H2020-INFRAEDI-2018-2020

The European Centre of Excellence for Engineering Applications

Project Number: 823691

D5.3
Services, Training and Portal Report

The EXCELLERAT project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823691

Workpackage:	WP5	Centre implementation	
Author(s):	Claudio Arlandini		CINECA
	Eric Pascolo		CINECA
	Lorenzo Zanon		USTUTT
	Francesco Salvadore		CINECA
	Ivan Spisso		CINECA
Approved by	Executive Centre Management		2020-01-31
Reviewer	Thomas Gerhold		DLR
Reviewer	Niclas Jansson		KTH
Dissemination Level	Public		

Date	Author	Comments	Version	Status
2019-11-12	C. Arlandini	Initial draft	V0.1	Draft
2019-12-16	All	Complete draft for QA	V1.0	Draft
2020-01-08	All	Revision after 1st round QA	V2.0	Draft
2020-01-08	L. Zanon	Additional revision	V2.1	Draft
2020-01-09	All	Final	V3.0	Final

List of abbreviations

ACSI	*American Customer Satisfaction Index*
API	*Application Programming Interface*
CFD	*Computational Fluid Dynamics*
ChEESE	*Centre of Excellence in the domain of Solid Earth*
CMS	*Content Management System*
CoE	*Centre of Excellence*
CSA	*Coordination and Support Action*
DN	*Distinguished Name*
ERD	*Entity-Relation Diagram*
GDPR	*General Data Protection Regulation*
HiFi-Turb	*High-Fidelity LES/DNS Data for Innovative Turbulence Models*
HPC	*High Performance Computing*
ISV	*Independent Software Vendor*
JSON	*JavaScript Object Notation*
KPI	*Key Performance Indicator*
LDAP	*Lightweight Directory Access Protocol*
MS	*Milestone*
ORM	*Object Relational Mapper*
PRACE	*Partnership for Advanced Computing in Europe*
REST	*Representational state transfer*
RT	*Request Tracker*
SQL	*Standard Query Language*
SWAN	*System for the Worldwide Exchange of Engineering Data*
WP	*Work Package*

Executive Summary

This deliverable D5.3 "*Services, Training and Portal Report*" presents the outcomes of the first year of Work Package (WP) 5 activities for all its five tasks and in particular for what concerns the architecture and implementation of the service portal, opened to the public at Project month 13 (M13) as described in D5.2 "*Portal Release*". Year 1 KPIs are also presented and commented.

After an Introduction Section, Section 2 "*Service portal implementation*" presents the architecture and implementation details of the alpha release of the service portal, as derived from the effort of Tasks 5.1 and 5.2. Section 3 "*Further applications*" describes Task 5.3 activities to prepare the service devoted to the onboarding of new applications in the Project. Section 4 "*Training*" describes Task 5.4 activities in preparing and deploying an EXCELLERAT training portfolio, while Section 5 "*HPC Provisioning*" accounts for Task 5.5 effort in providing the HPC resources the CoE activities rely on. Section 6 draws some conclusion and outlines future activities.

Table of Contents

Table of Figures

Table of Tables

1 Introduction

The objective of WP5 "Centre Implementation" is ensuring the management of the day-by-day operation of EXCELLERAT as well as supporting the implementation of the defined services. WP5 also manages services, users' access, and related issues; both for internal services and for services towards the users.

This includes:

- Services for end-users as designed in WPs 2,3,4 with a particular focus on the industrial users,
- Training and education services,
- Centre of Excellence (CoE) internal administrative services, and
- Support for new application codes.

Another key aspect is the management of HPC resources provisioning (CPU hours, storage, etc.), including the integration of infrastructure services, for EXCELLERAT activities, to be coordinated with the provision services of the HPC centres.

The main output of this work package is to provide a single-entry point to its ecosystem of services, thus, building up an online access point – the portal – which will be incrementally extended by new services. The alpha version of the portal was released in M13, as described in D5.2 "Portal Release".

This deliverable "Services, Training and Portal Report" presents the outcomes of the first year of WP5 activities for all its five tasks and in particular for what concerns the architecture and implementation of the above-mentioned portal.

The determination of the KPIs for Year 1 as defined in D5.1 "*Initial Assessment of Training Needs and Services Building Plan*" to support a successful implementation is also presented.

2 Service portal implementation

2.1 Status

2.1.1 Requirements, architecture and implementation

Following the requirements defined in D5.1 "*Initial Assessment of Training Needs and Services Building Plan*" and summarized in an essential way in D5.2 "*Portal Release*", the EXCELLERAT service portal has two main aims:

1. the registration, presentation and selection of services offered by the consortium for the end user
2. the first interaction between the end user and the EXCELLERAT staff related to the provision of the service (service request).

These two fundamental functions are related to each other - for example, each service request clearly refers to a service presented by the portal - and also includes a series of detailed features.

The portal will be under development during the whole project, in constant evolution thanks to active users' feedback and the implementation of new services, but is already operational in the pre-production phase at its definitive public address[1].

The architecture of the work-flow of use of the portal from the point of view of the user (external or belonging to EXCELLERAT staff) is summarized in Figure 1.

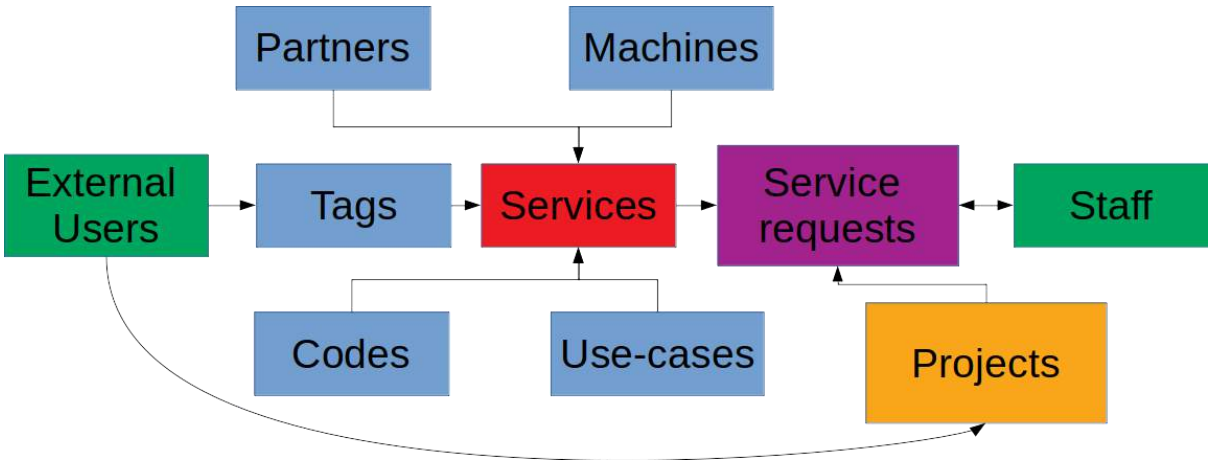

Figure 1: Architecture of the work-flow from the point of view of the user

The fundamental elements of the use of the portal itself have already been discussed in the deliverable D5.2, where contextually the method of realization of the requirements proposed in D5.1 "*Initial Assessment of Training Needs and Services Building Plan*" is also discussed. In this document we discuss the fundamental elements of the implementation choice. In this regard, the first choice concerns the possibility of using ready-made tools to configure for the needs. A wide variety of Content Management Systems (CMS) exists that are both widespread and open-source – e.g, WordPress [1] or Drupal [2] - and that allow a more or less significant possibility of customization with respect to, for instance, advanced types of research or management of content types and user access. In this way, we could essentially achieve the first objective mentioned above. Secondly, there are also consolidated tools for managing trouble

[1] *https://services.excellerat.eu*

tickets or issues, among which we have considered RT Issue Tracker [3] in more depth. For the EXCELLERAT service portal, however, we have noted the need to have a tool capable of integrating the complete tracking with the cataloging and search of contents, something not present in any of the available tools we have considered. We therefore opted to create a custom tool. This allows creating a portal that is truly adherent to the features required by the consortium. Furthermore, elements of greater specificity will be possible in the future concerning - for example - connections with computing machines or typical features in HPC that are difficult to manage with large-scale generic products.

