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

D3.2 – HEROES Deployment Suite

Work Package	**WP 3.2, Deployment tools, templates and scripts**
Lead Author (Org)	**Davide Pastorino (Do IT Systems s.r.l.)**
Contributing Author(s) (Org)	**Emanuele Viale (Do IT Systems s.r.l.)**
Reviewed by	**Jose E Torres (HPCNow!)** **Elisabeth Ortega (HPCNow!)** **Susana Hernandez (HPCNow!)**
Approved by	**Management Board**
Due Date	**01.03.2022**
Date	**06.04.2022**
Version	**V1.1**

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(s)	Notes
0.1	11.02.2022	Davide Pastorino (Do IT Systems s.r.l.)	TOC and V0.1
0.2	25.02.2022	Davide Pastorino (Do IT Systems s.r.l.), Emanuele Viale (Do IT Systems s.r.l.)	Executive Summary, Introduction, descriptive parts of each chapter. Technical examples.
0.3	28.02.2022	Davide Pastorino (Do IT Systems s.r.l.), Emanuele Viale (Do IT Systems s.r.l.)	Added descriptive parts to the use cases. Added process/task interaction diagrams.
0.4	11/03/2022	Davide Pastorino (Do IT Systems s.r.l.)	Fixed typos and made changes based on the review inputs. Added Conclusion.
1.0	21/03/2022	Davide Pastorino (Do IT Systems s.r.l.)	Cleanup & version upgrade.
1.1	06/04/2022	Corentin Lefèvre (Neovia)	Version approved by the Management Board
1.2	14/04/2022	Alysée Cibil & Corentin Lefèvre (Neovia)	Editing (proper versioning on front page and footer acknowledgement update)

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

List of Figures

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
API	Application Programming Interface
DB	Database
DNS	Domain Name System
DoA	Description of Action
EC	European Commission
GA	Grant Agreement to the project
HPC	High-Performance Computing
HRFS	HEROES pseudo-FileSystem
HEROES-RT	HEROES Runtime programs and tools
KPI	Key Performance Indicator
SSH	Secure Shell Protocol
UUID	Universally unique identifier
VPN	Virtual Private Network
HRFS	HEROES pseudo-FileSystem

References and Applicable Documents

[1] Terraform official website, https://www.terraform.io

[2] Ansible official website, https://www.ansible.com/

[3] Bash official website, https://www.gnu.org/software/bash/

[4] AWS, Amazon Web Services, https://aws.amazon.com/

[5] DynamoDB, NoSQL database service, https://aws.amazon.com/dynamodb/

[6] Ansible Reference Appendices, *https://docs.ansible.com/ansible/latest/reference_appendices/glossary.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.

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 Centers 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 document presents the tools, the templates and the scripts that are used by the HEROES platform component to perform deployment tasks of its underlying infrastructure, software requirements and related configurations.

Existing and new tools will be designed, integrated and operated to allow the HEROES platform to work with cloud infrastructures as well as on-premises infrastructures.

The development of this document is closely linked to the ongoing developments of WP2 and WP3 of the HEROES Project.

The scripts, templates and tools described in this document will be constantly reviewed and adapted to satisfy all emergent needs of the HEROES platform throughout the life of the Project.

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

1.1 Scope and Purpose

The purpose of this document is to describe, collect and provide guidance for future developments of the tools used by the HEROES Project to deploy, configure and control its own infrastructure and the remote HPC infrastructures that will be part of the HEROES platform.

The collection of these items will be referred as HEROES Deployment Suite.

1.2 Initial requirements & tool selection

The HEROES platform is made of several software and hardware components that need to operate in a consistent and predictable state.

This requirement alone implies that a strong configuration management solution, or multiple solutions, needs to be identified and developed to ensure the HEROES platform correct operations.

Starting from this assumption, the Project team identified a few Foundation Tools that could be built upon to satisfy this initial requirement and all emergent requirements that will be born during the development of the HEROES platform.

The selection was made with the following requirements in mind:

- The tool should be opensource
- The tool should be an industry standard, and/or widely used for the purpose of configuration management
- The tool should be simple but powerful
- The tool should be able to scale without incurring in complexity or performance problems
- The tool should require minimum footprint on remote servers, possibly just an SSH connection

Henceforth, the selected Foundation Tools are:

- Terraform [1] (infrastructure-as-code tool)
- Ansible [2] (task automation tool)
- Bash [3] scripts (“glue” to put everything together)

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 Deployment architecture overview

The HEROES Deployment Suite is composed of various scripts and tools that can be integrated at various levels in the HEROES platform.

