafana.com/

[12] Grafana Loki: https://grafana.com/oss/loki/

[13] Rook: https://rook.io/

[14] RabbitMQ: https://www.rabbitmq.com/

[15] Keycloak: https://www.keycloak.org/

[16] Argo CD: https://argo-cd.readthedocs.io/

[17] Gitlab: https://about.gitlab.com/

[18] Energy Aware Runtime Gitlab: https://gitlab.bsc.es/ear_team/ear

[19] Parallel ocean program. [Online]. Available: http://www.cesm.ucar.edu/models/ccsm4.0/pop

[20] The DUMSES website. [Online]. Available: https://github.com/marcjoos-phd/dumses-hybrid

[21] NPB-CUDA web page. [Online]. Available: https://www.tu-chemnitz.de/informatik/PI/sonstiges/downloads/npb-gpu/index.php.en

[22] HPCG benchmark [Online]. https://www.hpcg-benchmark.org/

[23] AFiD [Online]. https://github.com/PhysicsofFluids/AFiD

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.

List of Annexes

Videos:

1. HEROES Backend and APIs – https://youtu.be/_pQ4myKcrLI
2. HEROES Infrastructure – https://youtu.be/TKyQx5CnzKM
3. HEROES Workflow and Containerization – https://youtu.be/w17BeklNMKM
4. HEROES Energy Aware Runtime (EAR) – https://youtu.be/UfgTrvU6Jjg
5. HEROES Decision Module – https://youtu.be/0Z9PVvIAJjE
6. HEROES Web User Interface – https://youtu.be/aGQ4vy_-XpQ

HEROES APIs details: HEROES_API.html

The HEROES project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956874. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Spain, Italy.