It should be noted that the development of a custom portal should not be understood as a realization of a portal from scratch in a literal sense. The creation of the portal was instead set up and implemented so as to integrate state-of-the-art development frameworks and libraries, in order to combine efficient and robust tools. Figure 2 shows the diagram of the main components used in the implementation of the various layers of the portal.

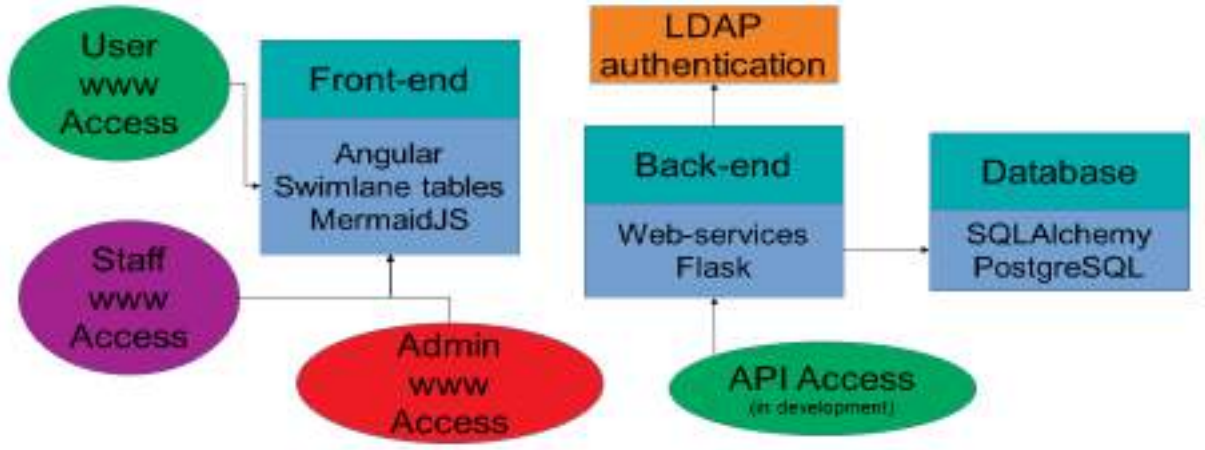

Figure 2: Diagram of the main portal components: implementation layers (rectangular boxes) and user access types (curved boxes)

The use of the individual components will be briefly discussed in the rest of the document.

2.1.2 Authentication, authorization and accounting

One of the primary aspects of portal architecture design is authentication management. In this regard, the first choice is between a local management - generally through the application database - or centralized through an external service possibly shared with other applications and / or related to a certain institution / body. After a discussion among the EXCELLERAT partners, it was decided to opt for a centralized solution that will allow at least potentially to share the authentication with different applications in addition to the service portal. We have chosen LDAP (Lightweight Directory Access Protocol) [4], a state-of-the-art solution that offers considerable robustness, documentation and configuration possibilities. We put up a specific instance for EXCELLERAT.

The LDAP instance manages only the user database through the *iNetOrgPerson* Object Class that we have configured in a minimal way to store only a minimum set of user data, with the user email adopted as Distinguished Name (DN). LDAP is the authentication mode used by the portal which then through the web-services and in particular using the python *ldap3* library verifies the correct insertion of user name and password. LDAP will contain both internal and

external users. At the moment the insertion takes place manually by EXCELLERAT staff who is contacted via the *support@excellerat.eu* support mailing list.

The LDAP instance is also configured with the standard *phpLDAPadmin* service [5] for intuitive web access to the service[2].

For security reasons, both access to the authentication LDAP service on port 389 and that of the web dashboard on port 443 are restricted by firewall to the IP addresses of the web-server of the portal that uses the service, and to the IPs of EXCELLERAT members that deal with the management of the service. In addition, both accesses enforce the usage of encrypted connections.

To allow a more user-friendly management of the password, a user manager on LDAP has been created having the privileges to change all user passwords, so that this functionality is active through the portal.

The LDAP service deals only with authentication, while the management of privileges is linked to the users belonging to specific groups. These groups are managed at the local portal database level. In this way, it will be potentially possible to have different services that use the same LDAP instance giving different privileges corresponding to different features of the services offered.

As described in D5.2, the services portal presents group-based privilege management, called in the portal *Organizations*, which are designed as a team of people who share the ownership of portal content. Users and organizations are in a many-to-many relationship. Some organizations are flagged as staff. In this case all the users who are part of it will have special privileges.

Depending on the membership of a user to one or more groups, a total of five authorization levels are configured in the portal:

- **anonymous user** (not logged in): can visit the public parts of the site in read only
- **logged-in user**: same permissions as anonymous user but can be included in a group
- **user included in at least one group**: can create projects and request services
- **user included in at least one group flagged as staff**: can access internal services and respond to service requests
- **admin user**: can edit the public and private parts of the site, as well as manage organizations and users.

The authorizations are managed both on the back-end at the web-services level - as necessary for security - and at the front-end level to improve usability, so that depending on the authorization level, different panels will be available and easily accessible features.

2.1.3 The data model

The data model is the main framework of the application. Its good design allows managing efficiently and error-prone the storage, extraction and search functions of the data offered by the portal.

First of all, it is necessary to define the fundamental entities that make up the data model. For our application, we have identified a first group of fundamental entities that refer to authentication, authorization and accounting:

- **users**, being the individuals internal or external who authenticate on the portal
- **organizations**, intended as work teams that share privileges in the portal

[2] *https://ldap.excellerat.eu*

A second group of fundamental entities refers to the actors of the project:

- **partners**, institutions that are part of the EXCELLERAT consortium
- **codes**, core-codes considered in the project
- **machines**, HPC clusters used in the project
- **use-cases**, the use-cases addressed in the project
- **services**, services offered by the consortium for external users
- **tags**, categories assigned to services

A third group of fundamental entities refers to the activity of users in the request for services, in particular therefore:

- **service requests**, created by users on behalf of teams they belong to
- **projects**, containers of service requests created by users to organize requests
- **messages**, related to a certain service request

The described fundamental entities are the main ones and actually constitute the framework of the data schema that the portal will manage. The entities are related to each other: for instance, a partner is associated with all the machines it manages (one-to-many) while a service is characterized by one or more tags (many-to-many). The detailed definition of these relationships leads to the so-called Entity-Relation Diagram (ERD) [6] shown in Figure 3. In the ERD are represented all the entities that we have used including the connecting entities for the many-to-many relationships. The data model was obtained after several iterations with the project partners to understand which were the most significant and important information to manage in the portal.

At the implementation level, data management takes place through the use of relational databases, a standard tool that guarantees a whole series of advantages with respect to manual data management, such as efficient queries for the most diverse functions and transparent transaction management. In practice, each entity corresponds to a database table and the relationships are implemented as usually through SQL Foreign Keys. In particular we have chosen to use PostgreSQL [7] for its robustness and excellent documentation. We also use the spatial features provided by PostGIS [8] at the moment only for the management of the addresses of the partners.

Access to the database always takes place via web-services and we have chosen to use an Object Relational Mapper (ORM) in order to simplify access to the database avoiding layers of pure SQL that would require laborious error and data conversion management. We have chosen in particular to use SQLAlchemy [9], a state-of-the-art ORM widely used in the Python environment. The use of SQLAlchemy allows a remarkable speed in writing web-services and their interfacing to the database. However, the use is not always intuitive and a certain level of SQL competence is required. For some heavier database interactions, it may be necessary to bypass the ORM and use SQL directly. In the application of the EXCELLERAT portal the low-level SQLAlchemy APIs were used only in the case of the most complicated query, which is that of faceted search for services based on the different types of tags and other related entities.

At the current stage, the database schema is fairly stable but it is not excluded that it will have to be modified in the future. A practical tool to manage any migration that will be necessary is Alembic [10].

Figure 3: Entity-Relation Diagram (ERD) for the EXCELLERAT service portal

2.1.4 Back-end web services

The first layer of the service portal is managed through the NGINX [11] http server. This layer sorts the requests to ports 80 and 443 and, through some reverse proxy redirections, allows the server to respond to different requests such as:

1. static content as images
2. content provided via web-services
3. content for front-end operation.

In this Subsection, we will deal in particular with the most complex request, i.e. web-services APIs. The APIs have been designed mostly according to the REST architecture [12], that is taking as input and returning output in the form of JSON dictionaries, but without necessarily implementing all the types of REST requests. The framework used is Flask [13], which is used together with several additional packages in addition to the aforementioned SQLAlchemy ORM.