For example, the same Ansible [2] playbook that controls the configuration status of a remote server can be called directly from the HEROES APIs during the enrollment process or within a Workflow that is taking care of a User's Job, to ensure that the Job has everything it needs to execute successfully.

The diagram in Figure 1 provides an overview of the HEROES Deployment Stack, with the different stages of deployment, the tools used at each stage and the level of the expected output artifacts and resources:

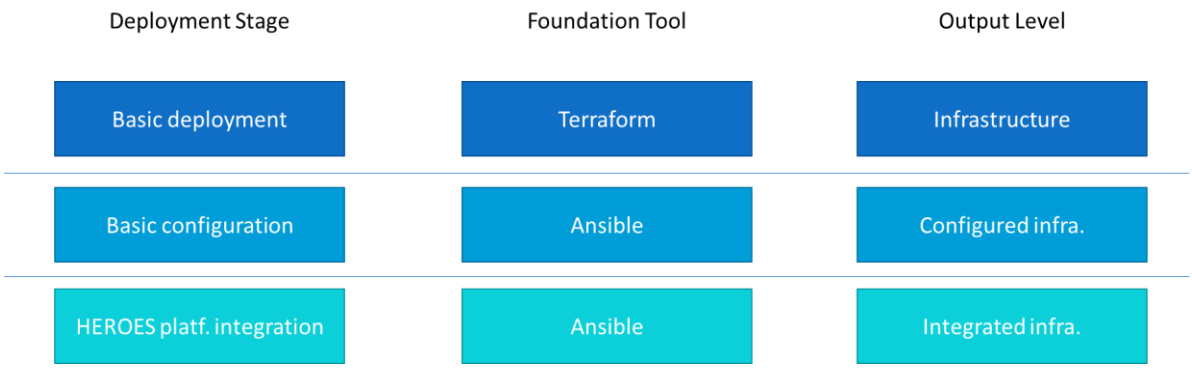

Figure 1 – HEROES Deployment Suite overview

2.1 Use Cases

The following are the current use cases for the HEROES Deployment Tools. As stated in the Executive Summary, these use cases are expected to grow in number and detail with the ongoing development of the HEROES platform.

2.1.1 HEROES Remote pseudo-FileSystem (HRFS)

The HEROES Remote pseudo-FileSystem (HRFS) is born from the requirement to integrate remote, on-premises HPC infrastructures that limit the HEROES platform access to a single service user (usually via SSH). This causes all remote files on the HPC infrastructure to be owned by the same service user.

To ensure that the HEROES platform user data is correctly segregated for each user of the platform, a particular directory structure needs to be:

- Created, at first access for the specific user

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.

- Checked (and possibly Corrected), for subsequent accesses of the same specific user

Since this is a basic Configuration Management task, the project team developed an Ansible role that takes the following inputs:

- User UUID
- User's Organization UUID

And sets up the User's folders on the remote HPC infrastructure.

The User's folders are then used as private mount points for the Containers that execute the User's Jobs.

This effectively segregates each User's own data from the others.

Additionally, since the directory structure uses UUIDs as directory names, all data is stored in anonymous paths: an authenticated session with the HEROES Identity Management module is required to match a User's UUID with his/her name.

2.1.1.1 HRFS Initial deployment

During the enrollment process of an HPC infrastructure, after the basic system configuration, the HRFS is created in the HEROES service user home directory. The basic HRFS structure includes the main directory, the HEROES management organization and the HEROES management and test users.

2.1.1.2 HRFS Sanity check

For every job that is run on a remote cluster, the HRFS is checked before the input data transfer step to verify that everything is ready for usage. If the job's organization directory and/or the user directory is missing, it is created according to the required UUIDs.

2.1.1.3 Interfaces with other modules and components

Figure 2 describes the scope of the HRFS use case:

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.

