Public
Copyright © 2020 Members of the EXCELLERAT Consortium

H2020-INFRAEDI-2018-2020

The European Centre of Excellence for Engineering Applications

Project Number: 823691

D1.6
Report on the Service Portfolio

Copyright © 2020 Members of the EXCELLERAT Consortium

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

Workpackage:	WP1	Report on the Service Portfolio
Author(s):	Ralf Schneider	USTUTT
	Johanna Lucha	USTUTT
	Amgad Dessoky	USTUTT
Approved by	Executive Centre Management	
Reviewer	Ricard Borrell	BSC
Reviewer	Adam Peplinski	KTH
Dissemination Level	Public	

Date	Author	Comments	Version	Status
14-10-2020	R. Schneider	Initial draft	V0.1	Draft
20-10-2020	A.Dessoky	Initial contributions	V0.2	Draft
03-11-2020	R. Schneider	Draft ready for QA – Round 1	V1.0	Draft
13-11-2020	A. Dessoky	Second version for QA – Round 2	V2.0	Draft
27-11-2020	R. Schneider	Updates from QA – Round 2	V2.1	Draft
30-11-2020	R. Schneider	Document accepted following review	V3.0	Final

Copyright © 2020 Members of the EXCELLERAT Consortium

List of abbreviations

CFD	*Computational fluid dynamics*
D	*Deliverable*
DOA	*Description of Action*
FPGA	*Field programmable gate array*
GA	*Grant Agreement*
GPU	*Graphic processing units*
HPC	*High performance computing*
UQ	*Uncertainty quantification*
WP	*Work Package*

Copyright © 2020 Members of the EXCELLERAT Consortium

Executive Summary

The deliverable (D) *1.6 Report on the Service Portfolio* describes the update on EXCELLERAT's approach to develop and evolve its service portfolio. The initial service portfolio was defined and described in D4.1 "Application Fields and Competence Map". The methodology how EXCELLERAT will develop this initial set of services towards a portfolio of fully marketable services was initially described in D4.2. This deliverable is considered as an update of D4.2.

With respect to the evolution of the service portfolio it is elaborated on, how the methodology applied to derive further service requirements from the work done within the use cases during the first 24 project months enabled the alignment of the service requests to the perspectives of the different actors involved in the development of the Exascale engineering cycle. Further on, the deliverable presents the explicit topics of the service requests formulated by the EXCELLERAT use-cases along with their mapping to different actor perspectives and service categories which finally forms the initial marketable service layout implemented via the EXCELLERAT portal. The current deliverable will also summarize the consulting topics that driven from the user requirements and will be explained particularly in D4.6.

Copyright © 2020 Members of the EXCELLERAT Consortium

Table of Contents

Copyright © 2020 Members of the EXCELLERAT Consortium

Table of Figures

Copyright © 2020 Members of the EXCELLERAT Consortium

Table of Tables

Copyright © 2020 Members of the EXCELLERAT Consortium

1 Introduction

As already explained in D4.2, it became obvious during the first project year that the development and evolution of the service portfolio of EXCELLERAT is misplaced in WP4. First of all, there is no task dedicated to this effort in this WP4. Second, based on the WP communication structure, WP4 is not in the position to steer the development of what will basically be the foundation of the entity EXCELLERAT. Therefore, and in order to ensure that the development of the entity EXCELLERAT will include the complete consortium, this work was incorporated into the coordination efforts located in WP1 as part of an amendment to the original contract in April 2020.

This deliverable presents the progress made with the evolution of the initial service catalogue, outlined in the description of action (DOA), towards the currently envisaged marketable service layout. To this end, beside the detailed description of the envisaged marketable service layout i.e. the detailed description of the current service portfolio of EXCELLERAT, it will be first of all elaborated on the procedure how the list of service requests was gathered. It will further on be shown how the service portfolio, developed from the gathered service requests, can be categorised to target the different perspectives by which the currently identified user groups of EXCELLERAT potentially approach the centre. These topics are covered in the following sections, while the details of the work performed towards these services are presented in D4.6.

Note:

Since we will use throughout this document the abbreviations for the core-codes and their respective use-cases as defined in Table 3 “Applications and their domains” in the Grant Agreement (GA) and the amendment, the table is repeated here (Table 1) in parts to recall the abbreviations used for the core-codes and the topics of the use-cases to the reader.