Particular attention has been paid to using a clean programming style that allows orderly management and good maintainability of the code. To this end, we make intensive use of the Python decorators that allow for each API:

1. validation of input arguments
2. automatic authorization management
3. automatic generation of Swagger [14] documentation (in development)

As an example, in Figure 4 we report the first layer of the code relating to the web service that implements the service search task.

```python
class SearchServicesSchema(BaseSchema):
    name               = fields.String(example="machine")
    name_like          = fields.String(example="mach")
    description_like   = fields.String(example="virtual")
    start              = fields.Integer(example=0)
    size               = fields.Integer(example=5)
    id_respuser        = fields.Integer(example=1)
    id_externalcontact = fields.Integer(example=1)
    id_taggroups       = fields.Dict()
    id_machines        = fields.List(fields.Integer(example=1))
    id_codes           = fields.List(fields.Integer(example=1))
    id_usecases        = fields.List(fields.Integer(example=1))
    id_partners        = fields.List(fields.Integer(example=1))
    include_disabled   = fields.Boolean(example=False)
@api_excellerat.route('/search_services', methods=['POST'])
@doc(description='Search services', tags=['Excellerat'])
@use_kwargs(SearchServicesSchema)
@my_jwt(SearchServicesSchema, "optional")
def api_search_services(**arg_json):
    if request.method == 'POST':
        print("arg_json wit auth: ",arg_json)
        n_all_services, services, n_partials = excellerat.search_services(arg_json)
        ret_dict = dict()
        ret_dict["results"] = dict()
        ret_dict["results"]["services"] = services
        ret_dict["results"]["n_all_services"] = n_all_services
        ret_dict["results"]["n_partials"] = n_partials
        return jsonify(ret_dict)
```

Figure 4: Flask API implementing the service search web-service (first layer)

The second layer that implements the service search functionality, which we do not report here for space reasons, implements the search according to different parameters and acts downstream of all the basic functions implemented by decorators in the first layer.

The development of a custom back-end layer - even if based on state-of-the-art frameworks and libraries - allows to implement advanced functionalities and, if necessary, the interaction with computing machines or similar advanced functions.

2.1.5 Front-end application

Within the multi-layered architecture and implementation used for the portal, the front-end is the direct access point with the user. In fact, even if the web services API can be called directly through command-line tools or specific applications, for a web application like this portal the external use through API appears, in the vast majority of cases, not optimal.

For the implementation of the front-end we first of all used the Angular framework [15], a state-of-the-art solution capable of achieving high usability dynamic interfaces and maintaining an orderly, robust and maintainable code structure. In terms of usability we focused on the implementation of static and dynamic messages that make the use of the portal by the user particularly intuitive.

As mentioned, the Angular implementation allows to easily manage the privileges management also in the front-end both because some routes are accessible only for certain authorization levels and because the same route can present different contents depending on which user is logged in.

Figure 5: Test project dashboard

An important part of the portal front-end is the dashboards that collect user-created content, i.e. projects and service requests. To this end we have used a special library, *ngx-datatable* [16],

for managing interactive data tables. An example of a test project dashboard is provided in Figure 5.

Figure 6: Small screen view of responsive service portal layout

Another interesting aspect of the front-end is the possibility of defining a typical workflow within the individual services, and for this type of content we use the *mermaidjs* library [17] which provides an extension of the Markdown language designed for the simple writing of diagrams to blocks.

Even the graphic style of the web pages is of fundamental importance to improve the usability of the portal. The Bootstrap toolkit [18] was used for the portal with an additional library for integration with Angular. Bootstrap allows, among other things, to implement the so-called responsiveness, i.e. the ability of the page to adapt to the device that requires it. In figure 6, the adapted view of the portal on a mobile device screen is provided. We have chosen to implement a hybrid layout between fluid and fixed, so that it adapts to the entire screen size up to a certain width and then becomes fixed for large screens so that the dashboards are still easy to read.

Regarding the use of cookies and compliance with the GDPR legislation [19], the portal currently uses only essential technical cookies relating to the session and authentication. The footer of each page of the portal includes a clear indication and link to the CINECA privacy policy and for this reason it is not currently necessary to use a specific banner to inform the user of the use of cookies. However, we have prepared the integration with the *Cookie Consent* library [20] that in the future will allow the user, if necessary, to accept or reject non-essential cookies that will be considered useful.

2.1.6 Code management and deployment

The portal code includes different languages and different libraries. A summary of the statistics of the used languages is available in Figure 7. The statistics refer only to the source code

developed for the project, while the files generated automatically during the compilation phases are excluded, as well as the files supplied by the used libraries.

Language	Files	Lines	Code	Comments	Blanks
Dockerfile	1	271	159	73	39
HTML	38	3651	2960	311	380
INI	2	70	39	13	18
JSON	1	66	66	0	0
Markdown	2	81	81	0	0
Python	19	5962	4451	752	759
Sass	38	1021	650	247	124
Shell	7	113	47	48	18
TypeScript	48	5035	4139	357	539
YAML	2	68	52	0	8
Total	158	16330	12644	1801	1885

Figure 7: Statistics of languages used

The management of such a source requires orderly structuring and advanced development tools to achieve high quality of the result. To this end we have used Git [21] as a tool for code versioning and Visual Studio Code (VSCode) [22] as a tool for code editing. In particular, VSCode offers numerous aids to improve productivity in terms of code writing (e.g., linting, IntelliSense), launching development servers, and Git integration. VSCode also offers a very useful debugger that greatly reduces the time to find and fix bugs.

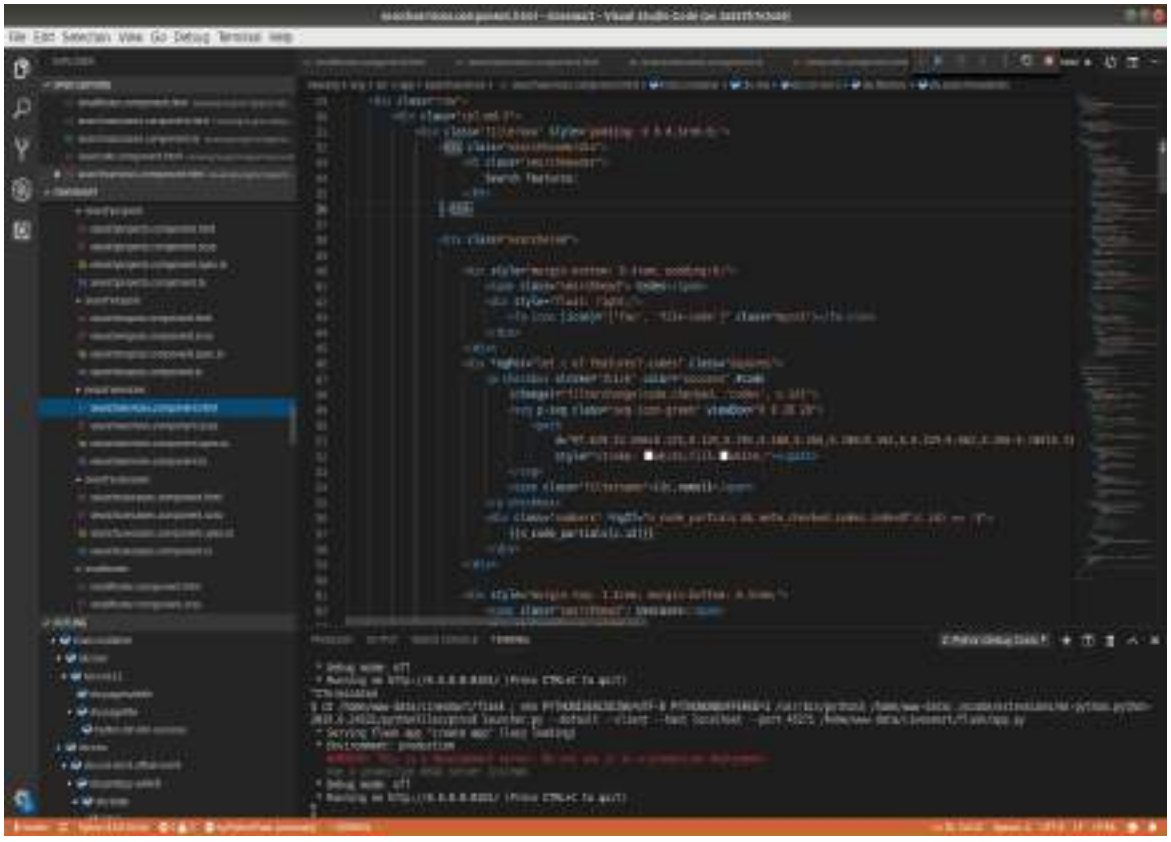

Figure 8: Visual Studio Code screenshot from portal source code

To facilitate the installation and management of the portal platform, all the portal components are containerized via Docker [23] containers - one configured for development instances, one for the production instance - and this allows to minimize the effort and time of deployment of the portal on a new architecture.