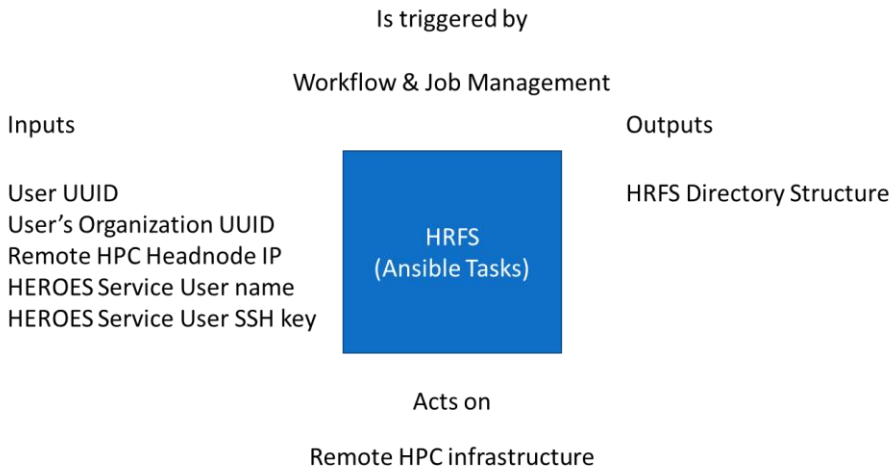

Figure 2 – HEROES pseudo-FileSystem process interaction diagram

2.1.2 HEROES Runtime

The HEROES Runtime (shorthand: HEROES-RT) is a collection of software and tools that needs to be available on all remote execution hosts that are part of the HEROES platform.

The tools are stored in a specific .tar.gz archive that needs to be transferred, unpacked and configured for each remote HPC infrastructure, cloud or on-premises.

To satisfy this use case the project team developed an Ansible role that:

- Recovers the .tar.gz archive from the storage location
- Transfers the archive to the remote host
- Checks the consistency of the archive to ensure it has not been changed or corrupted during the transfer operation
- Unpacks the archive in the pre-defined location for that specific host

The HEROES Runtime can then be accessed by the User's workflows and jobs in the expected location for the remote host that was selected for execution.

2.1.2.1 HEROES Runtime initial deployment

During the enrollment process of a HPC infrastructure, after the basic system configuration, the HEROES-RT is installed in a specific directory, that must be shared with all the execution nodes within the HPC infrastructure. Usually, this directory defaults to the HEROES service user home directory.

2.1.2.2 HEROES Runtime sanity check

For every job that is run on a remote cluster, the presence of the HEROES-RT is checked before the input data transfer to verify that everything is ready for usage. If the HEROES-RT is not present as expected, the procedure executes a new deployment.

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.

2.1.2.3 Interfaces with other modules and components

Figure 3 describes the scope of the HEROES-RT use case:

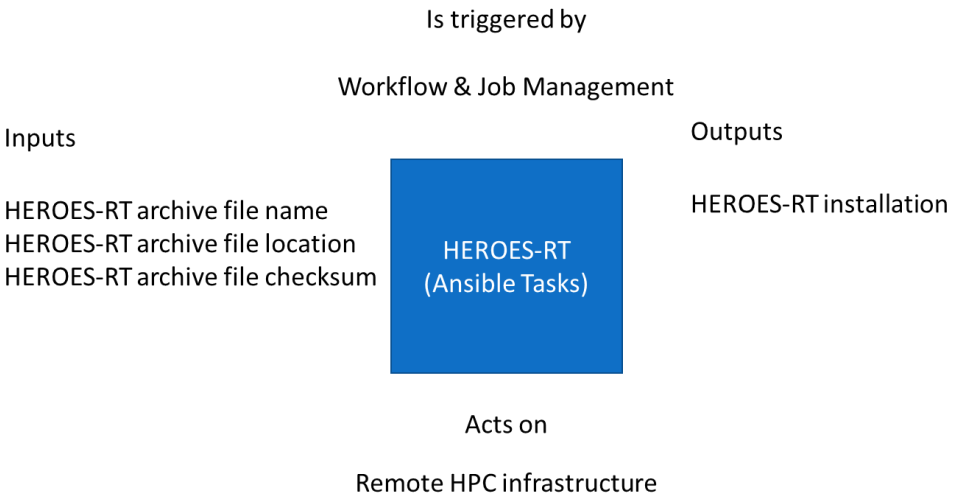

Figure 3 – HEROES Runtime deployment process interaction diagram

2.1.3 HEROES Public Cloud HPC Basic Kit

The HEROES platform includes a few major public cloud platforms as execution hosts. Each cloud platform has its own API and its own way of interaction with specific commands and configuration. The HEROES platform, however, needs to interact with these cloud platforms in a consistent way to ensure that the same HPC infrastructure can be deployed on each of them, keeping differences between cloud platforms at a minimum.