Nek5000 (C1)	**C1U1**: Aerospace - Flow around aerofoil with rounded wing tip
	C1U3: Aerospace - High fidelity simulation of rotating parts
Alya (C2)	**C2U1**: Automotive/ Aerospace - Emission prediction of internal combustion and gas turbine engines
	C2U2: Aerospace - Active flow control of aircraft aerodynamics including synthetic jet actuators
	C2U3: Transport systems - Coupled simulation of fluid and structure mechanics for fatigue and fracture
AVBP (C3)	**C3U1**: Aerospace and energy - Combustion instabilities and emission prediction
	C3U2: Safety applications - Explosion in confined spaces
TPLS (C4)	C4U1: Bi-directional two-phase thin film flow – chemical engineering processes
FEniCS (C5)	**C5U1**: Aerospace and Automotive - Adjoint optimization in external aerodynamics shape optimization
FLUCS (C6)	**C6U1**: Aerospace - Design process and simulation of full equipped aero planes
	C6U2: Aerospace – computational fluid dynamics (CFD) coupling with computational structural mechanics including elastic effects

Table 1: Applications and their respective use-cases

Copyright © 2020 Members of the EXCELLERAT Consortium

2 Methodology to Evolve the Service Portfolio

To be able to evolve EXCELLERAT's service portfolio in a targeted way it was first of all necessary to develop a methodology by which additional services could be extracted from the use cases. Additionally, a way had to be found to identify the different perspectives from which EXCELLERAT will potentially be approached by the actors within the high performance computing (HPC) engineering community.

To this end, the first step was to translate the use-cases into so-called user-stories to reveal the relationship between the different actors potentially taking part in the realization of a use-case. How this was done is described below in Section 2.1 "Deduction of user stories ".

In a second step the collected user-stories were analysed with respect to the different perspectives from which EXCELLERAT will potentially be approached by the actors within the HPC engineering community. This analysis resulted in the identification of four user-groups each with a different perspective towards the services provided by EXCELLERAT. The four user-groups, i.e. perspectives are described in Section 2.2 "Service perspectives".

The third step was to analyse the collected user-stories once again, this time with respect to the service requests posed in the user-stories. This second analysis led to a list of service request that could be organized into different service categories. These categories are described in Section 2.3 "

Categories of marketable services".

The fourth step in the evolution of the service portfolio was to relate the derived service perspectives of the different user-groups with the derived service categories. This exercise finally led to the initial layout of service perspectives and service categories, which can be implemented via EXCELLERAT's service portal and that is described in Section 3 "Initial Service Layout".

2.1 Deduction of user stories

As mentioned above, the first step in the evolution of the service portfolio was to translate the use-cases into so-called user-stories. While the use-cases focus on specific technical problems to be solved, the user-stories focused on the stakeholders engaged in the realization of a use-case and their relationships. Typical stakeholders are e.g. the (end-) user, and the technology (or know how) provider. This approach allowed the members of the consortium, some of whom contributed both, the use cases and the core codes used in them, to put themselves in the shoes of EXCELLERAT's potential external customers.

As an example, for this approach we will consider C1U3 - High fidelity simulation of rotating parts with NEK5000. The constellation of actors participating in the use-case as it was originally set up is depicted in Figure 1. Shortly the interaction of the three parties can be described as follows: A helicopter manufacturer is looking for a solution to execute high fidelity simulations of helicopter rotors. CINECA who has experience with NEK5000 would be able to come up with such a solution even though there is a lack of experience with respect to some needed features like e.g. adaptive mesh refinement and high order meshing for complex geometries which KTH, as one of the code owners and main developers of NEK5000, can provide.

Copyright © 2020 Members of the EXCELLERAT Consortium

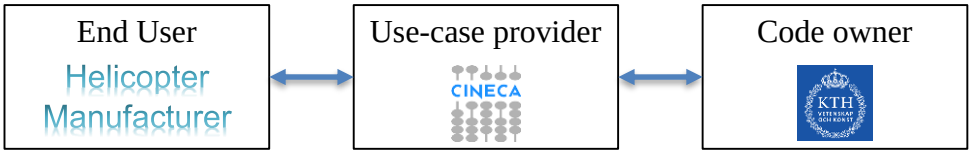

Figure 1: Initial actor constellation and information flow in C1U3.