The portal's production instance is currently deployed on a 4 vCPU Virtual Machine and 12GB of RAM on CINECA's OpenStack platform [24]. The physical hardware of the machine consists of Intel (R) Xeon (R) CPU E5-2670 v2 @ 2.50GHz CPUs.

As for the LDAP instance it is also in a Virtual Machine of the same environment but 1 vCPU and 3GB of RAM.

2.1.7 Internal services implementation

In D5.1 "*Initial Assessment of Training Needs and Services Building Plan*" we discussed the services identified for CoE operations, collected through questionnaires, and prioritized using the MoSCoW methodology [25]. The highest priority requests for services to be implemented were: document sharing, wiki, code repository and automated testing tools.

For the first portal release, the tools for internal services implemented into the service portal are:

- BSCW [26] for document sharing, integrating the HLRS instance deployed in the first phase of the project;
- A HLRS media wiki instance
- GoToMeeting [27] teleconferencing tool for project communication
- Etherpad [28] for collaborative information exchange
- OpenProject [29] for activity tracking

The choice was the result of a refinement of the analysis performed for D5.1. Etherpad and GoToMeeting were included already in the first portal version although they scored medium and low priority at the beginning, because with the start of the work on the reference codes more frequent and deeper interactions among partners were needed, and their usefulness was recognized. At the same time, it was decided not to include (at least for the moment) a common code repository and automatic testing tools environment, since each partner had already their own in use.

2.2 Activities planned

At present, it can be briefly stated that the portal is in a fairly advanced stage of development both from the point of view of architecture design and from the implementation point of view. However, several aspects still require significant improvements and others need to be fine-tuned. Below we indicate some points - not exhaustive - on which we will work later.

A first point on which we intend to work in the continuation of the project concerns the adaptation of the architecture to the cataloguing and presentation of the services that will be detailed by the other WPs. In particular, following the indications of the EXCELLERAT partners, we intend to improve the presentation of the catalogue of services minimizing the effort of the user in the search for a service that is significant for his own use case. In this sense, the description of a service will be mainly linked to the presentation of the use-case which shows in concrete cases results obtainable with the service in question. In this perspective, a reference site that we can consider is, for instance Carbone4[3]. In this regard, a particularly clear

[3] https://www.carbone4.com/services_/?lang=en

way of categorizing services is the one presented in D4.2 "*Report on the service portfolio*" in which external services are divided into three categories "Community perspective", "Solution evolution" and "Application / Code evolution ". Our tag management can be adapted to manage services following this first subdivision.

A second essential point for the success of the portal is the population of the contents. A well-engineered portal cannot succeed unless it has significant and continuously updated content. In this sense, we will do in such a way as to have a high richness of the contents, in particular those that are more in evolution, or use cases and services. Also, with regard to users, we will complete the insertion of internal users and manage the addition of external users who will arrive after the official opening of the portal.

A third important activity of the portal concerns the testing of the implemented functions as well as the evaluation of the robustness and usability of the platform. Particularly important is the testing of the back-end API so we will choose the most complete methods possible. As for the robustness we will monitor the percentage of uptime that is the time in which the portal is correctly up and running. As for usability, we will collect feedback from both project partners and external users. Also based on their suggestions, we will make the most appropriate changes.

The fourth activity we foresee in the future of portal development is the preparation of a user guide. Although the portal will be as intuitive as possible to use, a user guide can help you make the most out of the features offered. The guide will then be put online using consolidated development tools and will be easily accessible from the portal itself.

2.3 KPI analysis

Concerning the KPIs collection, we will start from those defined in D5.1. In that report, a distinction was made between the KPIs for internal services, defined by Task 5.1, and those for external services, defined by Task T5.2.

KPI number	Description	Definition	Target	Score at M13
5.1.1	Percentage of the number of internal services activated compared to the number of services identified for the present implementation release.	*N_activated / N_identified * 100*	>75%	92% (6,5 implemented / 7 high priority)
5.1.2	Number of internal active users	*N_active_monthly*	≥15 for the first year after MS4 completion, with a 20% average increase later	20 (estimated by telco and BSCW interaction)
5.1.3	ACSI score: user satisfaction for internal services, annually	N_{ACSI}*_int_services*	≥80	N/A

Table 1: KPIs for Task 5.1

Table 1 gives an overview of the KPIs for internal services. Since the portal has been just made available to the users, it is too early to evaluate their satisfaction, hence KPI 5.1.3 score will be presented in the next iteration of the deliverable.

KPI number	Description	Definition	Target	Score at M13
5.2.1	Number of external	*N_ext_users*	≥20 at the end of the first year after MS4 completion, ≥30 at the end of the project	N/A
5.2.2	Number of requested services (for the present implementation release)	*N_services*	>1	N/A
5.2.3	Percentage of the number of external services activated compared to the number of services identified for the present implementation release.	*N_activated / N_identified * 100*	>75%	N/A
5.2.4	Uptime	*Time of portal availability /Time from first production release*100*	>70% for the first year after MS4 completion >85% for the second, >95% later	N/A
5.2.5	ACSI score: user satisfaction for external service, annually	N_{ACSI}*_ext_services*	≥80	N/A

Table 2: KPIs for Task 5.2

Table 2 gives an overview of the KPIs for external services. Since the portal has been just made available and only to internal users, it is too early to evaluate all those metrics, hence they will be presented in the next iteration of the deliverable.

3 Further applications

3.1 Status and activities performed

In D5.1 "*Initial Assessment of Training Needs and Services Building Plan*", several ways have been foreseen for EXCELLERAT to attract new applications and, thus, potential new end-users/customers. The main two, interaction with the Interest Groups and portal development, are described in detail in the next Subsections.

A third one is the outreach of EXCELLERAT members, by means of which an interesting candidate application arose, namely the code for fluid-dynamic applications in "two-way coupling" regime by Prof. C.M. Casciola at the University "La Sapienza" (Rome, Italy). A first contact has been established by CINECA members with Prof. Casciola and Dr. Battista, whereas details of an upcoming collaboration will be defined in the next months.

Concerning recruitment through awareness creation channels and contacts at conferences and workshops, no substantial achievements can be mentioned. Therefore, these points will be mentioned in Subsection 3.2.3 (outlook).

3.1.1 Interest Groups

The first contact through "letters of support" marked the kick-off of our interaction with the Interest Groups. This endeavour, described in D5.1, has been now intensified in two ways.

First, with an official invitation, the Interest Groups have been asked to fill in and sign a consent form, requesting information about their will to become member of the EXCELLERAT Interest Groups, to appear in EXCELLERAT dissemination material, and to subscribe to the Interest Group mailing list.

Next, a first webinar hosted by HLRS and dedicated to involve the Interest Groups in the EXCELLERAT activity took place on Nov. 8, 2019. There, their role has been highlighted, as actors in communicating, monitoring and validating the project's overall goal and its industrial, as well as technological relevance for the user communities. A short overview of the project has been given, and three applications have been presented by the respective developers (the data exchange platform SWAN, visualization with Vistle, the BigWhoop compression library).

Table 3 integrates the list in D5.1 with the Interest Groups' responses to the official invitation, and their participation to the first webinar. At the moment, no further feedback has been received after the letter of support from the following Interest Groups:

- Industrial End Users: PORSCHE, SAFRAN Helicopter Engines, ROLLS ROYCE plc, DAIMLER AG, REPSOL, SCANIA, FESTO AG & Co. KG, TECOSIM GmbH, AIRBUS;
- Code-Developers/ISVs: ANSYS Inc., OPENCFD Ltd;
- Scientific Experts: GENCI;
- Technology Providers: IBM, NEC, LENOVO, Sugon.