To address this need, the project team developed a series of Terraform [1] templates.
These templates describe the objects that compose the "HEROES Public Cloud HPC Basic Kit":

- VPN
- cloud network subnet
- jumphost server, a host for administrative purposes
- scheduler & login server
- shared home storage
- base compute host image
- persistent compute nodes (if any)

The Terraform [1] template is used during the enrollment process of an Organization that choses to have its own private HPC cluster in the public cloud, to effectively setup all the required infrastructure.

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 project team is also developing various Ansible [2] tasks to install and configure the basic software that is needed to operate the HPC cluster, such as the scheduler, and all related infrastructural services.

2.1.3.1 Initial deployment

During the creation of a new HPC infrastructure on a public cloud, the procedure sets up all the basic components that are required to connect the new infrastructure with the HEROES platform and configures them with the parameters for Organization that requested the deployment.
At the end of the deployment all information about the newly created infrastructure is sent back as output to the HEROES platform.

2.1.3.2 Interfaces with other modules and components

Figure 4 describes the scope of the HEROES-RT use case:

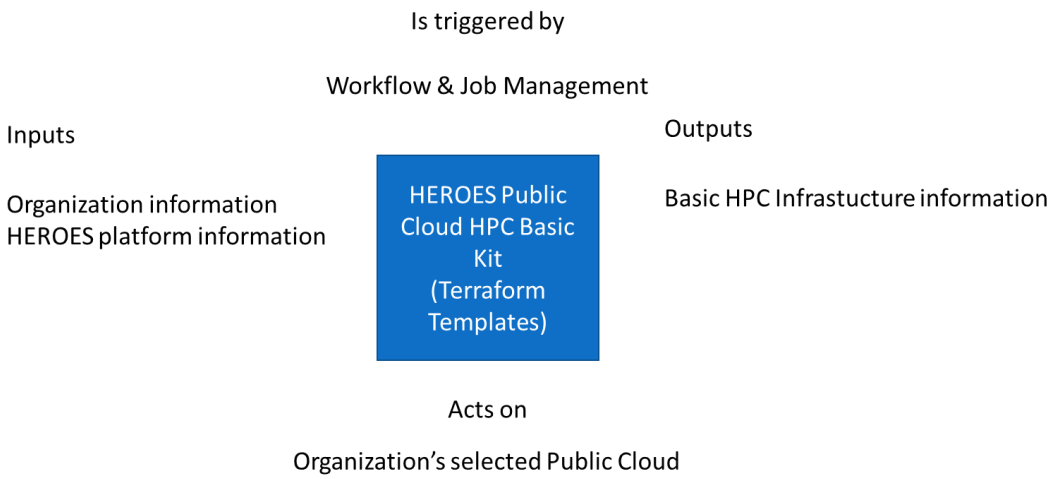

Figure 4 – HEROES Public Cloud HPC Basic Kit deployment process interaction diagram

Please note that, at the time of writing this document, the full details of the information and interfaces that will be part of this use case are not yet fully defined. The full procedure will be completed during the ongoing work of the HEROES Project.

3 Foundation tools

The following tools were chosen as foundations to build the HEROES Deployment architecture:

- Terraform [1]
- Ansible [2]

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 architecture also makes use of Bash [3] scripts for integration purposes and to overcome any limitation that have been (or could be) encountered with each of the foundation tools.

3.1 Terraform

Terraform [1] is an infrastructure as code tool that lets us define both cloud and on-prem resources in human-readable configuration files that we can version, reuse, and share.

It can be used to implement a consistent workflow to provision and manage all the Project's infrastructure throughout its lifecycle. Terraform [1] is able manage low-level components like compute, storage, and networking resources, as well as high-level components like DNS entries and other Software-as-a-Service features that are commonly used on public cloud providers.

3.1.1 Role & Scope

Terraform is used by the HEROES Project to deploy cloud HPC infrastructures that have a consistent configuration between different cloud providers. The tool scope is focused on the deployment of a raw infrastructure that is further customized using Ansible (more about Ansible roles in the next chapter).