After starting to work on the use-case and discussing with the EXCELLERAT consortium the different perspectives from which one is able to shed light onto this use-case from an outside point of view, CINECA was able to develop seven user stories with whose problems and resulting requirements EXCELLERAT could potentially be confronted once it is established to be the European knowledge hub for Exascale engineering applications. To illustrate the user-stories three examples will be given:

- First of all, it would be possible for EXCELLERAT to be approached directly by the industrial end-user. This end user would like to execute high fidelity simulations of rotating parts with an open-source simulation code. Moreover, he is looking to find a suitable workflow to execute this Exascale engineering problem efficiently.
- Secondly, it would be possible for EXCELLERAT to be approached by the use-case provider who is already collaborating with the industrial end-user outside the framework of EXCELLERAT to provide the requested workflow and is now searching for expertise in robust strategies for rotating parts modelling at high Mach numbers in spectral element codes since NEK5000 is proven only at low Mach numbers.
- The third example shows how the change of perspective revealed further collaboration opportunities within the EXCELLERAT consortium and enabled the generalized extension of the service portfolio beyond the original goals of the use-cases. Thus, it is possible that an end-user, building on the already successfully implemented original use-case, also performs the simulation of rotating parts and finds in this case that it would significantly increase the efficiency of the method if a special evaluation program could be coupled in-situ to NEK5000. The search of the end-user for required expertise would expand the circle of actors to include the code developers of a corresponding visualization program as well as the developers of a data analytics library to be used, who would be able to provide the corresponding know-how.

By these three examples it can already be recognised that on the one hand, the problems and requirements derived from the user-stories are not necessarily focused to the narrowed scope of the use-cases but can be formulated to fit in a more generalized scope in which a service can be provided that targets the engineering community. On the other hand, the requirements can be categorized and associated with different user-groups like application end-users and code developers.

2.2 Service perspectives

From the analysis of the user-stories, as the second step in the evolution of EXCELLERAT's service portfolio, it was possible to identify four user-groups. Each of this user-groups represents a perspective by which EXCELLERAT can be potentially approached.

To illustrate this, the example of service requests for consulting on in-situ visualization is taken. The analysis of the user-stories showed, that consulting on in-situ visualization was requested several times within different user-stories. On the one hand, there was a detailed service request regarding implementation methods and efficient procedures for the coupling of in-situ libraries to a code with elements of high order. On the other hand, the request was formulated with

respect to the handling and concrete implementation of a special application case with an already in-situ capable code. This means that from this example one can already extract two service perspectives. The first request is obviously posed from the perspective of the code-developer while the second one is obviously posed from the perspective of the end-user.

By analysing all user-stories following the strategy demonstrated by the example above, the following four perspectives from which services can be requested from EXCELLERAT could be extracted:

1. **The perspective of the application end-user.**

   If EXCELLERAT is approached from this perspective, it should provide services that deliver solutions to simulate a given engineering problem by executing an Exascale engineering cycle. This perspective will further on be named "**Solution evolution**". It should be explicitly noted that at this point solution means the processing of an engineering problem along the entire Exascale engineering cycle and not just the solution of a discretized problem in the sense of executing a simulation code or a numerical solver. Under this perspective, all services are united that will evolve the solution of engineering problems towards Exascale.

2. **The perspective of the code-developer.**

   If EXCELLERAT is approached from this perspective, it should provide services that deliver expertise in areas that enable the code-developer to evolve engineering software packages towards extreme-scale applicability. This perspective will further on be named **"Code / Application evolution"** with application meaning all software components being used in implementing the engineering cycle. Under this perspective, all services are united that will evolve engineering codes / applications towards Exascale by means of massive parallelism and extreme scalability while maintaining the applicability to real-world problems and improving the ease of use for the end-user.

3. **The perspective of the vendor.**

   In this context vendor refers to the group of system integrator, hardware developer, system-software developer or hardware vendor. If EXCELLERAT is approached from this perspective, it should provide services that deliver input and discussion partners to evolve the future HPC-systems to be ready to efficiently execute the Exascale engineering cycle.

4. **The perspective of the HPC engineering community.**

   In this perspective any other actor who is already or wants to become part of the HPC engineering community is included.

   From this perspective EXCELLERAT's consortium will provide training up on request related to each phase of the simulation workflow towards Exascale computing. Moreover, events like conferences, and symposiums will be regularly organized to keep all the stakeholder always together, ensure the transfer of knowledge and avoidance of duplicated efforts.