Interest Group	Typology	Membership	Media	Mailing List	Webinar
ESI-GROUP	Code-Developer/ISV	Yes	Yes	Yes	No
FRAUNHOFER IGD	Scientific Expert	No	No	No	No
Politecnico di Milano	Scientific Expert	Yes	Yes	Yes	Yes
University of Ljubljana	Scientific Expert	Yes	Yes	Yes	No
University of Rome "Tor Vergata"	Scientific Expert				Yes
AMD	Technology Provider				Yes
ARM	Technology Provider	No	No	Yes	Yes
ATOS	Technology Provider	Yes	Yes	Yes	Yes
CRAY	Technology Provider	No	Yes	Yes	Yes
Hewlett Packard Enterprise	Technology Provider	Yes	Yes	Yes	No
INTEL	Technology Provider				Yes
Westinghouse Electric Sweden	Technology Provider	Yes	Yes	Yes	Yes

Table 3: Interest Groups (official invitation and webinar attendance)

3.1.2 EXCELLERAT Portal Application Form

This Subsection will describe in more detail the aspects of the Portal described in Section 2 regarding the task "further applications".

Based on the assessment of portals offering similar services (in D5.1), a workflow for onboarding external codes has been elaborated and graphically diagrammed with Bizagi modeler [30]. This contains in detail the various aspects of the onboarding of a new application, also distinguishing among the different responsible actors and their interaction (EXCELLERAT portal, panel, service-provider, administration). A partial representation of such a workflow has been implemented in the portal, and is accessible to registered users. Whether this workflow will be fully put into practice will depend on the evolution of the portal (Task 5.2).

In D5.1, the structural requirements for the portal to attract external codes have been listed. Some of them have already been fulfilled:

- *A description of the possible target audience*: The "Search services" page of the current portal implementation is devoted to attract external engineering applications of different fields that could benefit from EXCELLERAT achievements, here displayed in terms of available "services". The tag "custom" in particular explicitly defines a service (also) dedicated to external codes or machines.
- *The requirements they [the target audience] have to fulfil*: In this initial period of availability of the portal and of the onboarding process, we do not expect a number of applications exceeding our review capabilities. Therefore, the selection of successful external applications and codes is delegated to the contact persons of the requested service.
- *A clear description of the services that can be provided*: In addition to the services already mentioned and described in Section 2.1.1 and D4.2, a service "New challenges" has been especially conceived. This will deal with issues that could still be solved within EXCELLERAT, even though not directly corresponding to an already available service. Moreover, the specific workflow for the onboarding of external applications is represented in this service.
- *A direct link to a simple contact form*: The home-page of the portal displays an email address enabling (new) users to directly contact the EXCELLERAT consortium[4]. At the current state of implementation, visitors are strongly encouraged to get in touch with the support team, since the registration must be carried out by interacting with a portal administrator.
- *A direct link to the service/support request form*: Once a portal account has been activated for a new user, the logged-in applicant who would like to select a service is directed to a "Create Project" form. Here, new users should provide a description of their issue and select the organization to which this belongs.

3.2 Activities planned

3.2.1 Interest Groups

Further recruiting of additional Interest Groups is not planned yet, whereas a new Interest Group meeting will take place in Year 2, most probably again in form of a webinar to facilitate a larger participation. There, the Interest Groups will be updated regarding the achievements of the project, and encouraged to submit new applications, with a major stress on the now accessible portal.

The Interest Groups that agreed to be involved in dissemination material will appear in the updated EXCELLERAT website as well as in the portal (see Subsection 3.2.2).

The Interest Groups should also be a mean to identify promising applications that could be boosted by EXCELLERAT. A first contact in this direction has been established with ESI-GROUP and Westinghouse Electric Sweden, even though further interaction is needed. Strategies to efficiently single out promising applications and codes within the Interest Groups collaboration must also be devised.

There hasn't been any interaction between Interest Groups and FocusCoE [31] yet (as mentioned in D5.1). Even so, this endeavour will be integrated e.g. in the cross-over of EXCELLERAT Task 5.4 and FocusCoE Task 4.2, within the assessment of requirements for industrial training.

[4] support@excellerat.eu

3.2.2 EXCELLERAT Portal Application Form

Some of the structural requirements for the portal listed in D5.1 will be addressed in conjunction or directly after MS4 (M14):

- *Documentation to guide interested parties through the sections of the portal*: Every section of the portal is introduced by a header describing the page function. Nevertheless, as the portal will grow in complexity, additional documentation will be provided.
- *The requirements they [the target audience] have to fulfil*: This item will be elaborated (compared to Subsection 3.1.2) also according to the number of applications submitted to the portal. The criteria listed in D5.1, Subsection 4.2.3, will be carefully considered.
- *Information about Interest Group members* will be added in the "Partners" page, which will be extended with a "switch" button (or in an analogous fashion) to show/hide Interest Groups that agreed to be cited on the EXCELLERAT dissemination platforms.
- *Further material that is helpful for engineering applications* will be discussed in cooperation with Task 5.2 and WPs 2,3,4, according to the evolution of the flagship applications.

3.2.3 Planned work

A Gantt chart has been produced to guide the evolution of Task 5.3 and is available to all internal members of EXCELLERAT. Among the main objectives in there, outside the scope of Subsections 3.2.1 and 3.2.2, we can highlight:

- *Second and third review of additional services to offer, Preparation of services to offer for external applications*: Following the evolution of the reference applications, the core services already available in the portal will be enriched with more technical details, and their number possibly increased. This must be achieved in cooperation with WPs 2,3,4, e.g. by joining the technical calls of these WPs.
- *Promotion for additional applications*: Awareness creation channels and presentations by EXCELLERAT members at conferences and workshops should foster the onboarding of further applications and the dissemination of their evolution story. The cooperation with WP7 should be certainly strengthened, e.g. contributing to the newsletter and the podcast currently in preparation (HLRS in collaboration with TERATEC). WP5 and WP7 additionally encouraged the participation of EXCELLERAT members in a mini-symposium on "Open Source software for real-world applications" at the conference WWMC/ECCOMAS 2020 [32], as an opportunity to raise the interest for EXCELLERAT by developers of new applications from academia and industry.

3.3 KPI analysis

In D5.1, four functional and quality-related KPIs have been proposed, which can be seen in Table 4. The quality-related KPIs refer to the ACSI Score (American Customer Satisfaction Index), defined in Subsection 2.3 of D5.1.

Such KPIs were intended to be measured after the portal implementation in M14. In fact, before the portal release, a formal onboarding of external applications could not be properly framed. Even external codes recruited through other means, as mentioned in Subsections 3.1 and 3.2.3, must eventually be channelled into the portal before the start of a proper collaboration. For these reasons, KPIs regarding further applications will be measured after MS4 (M14).

KPI number	Description	Definition	Target	Score at M13
5.3.1	Number of external entities sending an application (quarterly)	*N_app_sent_quarterly*	≥1	N/A
5.3.2	Number of selected applications to interact with (quarterly)	*N_app_interacting_quarterly*	≥1	N/A
5.3.3	ACSI score: Satisfaction of the on-boarding (candidature) process for a service (e.g. via survey/feedback sheet) as average per year	N_{ACSI}*_onboarding*	≥80	N/A
5.3.4	ACSI score: Satisfaction of the processed applicants (e.g. via survey / feedback sheet) as average per year	N_{ACSI}*_applications_per_year*	≥80	N/A

Table 4: KPIs for Task 5.3

4 Training

4.1 Status and activities performed

As described in D5.1, training activity within EXCELLERAT must aim at filling specific gaps in the present offer of the involved partners. On the one hand, we must take into account the audience demand and the results of our training assessments. On the other hand, EXCELLERAT must be able to provide application-specific training, related to the core-codes fostered by the project. Since the production of training material, including tutorials and best-practice guidelines, presumes an advanced technical development of the codes, we expect a more significant effort in this direction in Years 2 and 3 of the project.

Training material for self-learning is an important part of our offer and is being collected through a constant exchange with the developers of EXCELLERAT applications. Announcements and links to available material are published on the EXCELLERAT website. After MS4, such material will be accessible in a detailed and structured way in the portal. There, training is present as an EXCELLERAT service: In detail, training services can be approached from the community perspective as one-on-n consulting services (i.e., a community addressing a specific provider). The portal must act as both a repository of collected material and as a mean to submit requests for special training (see also D4.2 *"Report on the Service Portfolio"*, Figs. 2 and 5).