3.1.2 Example

The following is an example (Figure 5) of a Terraform directory structure required to deploy the HEROES Public Cloud HPC Basic Kit infrastructure in the AWS [4] cloud provider:

```
aws
├── bootstrap
│   ├── main.tf
│   ├── outputs.tf
│   ├── provider.tf
│   ├── README.md
│   ├── state.tf
│   ├── terraform.tfvars
│   └── variables.tf
└── project1
    ├── backend.tf
    ├── customer_gateway.tf
    ├── instances.tf
    ├── internet_gateway.tf
    ├── key_pair.tf
    ├── main.tf
    ├── outputs.tf
    ├── provider.tf
    ├── README.md
    ├── routes.tf
    ├── security-group.tf
```

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.

```
    ├── subnet.tf
    ├── terraform.tfvars
    ├── variables.tf
    ├── virtual_private_gateway.tf
    ├── vpc.tf
    └── vpn.tf
```

Figure 5 – Directory structure of the HEROES Public Cloud HPC Basic Kit Terraform template

The project1 directory contains the actual files to deploy the infrastructure; the bootstrap directory contains the necessary files to create a locking database (in AWS, for example, we used DynamoDB [5]) that will manage the necessary locks to enable the project team to work together on the same configurations.

Each file in the project1 directory contains the declarations of the single objects that will be created on the AWS cloud provider by running the `terraform apply` command.

All the relevant variables that are used to customize the basic deployment (IP, subnet, VPN endpoints, host name prefix, etc.) are stored in project1/terraform.tfvars. This file can be automatically generated for each HEROES Organization at runtime by getting the necessary values from the HEROES central DB.

3.2 Ansible

Ansible [2] is the simplest solution for configuration management available. It's designed to be minimal in nature, consistent, secure and highly reliable, with an extremely low learning curve for administrators and developers. Ansible also requires nothing more than a SSH key to manage systems, so it can manage them without installing any agent software.

3.2.1 Role & scope

Ansible role in the HEROES Project is to perform tasks related to Configuration Management, to ensure that all the components of the HEROES platform are configured in a consistent and repeatable way. Ansible features let the project team define the desired state of a host or a software configuration within a role. When the role is “run” against a target, it can work as “first-time setup” and as “consistency check” at the same time: Ansible will ensure that the target is configured as described and apply all relevant changes to bring it to the desired state; if the target is already in the desired state, nothing will be changed and the outcome will register the target as “configured as expected”.

3.2.2 Example usage

The following is the directory structure of the tasks used to apply the User segregation structure of the HEROES pseudo FileSystem (HRFS) on an on-premises HPC infrastructure:

```
ansible/
```

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.

```
├── ansible.cfg
├── hosts
│   ├── production
│   └── staging
├── main.yml
├── README.md
├── requirements.yml
├── roles
├── tasks
│   └── segregation_check.yml
└── vars
    ├── ansible.yml
    ├── heroes.yml
    ├── organizations.yml
```

Figure 6 – Directory structure of the HEROES pseudo-FileSystem Ansible task

The tasks file defines all the actions needed to (re)create the HRFS structure on the remote host. The content of many other files (ansible.cfg, heroes.yml, organizations.yml, hosts/production, hosts/staging) are generated at runtime with the appropriate values when the ansible task is called by the HEROES platform to perform the configuration step.

As previously stated, the tasks themselves are idempotent, so they can be called at any time to both serve as initial configuration and to subsequently check the consistency of said configuration.

(*An operation is idempotent if the result of performing it once is exactly the same as the result of performing it repeatedly without any intervening actions. [6]*)

3.3 Bash scripts

BASH [3] is a system shell first released in 1989 and has since been used as the default user environment for most Linux and Unix-like operating systems. This ensures that all scripts developed by the HEROES project team will be able to run on all the systems that will be enrolled in the HEROES platform.

3.3.1 Role & Scope

The role of the Bash Scripts within the HEROES project is to act as “glue” between other software and components of the HEROES platform, and also to manage the basic user environment that is expected to be found on the on-premises HPC infrastructures that will join the platform (for example, the BSC HPC cluster).

3.3.2 Example Usage

At the current revision of this document no bash scripts have been required by the project.

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 Conclusions

In this document we presented the current deployment tools that have been developed to support the HEROES platform during the initial phases of development. These tools, as well as new tools that have yet to be created, will evolve throughout the development process of the HEROES platform.

The HEROES Deployment Tools already cover a number of key use-cases and provide help to manage file-level security within the project platform and to verify (and correct) the software-level compliance of the HPC infrastructures integrated in the HEROES platform.

Additionally, a framework is provided for the creation and management of a consistent basic HPC infrastructure on various public clouds.

This document will be updated to reflect the developments made in the HEROES Deployment Suite during the development process of the HEROES platform.

We expect to release a new version of this document at least after each future project Milestone.

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.