The development of the services perspectives described above will help to develop a strategy for the targeted user guidance within the EXCELLERAT portal i.e. help to develop a strategy for the targeted presentation of the services to EXCELLERAT's user community in general. This must be seen as particularly important as it will help to conduct a more fine-grained and targeted market analysis that will significantly contribute to the development of a sustainable business model.

Copyright © 2020 Members of the EXCELLERAT Consortium

2.3 Categories of marketable services

As already mentioned in the introduction of Section 2, the third step towards the evolution of EXCELLERAT's service portfolio was to analyse the user stories with respect to the service requests posed in the user-stories. This second analysis led to a list of service requests that could be organized into six different service categories. Four of these categories concern different types of consulting, one targets the provisioning of tools for the Exascale engineering cycle and another one the provision and hosting of datasets. The detailed description of the service categories is given in Sections 2.3.1 - 2.3.3 below. The service requests posed from each of the four service perspectives in the respective service category are afterwards listed in Section 3.

2.3.1 Consulting service categories

The most prominent service category for which all user-stories contained requests was consulting i.e. know-how transfer in its different appearances. The currently identified different appearances of consulting were further on separated into four different types of consulting:

1. **Active one-on-one consulting**

   In this type of consulting an expert or a group of experts helps the service inquirer to develop or execute / implement a feature, problem, method etc.

2. **Passive one-on-one consulting**

   With this type of consulting it is referred to best-practise guidelines which deal in a technically and detailed way with the execution / implementation of a feature, problem, method etc.

3. **One-on-n consulting**

   This type of consulting refers to training courses.

4. **n-on-n consulting**

   In the form of access to community events like expert networks, specific workshops or targeted symposia.

In turn, these four types of consulting could in part be related to all of the perspectives described in Section 2.2 even though additionally a relation between the requested type of consulting and the inquirer's background i.e. affiliation became visible. E.g. requests originating from actors with academic affiliation and the code-developer perspective mostly targeted consulting requests in form of best practise guidelines, training courses and dedicated workshops whereas requests from actors with an industrial affiliation and the end-user perspective mostly targeted one-on-one consulting and specific training courses.

Copyright © 2020 Members of the EXCELLERAT Consortium

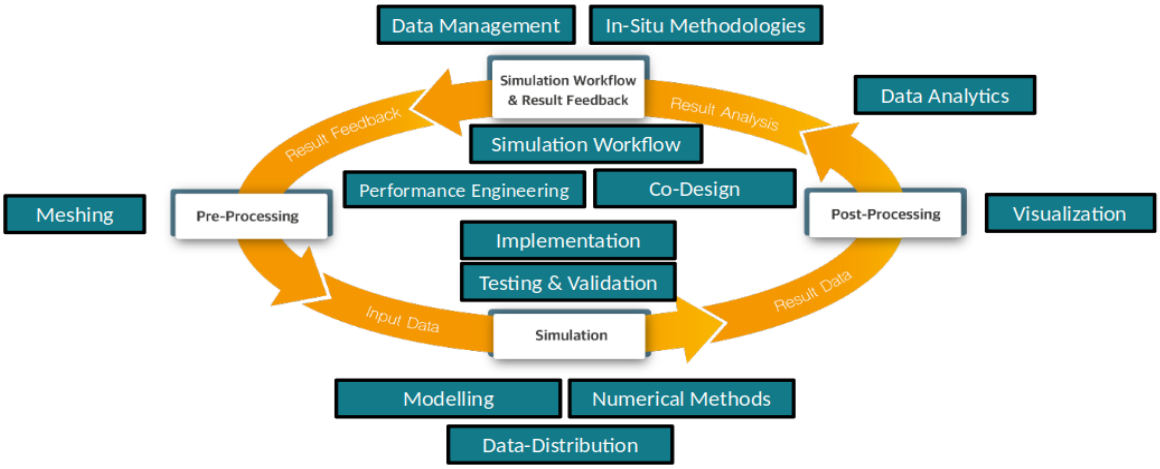

Figure 2: Consulting service scheme

Here we will give overview on the consultant topics that will be discussed in details in D4.6. Driven from the use-cases and based on the user requirements for the entire simulation workflow, the consulting scheme is considered as shown in Figure 2. Therefore, EXCELLERAT provides expertise and consulting for the efficient execution of the complete large scale engineering simulation cycle. This includes not only the three classical process steps "Pre-Processing", "Simulation" and "Post-Processing" but also ranging over external as well as intrinsic optimization and uncertainty quantification methodologies and multi code coupling methodologies up to advanced data analytics methodologies in engineering simulation as the following:

1-Meshing

EXCELLERAT provides expertise and consulting for meshing techniques, methodologies and software packages, including mesh adaption methods for structured and unstructured grids, front tracking and mesh generation for high order meshes. Alongside the mesh generation expertise can also be provided with respect to efficient CAD-geometry generation e.g. by OpenCASCADE.

2-Visualization

EXCELLERAT provides expertise and consulting with respect to the setup, implementation and execution of visualization methodologies, tools and software packages for result data of largescale engineering simulations including not only classical post-processing but also in-situ methodologies. Featured tools and libraries are currently, Vistle, ParaView, and Catalyst.

3-Data Analytics

EXCELLERAT provides expertise and consulting in data analytics tailored to the field of engineering. This includes interfaces between general Machine Learning libraries (e.g. skit-learn and Tensorflow) as well as proprietary tools developed by the partners of the consortium and the engineering workflow. Additionally, expertise can be provided with respect to different kinds of data representations e.g. wavelets, spectral bases or data-driven decompositions along with their applicability as input features for Machine Learning and their ability to separate physical effects.

4-Data Management

EXCELLERAT provides expertise and consulting for management and storage of large scale data sets originated from large scale engineering simulation workflows. This includes expertise in parallel Input and Output (I/O) strategies, handling of fine grain large scale datasets as well as advanced data compression methods for simulation result data on structured and unstructured

grids. Additionally, consulting in the area of the customer to data center interface can be offered including efficient piecewise update and data encryption based on Merkle Tree technologies.

5-Implementation

EXCELLERAT provides expertise and consulting for efficient code implementation strategies and implementation strategies for efficient, Exascale ready code. This includes not only implementation from scratch and porting of legacy applications to modern hardware but also the efficient programming for accelerated and heterogeneous system architectures with featuring hierarchical memory architectures and storage systems.

6-In-Situ Methodologies

EXCELLERAT provides expertise and consulting for efficient and universal implementation and execution of in-situ methodologies for visualization, data analytics and also multiphysics coupling. The currently targeted framework which is generalized by the EXCELLERAT In-Situ Methods Working group is Catalyst as discussed in D4.6.

7-Data-Distribution

EXCELLERAT provides expertise and consulting in advanced data-distribution and load-balancing strategies for various discretization methods like Finite-Elements, High-Order Spectral-Element methods or Finite-Volume methods both on structured as well as unstructured grids. Additionally, consulting with respect to dynamic mesh-refinement and the efficient redistribution strategies can be offered including graph based methods, kd-trees and space filling curves.

8-Modelling

EXCELLERAT offers expertise and consulting in all aspects of modelling and simulation of engineering problems. Starting with the description of a given physical problem by corresponding equations, over the discretization of partial differential equations up to the technical realization of a simulation task by means of a given software package.

9-Solution Methods

EXCELLERAT offers expertise and consulting in all aspects of numerical treatment of simulation problems in engineering. This includes expertise in linear and nonlinear solution procedures like direct-sparse solvers, Krylov-Subspace methods, Multi-Grid methods and Newton methods among others. Additionally, expertise can be offered in the efficient numerical treatment of adjoint optimization, muti-physics coupling and explicit and implicit numerical solution procedures.

10-Testing and Validation

EXCELLERAT provides expertise and consulting with respect to testing and validation of simulation codes, implementation of pre- and post-processing methodologies and complete simulation workflows. This includes access to datasets for the purpose of validation and verification, technology testing and benchmarking as well as consulting with respect to configuration setups, benchmark configurations and result verification.

11-Co-Design

The EXCELLERAT consortium provides consulting for software and systems co-design by means of its cross-competence co-design working group as discussed in D4.6. On the software co-design side, based on the reference codes and other relevant applications and libraries of the engineering workflow, the operations, kernels and algorithmic features are characterized that are common and widely used in engineering applications and which demand large amounts of

computational time. On the hardware side EXCELLERAT is cooperating via its interest groups with original equipment manufacturers and system integrators considering the complete hardware bandwidth from standard x86_64 architectures over ARM, graphic processing units (GPUs) and NEC Vector processors down to field programmable gate arrays (FPGAs).