For M1-13 the training effort below has been conducted. We also indicate the respective training category according the definition introduced in Subsection 5.2 of D5.1:

- "Introductory Webinar on OpenProject" [OTH] (2019-08-02): a webinar on an internal service, hosted by HLRS, recording available on the website.
- "SWAN webinar" [DAT] (2019-08-26): a webinar on the system for the worldwide exchange of engineering data developed at SSC-Services, recording and link to documentation available on the website**.**
- "NEK5000" [CFD] (2019-12-06): a hands-on session and use-case presentation by KTH personnel on an application developed within EXCELLERAT, at the "School on Numerical Methods for Parallel CFD" [33] ("CFD School") at CINECA in Rome. The event has been announced on the website, additional material will be made available on the portal.
- "BigWhoop compression library" [DAT] (2019-11-08): a short presentation/demo by HLRS within the EXCELLERAT Interest Group meeting (Subsection 3.1.1) on an efficient tool for effectively reducing the size of large numerical data sets. Recording is available on request.

Moreover, within the Focus CoE Coordination and Support Action, an OpenProject webinar has been jointly conducted by one member of EXCELLERAT and one of the CoE HiDALGO [34]. Also in a CoE collaboration landscape, announcement and recording of the HiDALGO webinar "Best Practice Guide for Git & Jenkins" (2018-10-31) have been published on the EXCELLERAT website. Finally, the "CFD School" included not only an EXCELLERAT session, but also a joint introduction of two additional CoEs for computing applications (ChEESE [35] and HiFi-TURB [36]), and a use-case presentation by ChEESE.

4.1.1 Status of assessment of training activities

The assessment of training activities carried out in January 2019 and thoroughly described in D5.1 encompassed events organized, held, or advertised by the organisations participating as members of EXCELLERAT. Such an assessment has not been repeated in Y1, since the first issue already considered all the year's scheduled activities, representing at the same time a substantial endeavour for the EXCELLERAT partners. It will be by all means repeated in Q1/2020.

Another important assessment activity took place at the Focus CoE Stakeholder Workshop (EC, Bruxelles, Oct. 8, 2019). There, EXCELLERAT has been represented along with other 9 CoEs, such that a comprehensive overview of various aspects of HPC training could emerge. Not only users', but also developers' and trainers' needs, and the education landscapes have been analysed.

Among the common aspects in both assessments, an under-representation of training in the domain of data management / data analysis ([DAT]), visualization ([VIS]) and software engineering (e.g. [CLU], [PRF]) could be highlighted. This is being taken into consideration by EXCELLERAT, for example with the BigWhoop and SWAN events ([DAT]), the Vistle effort at HLRS (a visualization tool for distributed parallel visualization, [VIS]), and the FRAUNHOFER activity on Data Analysis ([DAT]) listed in Subsection 4.1.2. A (higher) actual demand in the fields of cluster usage and administration, as well as in performance optimization and debugging must still be assessed.

We can also correlate the under-representation of the online offer emerged in the EXCELLERAT assessment, and the identification at the Workshop of the "mix-and-match" approach as a successful training typology. That is, in-class and online events should complement each other, the latter ideally involving an interaction among participants and training providers. The current EXCELLERAT offer in terms of webinars is a promising start in addressing this issue.

Some additional aspects of the training demand emerged at the Workshop and are relevant for EXCELLERAT:

- HPC training should be possibly provided at an early career stage (e.g. to students), and in close connection to the audience field of work/study;
- Industrial end-users appreciate goal-oriented sessions;
- Quality trainers are needed (with a stress on quality assurance, lifelong learning, and up-to-date personnel and material). In this regard, technology providers have been indicated as key contributors of valuable application-specific HPC training (in terms of best practices and up-to-date technology). For this reason, a closer cooperation with our Interest Groups (Subsection 3.1.1) is advocated. Training offer based on EXCELLERAT codes and applications can also be regarded as an action in this direction (also aimed at filling the gap in application-specific training according to the EXCELLERAT assessment).

The collaboration with FocusCoE will continue by all means after the workshop, since a comprehensive assessment – similar to that carried out by EXCELLERAT – will take place at the FocusCoE level and will provide a much broader overview, since it will involve a wider spectrum of CoEs and possibly their related partners.

Finally, further training requirements could be assessed by the analysis of the user-stories connected to the six EXCELLERAT reference applications. The emerged topics are listed in D4.2, and encompass specific topics in the fields of (hybrid) programming, CFD, and advanced

meshing techniques. The present or scheduled offer of software-specific training (NEK5000, FEniCS) moves already in that direction.

4.1.2 Planned training and workshop activities

Analysis of offer and demand among partners and end-users will be intensified. In addition to the bi-annual assessment, a competence map of topics and involved partners will be produced, as by D5.1, also based on the collected abstracts of all training material. From the users' perspective, their wishes in terms of training content and typology must be assessed. To this purpose, the satisfaction surveys of the upcoming activities can be extended with a dedicated field, as it is customary in other EXCELLERAT partnerships, e.g. PRACE. By intensifying the collaboration with WP 2/ 3/ 4, additional gaps in technical knowledge and need for training can be identified, to complement the requests already arising from the user-stories (see D4.2).

Even though webinars and online activities are capable of reaching a wide audience with ease of participation, in-class events must not be overlooked, considering their incomparable effectiveness in know-how transfer. Particularly efficient would be the integration of EXCELLERAT-specific training sessions within the partners' established curricula of courses and workshops. For example, the "Scientific Visualization" course at HLRS (two-day course, twice a year) could be enriched with a session based on the application Vistle, developed within EXCELLERAT. Plans for Year 2 foresee the inclusion in the present course of a dedicated session on distributed parallel visualization.

Similarly, the EXCELLERAT partner FRAUNHOFER plans to extend their current trainings on "Data Scientist" and "Big Data Architecture" within the "Fraunhofer Academy" with training modules on Data Analysis and Engineering Simulation Data for industrial end-users, as soon as the technical application reaches maturity (see also D4.2).

As a follow-up of the workshop at the "CFD School" (Subsection 4.1), a two-day workshop hosted at CINECA with a similar format (hands-on and use-case presentation) is planned for Year 2. The flagship applications NEK5000 and the Finite Element solver FEniCS will be presented by KTH personnel. Further contacts have been established with FEniCS experienced users (e.g. at SISSA mathLab, Trieste, Italy, and the Department of Engineering of the University of Cambridge), who could provide contributions related to their field of investigation.

Finally, we schedule to publish in the portal repository (as specified in Subsection 4.1) training modules for self-learning or complementary to training sessions. These include tutorials on Vistle, OpenProject, SWAN, as well as material complementary to the in-class events on NEK5000 and FEniCS.

4.2 KPI analysis

KPI number	Description	Definition	Target	Score at M13
5.4.1	Number of participants in EXCELLERAT training (per year)	*N_participants_per_year*	≥20 for the first year after MS4 completion, +20% for the final period	72
5.4.2	Size of developed material in number of lecture units	*Size_training_material*	≥30 at the end of the first year after MS4 completion, +50% at the project conclusion	6

5.4.3	Gender ratio of the participants	*R_female_male_ quotient*	50%	14%
5.4.4	ACSI score: Satisfaction about EXCELLERAT training outcome (via survey/feedback sheet) as average per year	N_{ACSI}*_training_ per_year*	≥80	71
5.4.5	ACSI score: Quality and uptake of developed EXCELLERAT material also by non-EXCELLERAT stakeholders as average per year (via survey/feedback sheet)	N_{ACSI}*_G_quality _material*	≥80	66

Table 5: KPIs for Task 5.4

In D5.1, five functional and quality-related KPIs have been proposed, which can be seen in Table 5. The quality-related KPIs refer to the ACSI Score (American Customer Satisfaction Index), defined in Subsection 2.3 of D5.1.

Although these KPIs are meant to be evaluated after the release of the portal (MS4 at M14), it has already been possible to get first results for M1-13. The "Number of participants" (KPI 5.4.1), referring to three webinars and the "CFD School", is largely satisfactory, even though these events could have profited in terms of audience size and diversity from an additional announcement on the portal. Nevertheless, the organizers of the activities have been able to exploit their own dissemination channels, as well as the EXCELLERAT website.