12-Performance-Engineering

EXCELLERAT provides expertise and consulting in performance engineering for Eexascale engineering applications in all areas of the engineering simulation workflow. This means we not only provide performance engineering in terms of floating point operations per second in selected kernel operations but also in terms of efficient execution of production stage applications executing real world use-cases. This includes also I/O bandwidth evaluation and optimization, data transfer performance and efficiency and performance with respect to data analytics. For performance analyses we team up with POP2 the Center of Excellence in Performance Engineering.

2.3.2 Provisioning of tools

Besides consulting services, the analysis of the user stories with respect to the service requests revealed that a demand for provisioning of dedicated tools for the Exascale engineering cycle exists. As already defined in the grant agreement [1] these tools complete the Exascale engineering cycle with respect to pre- and post-processing, data management and usability. The persisting request for these tools confirmed the consortium in its decision to grant the enhanced services, developed in Task 4.1 – Task 4.5 the envisaged prominent position within the evolution of the service portfolio. The work conducted in the tasks of WP 4 along with their envisaged relation to the service category “Provisioning of tools” is explicitly described in D4.6.

2.3.3 Dataset hosting

From the analysis of the user-stories with respect to the service requests as well as from ongoing discussions within the consortium the service request to host and provide large data sets was extracted. Even though the specific purpose of the datasets to be hosted was not mentioned in the user-stories from the perspective of the end-users, the requests posed from the code-developers’ perspective in that respect were more clear. Since in almost every use-case a dataset for validation or verification is used this led to the conclusion that in the first place such datasets should be targeted. These datasets, either from reference simulations or measurements, used to proof the correctness of the implemented methods, models, algorithms and approaches will be beneficial for code-developers as well as for end-users, trying to verify and validate a model setup.

The overview about the data-sets currently used by the use-cases as input and validation data is given in the list below.

- NASA CRM Modell at stall. [2]
- JAXA high lift configuration standard model. [3]
- Injection problem - Menard et.al 2007 [4]
- Configuration of gas turbine combuster [5]
- NACA0012 [6]

The current idea is to set up a repository for datasets and data sources which can be directly hosted by EXCELLERAT or linked into the repository from external sources. In addition to the

Copyright © 2020 Members of the EXCELLERAT Consortium

datasets directly used by the use-cases, standard cases like e.g. the Taylor Green Vortex [7] can be added to the repository.

Further on the idea exists to provide access via the repository to reference result datasets produced in the frame of EXCELLERAT even though the discussion whether it is reasonable to host complete datasets is still ongoing since this approach would need a significant amount of resources.

Copyright © 2020 Members of the EXCELLERAT Consortium

3 Initial Service Layout

After having elaborated on the perspectives under which services are requested in Section 2.2 and on the service categories that resulted from the complete list of service requests in Section 2.3, this section introduces the currently selected layout for the service implementation which is shown in Figure 3. The service portal is reachable under https://services.excellerat.eu/. Its implementation details along with the description on how the theoretical layout presented in is technically realized can be found in D 5.2 and its update D5.4.

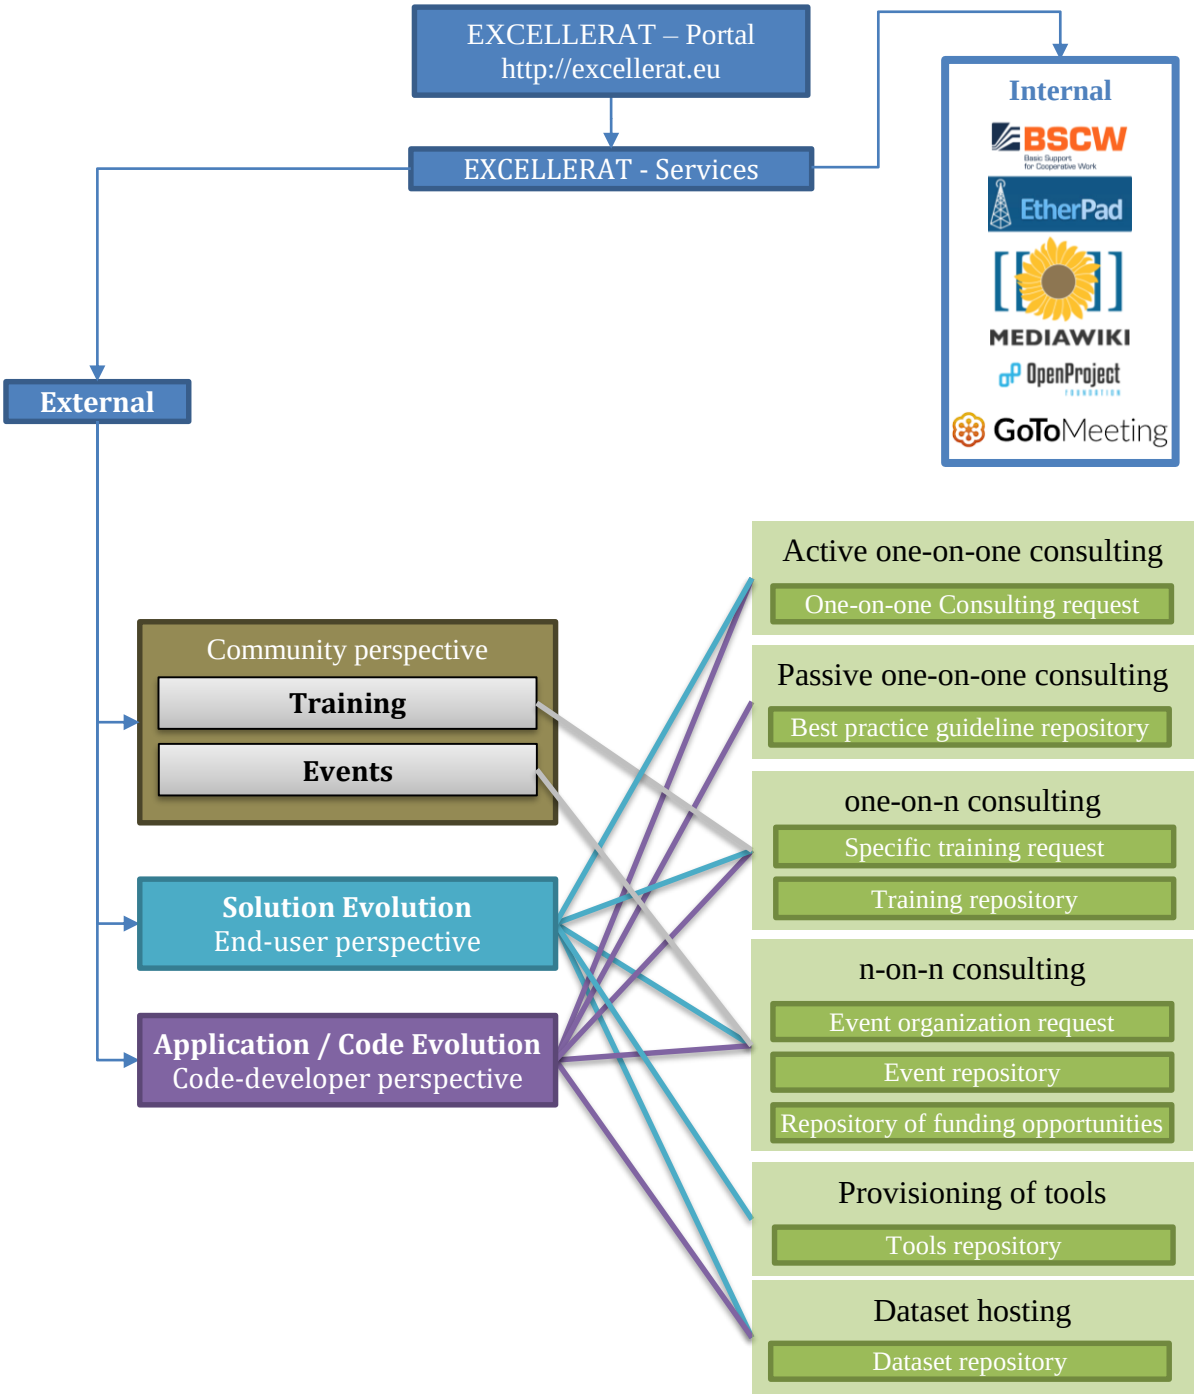

Figure 3: Initial service layout

As can be seen, the EXCELLERAT service portal will be accessible form the EXCELLERAT homepage. From the service portal the external and the internal services can be accessed.