The goal of ≥50% not-male participants seems far-fetched but will be addressed in the future planning of events by means of properly tuned dissemination, and by stressing on gender-equality within all EXCELLERAT future initiatives. These actions will have to be taken in greater cooperation with WP7 and FocusCoE, and by properly designing the training section of the portal (together with Task 5.2).

The training service of the portal, also acting as a repository, will be instead critical to obtain a comprehensive and structured overview of the developed training material, in order for KPI 5.4.2 to get an exact measurement. During the first year, we only considered the complementary material (slides and exercise tutorial) of the NEK5000 event, and the material available online for the platform SWAN, also accessible from the EXCELLERAT website. We therefore expect a steep increase of lecture units and to reach the target of KPI 5.4.2 after MS4.

Although the foreseen target for KPIs 5.4.4 and 5.4.5 could not yet be met, we observed sufficient results in terms of satisfaction of the training outcome and the developed material. We must stress that these KPIs solely refer to the "CFD School". CINECA personnel stressed how this result suited the outcome of a "first edition", that the content itself of the school was largely appreciated, and that they will improve the organization of future events based on this experience. We also remark that a PRACE-questionnaire has been used to obtain the necessary information. An EXCELLERAT-specific survey (online and on paper) must be conceived for

the project's events, not only to assess the participants' satisfaction, but also as a tool to inquire on training needs and wishes (see Subsection 4.1.2).

Other useful indicators have been measured and will inspire future reflection. First, they refer to the percentage of women involved in delivering training in the technical areas of the CoE. Even though the training events offered so far have been conducted by male personnel, among the already planned training offer, the gender ratio is as follows (% of not-male personnel involved in training):

- Visualization (Vistle at HLRS): 50% (2 main presenters, one male, one female)
- FRAUNHOFER: 33% (6 main presenters, 4 males, 2 females)
- NEK5000: one male presenter
- FEniCS: one male presenter

In terms of gender, stakeholder (industry/academia), and country, we extracted the following numbers from the participants taking part in the project's training activities:

- Gender: 14% female (i.e., KPI 5.4.3)
- Stakeholder: 87.5% academia
- Country: Not measured in detail but almost exclusively EU. Will be measured in Years 2 and 3.

Further on gender equality, the trainings produced and executed by EXCELLERAT will encourage the participation by non-male participants, which will also be reflected in the material proposed. Communication channels dedicated to gender equality (e.g., gender-equality mailing lists at universities) will be further identified and used to spread the messages. Once established, the CSA for the National Competence Centres, will be taken in the loop to distribute the EXCELLERAT activities to the diverse European countries and stakeholders.

5 HPC Provisioning

Task 5.5 "HPC Provisioning" implements the CoE hub for internal HPC resources provisioning, to support other WPs' activities. The task compiles and updates the list of available services, provides a point of contact for all CoE members to access specific services, support and documentation, and makes sure that the requests coming from the CoE are addressed by the service provisioning functions of the HPC centres. In other words, this task implements a link between the CoE and the European HPC infrastructure, including both PRACE members and EuroHPC pre-exascale prototype owners.

5.1 Provisioned resources and allocation

Computational resources have been provided to the Consortium, specifically by PRACE [37]. 0.5% of the total resources available for each PRACE call are reserved for the CoEs as selected by the European Commission under the E-INFRA-5-2015 and INFRAEDI-02-2018 calls for proposals. EXCELLERAT asked for the computational resources available for the Call 18 and 19. The assigned resources are reported in Table 1 below. For details regarding the architecture's specifications of the involved HPC systems, please see Table 3 in Annex 1: Internal Resources.

Awards (cores/hours)	PRACE 18	PRACE 19	Total
Marconi BDW	80,000	45,000	125,000
Marconi KNL	1,000,000	750,000	1,750,000
HAWK[5]		1,150,000	1,150,000
JUWELS	100,000	175,000	275,000
Joliot Curie AMD		1,715,000	1,715,000
Joliot Curie KNL		150,000	150,000
Joliot Curie SKL	350,000	180,000	530,000
MareNostrum4	700,000	240,000	940,000
Piz Daint	450,000	850,000	1,300,000

Table 6: HPC Resources allocated to Call PRACE 18 and 19

The accounts are open on whole the different clusters and the access is granted on request to the various partners.

It is still too early to present the number of utilized resources with respect to the allocated.

[5] *The planned installation of HAWK is some weeks delayed and therefore HAWK will only be available for production from 1st February 2020 on.*

5.2 KPI analysis

Table 7 presents the KPIs for this task, as defined in D5.1.

KPI number	Description	Definition	Target	Score at M13
5.5.1	Percentage of approved proposals in competitive calls for resources for EXCELLERAT projects	*N_prj_appr / N_prj_submitted * 100*	>50%	N/A
5.5.2	Number of yearly allocated cpu_hours for EXCELLERAT projects	*N_cpu_hours*	$\geq 10^6$ for PY2, $\geq 10^7$ for PY3	7,935,000

Table 7: KPIs for Task 5.5

The KPI 5.5.1 cannot be evaluated for the first year, because the Consortium has not applied for any proposal in competitive calls for resources, being the first year mainly devoted to development activities.

6 Conclusion and Time Plan

With the successful submission of deliverable D5.3, last of a long series of preparatory deliverable reports, Milestone MS3 "*First Portal Release and information about Services available*" is achieved.

In year 2 the work will continue towards the achievement of Milestone MS4 "*First verification and update phase of the EXCELLERAT services has been performed. Success stories available*" at M25. This means in particular that the service portal will evolve according the feedback received by partners and users (November 2019 Bologna All-hands meeting has already provided important insight), and deploying the services that will be progressively prepared by WP5 and all the other work packages.

7 References

[1] “Wordpress CMS website,” [Online]. Available: https://wordpress.com/. [Accessed 16 12 2019].

[2] “Drupal CMS website,” [Online]. Available: https://www.drupal.org/. [Accessed 16 12 2019].

[3] “RT issue tracker website,” [Online]. Available: https://bestpractical.com/request-tracker. [Accessed 16 12 2019].

[4] “LDAP descriptiom,” [Online]. Available: https://en.wikipedia.org/wiki/Lightweight_Directory_Access_Protocol. [Accessed 16 12 2019].

[5] “phpLDAPadmin tool website,” [Online]. Available: http://phpldapadmin.sourceforge.net/. [Accessed 08 01 2020].

[6] “Entity-Relation Diagram definition,” [Online]. Available: https://en.wikipedia.org/wiki/Entity%E2%80%93relationship_model. [Accessed 16 12 2019].

[7] “PostgreSQL website,” [Online]. Available: https://www.postgresql.org/. [Accessed 16 12 2019].

[8] “PostGIS website,” [Online]. Available: https://postgis.net/. [Accessed 16 12 2019].

[9] “SQLAlchemy website,” [Online]. Available: https://www.sqlalchemy.org/. [Accessed 16 12 2019].

[10] “Alembic website,” [Online]. Available: https://alembic.sqlalchemy.org/en/latest/. [Accessed 16 12 2019].

[11] “NGINX website,” [Online]. Available: https://www.nginx.com/. [Accessed 16 12 2019].

[12] “REST architecture definition,” [Online]. Available: https://en.wikipedia.org/wiki/Representational_state_transfer. [Accessed 18 12 2019].

[13] “Flask website,” [Online]. Available: http://flask.palletsprojects.com/en/1.1.x/. [Accessed 17 12 2019].

[14] “Swagger website,” [Online]. Available: https://swagger.io/. [Accessed 17 12 2019].

[15] “Angular website,” [Online]. Available: https://angular.io/. [Accessed 16 12 2019].

[16] “ngx-datatable website,” [Online]. Available: https://swimlane.gitbook.io/ngx-datatable. [Accessed 18 12 2019].

[17] “Mermaidjs library website,” [Online]. Available: https://mermaid-js.github.io/mermaid/#/. [Accessed 16 12 2019].

[18] “Bootstrap toolkit website,” [Online]. Available: https://getbootstrap.com/. [Accessed 17 12 2019].

[19] “GDPR EU regulation,” [Online]. Available: https://en.wikipedia.org/wiki/General_Data_Protection_Regulation. [Accessed 17 12 2019].

[20] “Cookieconsent library website,” [Online]. Available: https://github.com/osano/cookieconsent. [Accessed 17 12 2019].

[21] “Git website,” [Online]. Available: https://git-scm.com/. [Accessed 17 12 2019].

[22] “Visual Studio Code website,” [Online]. Available: https://code.visualstudio.com/. [Accessed 17 12 2019].

[23] “Docker website,” [Online]. Available: https://www.docker.com/. [Accessed 17 12 2019].

[24] “OpenStack website,” [Online]. Available: https://www.openstack.org/. [Accessed 16 12 2019].

[25] “MoSCoW method description,” [Online]. Available: https://en.wikipedia.org/wiki/MoSCoW_method. [Accessed 17 12 2019].

[26] “BSCW website,” [Online]. Available: https://www.bscw.de/en/. [Accessed 17 12 2019].

[27] “GoToMeeting website,” [Online]. Available: https://www.gotomeeting.com/. [Accessed 17 12 2019].

[28] “Etherpad website,” [Online]. Available: https://etherpad.org/. [Accessed 17 12 2019].

[29] “OpenProject website,” [Online]. Available: https://www.openproject.org/. [Accessed 17 12 2019].

[30] “Bizagi modeler website,” [Online]. Available: https://www.bizagi.com/en/products/bpm-suite/modeler. [Accessed 17 12 2019].

[31] “FocusCoE CSA website,” [Online]. Available: https://www.focus-coe.eu/. [Accessed 17 12 2019].

[32] “WCCM ECCOMAS Symposium website,” [Online]. Available: https://www.wccm-eccomas2020.org/frontal/. [Accessed 17 12 2019].

[33] “Parallel CFD school website,” [Online]. Available: https://events.prace-ri.eu/event/929/overview. [Accessed 17 12 2019].

[34] “HiDALGO CoE website,” [Online]. Available: https://hidalgo-project.eu/. [Accessed 17 12 2019].

[35] “ChEESE CoE website,” [Online]. Available: https://cheese-coe.eu/. [Accessed 17 12 2019].

[36] “HiFi-TURB project description,” [Online]. Available: https://cordis.europa.eu/project/id/814837. [Accessed 17 12 2019].

[37] “PRACE RI website,” [Online]. Available: http://www.prace-ri.eu/. [Accessed 17 12 2019].

8 Annex 1: Internal Resources

HPC Centre / HPC Infrastructure		***ARCTUR Arctur-2***	***BSC MareNostrum***	***BSC CTE-KNL***	***BSC CTE-POWER***	***BSC StarLife***
Access		Grant/Restricted/on-request	Grant	On-request	Restricted	Restricted
System Type		Sugon				
Compute	*Processor type*	Intel Broadwell	Intel Skylake	Intel Knights Landing 7230	IBM Power 9	Intel Skylake
	Total n. of nodes	30	3456	16	52	
	Total n. of cores	840	165.888	1.024	2.080	
	N. of accelerators/node	4	0	0	4	
	Type of accelerator	Nvidia M60	-	-	GPU NVIDIA V100	

Memory	*Memory / Node*	128 GB / node in compute and GPU nodes 1024 GB/node in high-memory nodes	96 GB -DDR4	96 GB	512 GB	
Network	*Network Type*	Dual Ethernet	Intel Omni-Path	Intel Omni-Path	Mellanox EDR	
	Connectivity	2x25GbE copper	Full-Fat Tree			
Home file system	*type*	nfs	GPFS	GPFS	GPFS	
	capacity	20T				
Work file system	*type*	Ceph – on demand	GPFS	GPFS	GPFS	
	capacity	100T				
Scratch file system	*type*	nfs	GPFS	GPFS	GPFS	
	capacity	20T				
Archive	*capacity*	/	On demand	On demand	On demand	

HPC Centre / HPC Infrastructure		***CINECA Marconi-KNL***	***CINECA Marconi-SKL***	***CINECA Galileo***	***CINECA DAVIDE***	***CINECA Meucci***	***EPCC ARCHER***	***EPCC CIRRUS***
Access		Grant	Restricted	Grant	Restricted	Grant	Grant	On-request

System Type		Lenovo System Adam Pass	Lenovo System Adam Pass	IBM NexTScale	OpenPower NViDIA NVLink	Lenovo System Adam Pass	Cray XC30	HPE/SGI 8600
Compute	*Processor type*	Intel Knights Landing 7250	Intel Skylake 8160	Intel Broadwell E5-2697	IBM Power 8	Intel Broadwell E5-2697	E5-2697 v2 (Ivy Bridge)	Intel Xeon E5-2695 (Broadwell)
	Total n. of nodes	3600	1512 + 792	≈1000	90	216	4920	280
	Total n. of cores	244.800	72.576+38.016	≈3.600	720	7776	109.056	10.080
	N. of accelerators/node	n.a.	n.a.	2	4	n.a.	n.a.	2 nodes (Intel Xeon Gold 6148 (Skylake)) each with 4 GPU
	Type of accelerator	n.a.	n.a.	GPU NVIDIA K80	GPU NVIDIA Tesla P100	n.a.	n.a.	NVIDIA Tesla V100-SXM2-16GB (Volta)
Memory	*Memory / Node*	96 GB – DDR4 + 16 GB - MCDRAM	192 GB/node – DDR4	128 GB/node	256 GB/node	256 GB/node	64GB (small number of 128 GB)	256 GB
Network	*Network Type*	Intel Omni-Path Architecture 2:1	Intel Omni-Path Architecture 2:1	Intel Omni-Path	Infiniband EDR 100 Gb/s	Ethernet 25/100 Gb/s	Aries	Infiniband

	Connectivity	Fat Tree	Fat Tree		2xIB EDR, 2x1GbE		Dragonfly	
Home file system	*type*	GPFS	GPFS	GPFS	GPFS	GPFS	NFS	Lustre (single file system)
	capacity	200 TB	200 TB				218 TB	406 TB
Work file system	*type*	GPFS	GPFS	GPFS	GPFS	GPFS	Lustre	
	capacity	7.1 PB	7.1 PB				4.4PB	
Scratch file system	*type*	GPFS	GPFS	GPFS	GPFS	GPFS		
	capacity	2.5 PB	2.5 PB					
Archive	*capacity*	On demand	On demand	On demand			20.29 PB	
Minimum required job size	*N.. of cores*	2 040	-				24	

HPC Centre / HPC Infrastructure		***HLRS Hawk***	***HLRS Hazelhen***	***HLRS ARM ThunderX***	***HLRS OpenPower8+***	***KTH Beskow***
Access		Grant?	Grant	On-request	On-request	Grant
System Type		HPE AMD Cluster	Cray Cascade XC40	ARM ThunderX	IBM OpenPower8+	Cray XC40

Compute	*Processor type*	AMD EPYC™ processor code named Rome	Intel Haswell E5-2680v3 2,5 GHz	1	1	Xeon E5-2698v3 Haswell 2.3 GHz and Xeon E5-2695v4 Broadwell 2.1 GHz
	Total n.. of nodes	5000	7712	96	20	2060
	Total n. of cores	640.000	185.088	n.a.	2	67.456
	N.b of accelerators/node	n.a.	n.a.	n.a.	NVIDIA Pascal P100	n.a.
	Type of accelerator	n.a.	n.a.	128GB	256GB	n.a.
Memory	*Memory / Node*	128GB	128GB	40 GbE	Infiniband DER	64 GB Haswell nodes, 128 GB Broadwell nodes
Network	*Network Type*	Infiniband HDR	Cray Aries	n.a.	n.a.	Cray Aries
	Connectivity	Enhanced hypercube	Aries™ Interconnect and Dragonfly Topology	n.a.	n.a.	Aries™ Interconnect and Dragonfly Topology

Home file system	*type*	NFS	NFS	n.a.	n.a.	AFS
	capacity	30 TB	30 TB	n.a.	n.a.	50 TB
Work file system	*type*	Lustre	Lustre	n.a.	n.a.	Lustre
	capacity	26PB	ws8 2.4PB, ws9 9PB	n.a.	n.a.	5 PB
Scratch file system	*type*			n.a.	n.a.	
	capacity			n.a.	n.a.	
Archive	*capacity*	60PB	60PB	n.a.	n.a.	On demand
Minimum required job size	*N. of cores*	4080	4080	Only for testing and porting	Only for testing and porting	512

Table 8: Initial list of internal available HPC resources