Copyright © 2020 Members of the EXCELLERAT Consortium

Currently the internal services are only accessible for members of the EXCELLERAT consortium. Since the internal services are currently not considered as marketable services they are listed in this deliverable only for completeness. Currently the following list of tools for collaborative work and project management is provided to the members of the consortium:

1. BSCW: File sharing.
2. EtherPad: Collaborative text editing.
3. MEDIAWIKI: Internal documentation.
4. OpenProject: Project Management.
5. GoToMeeting: Conference Calls and Webinars.
6. Nextcloud installation with integrated OnlyOffice

As it for data protection reasons critical for several partners of EXCELLERAT's consortium to do collaborative document editing based on google docs it was decided to use the Nextcloud installation at HLRS with integrated OnlyOffice.

If a non-consortium user enters the service portal she or he will be guided to the external services where the user can decide from which perspective the services of EXCELLERAT should be approached.

As shown in Figure 3, if a user wants to approach EXCELLERAT from general perspective of the HPC engineering community she or he can further choose to approach the training, i.e. the one-on-n consulting service category or the events, i.e. n-on-n consulting service category. This further differentiation was done since training is the most advanced and elaborated form of know-how transfer. This was taken into account already during the proposal writing phase by creating the dedicated task 5.4 and now by implementing a dedicated approach perspective.

Further on the differentiation within the community's approach perspective was done since from the perspectives of the end-user, the code-developer and the vendor know-how transfer in form of n-on-n consulting was requested. This means all three groups that approach EXCELLERAT with focused service requests would like to interact with the community. Conversely, the conclusion was drawn that the community should be given the opportunity to interact with these groups via direct this access to the n-on-n consulting service category.

In Figure 3, it can further on be seen that not all service perspectives are connected to all service categories, e.g. the solution evolution perspective is not connected to the passive one-on-one consulting service category. This resulted from the fact that not from all service perspectives service requests were posed in all service categories. In fact, since all use cases are driven by the requirements of the user perspective, which in turn induce the requirements from the perspective of the code developers, the requests arising from the perspective of system evolution i.e. the vendor are at the moment relatively underrepresented. Due to that the implementation of this service perspective will be postponed and it is not shown at all in the initial service layout.

In addition to the connection of the service perspectives to the service categories in Figure 3 the services which are currently planned to be implemented under each service category are shown. It can be seen, that initially nine different services will be set up of which six are planned to be implemented as tagged repositories and three as service request forms.

After the presentation of the first service layout, in the following Sections we will elaborate on in more detail from which service perspective in which of the various service categories service requests were posed.

3.1 Solution evolution

As stated above the requests targeting the perspective of the end-user mainly concerned the processing of an engineering problem along the entire Exascale engineering cycle and not just the solution in the sense of executing a simulation code or a numerical solver. This means that the requests assume that there is already a method or toolchain for the given problem to be solved and that consulting is "only" needed in relation to the application of the method, the process chain or its components i.e. tools. If, on the other hand, the further development of individual components of the method becomes necessary i.e. implementation and code-development has to be conducted, the requests target one-on-one consulting by the code developer in which it is assumed that the developer does the actual development work until a ready to use solution exists. As examples for these statements, C1U3 and C5U1 can be taken into account. In these use-cases, CINECA as the use-case owner is provided consulting services by KTH – the code-owner – in terms of code usage and case-setup. Additionally, implementation work has to be conducted. In the case of NEK5000 adaptive mesh refinement methodologies have to be further developed and implemented and in the case of FEniCS matrix assembly routines have to be modified to make the code usable in the targeted Exascale uncertainty quantification scenario. Both implementations are not carried out by CINECA but as said by the code-owner KTH.

In addition to the insights discussed in the previous sections, the analysis of the user stories and subsequent discussions within the consortium indicated that the requests made from the end user's perspective are much more specific than those made from the other perspectives. Meaning that the requests are much more targeted towards the solution of given problems by specific codes and approaches rather than towards more generic methodologies. These considerations, as well as the assumption that later most of these request will originate from industrial end-users, lead to the conclusion that for this perspective best practise guidelines even though requested by the user-stories would have two major disadvantages:

1.) The production of best practise guidelines in advance would require a significant amount of person month to be invested now. But from these efforts guidelines with a relatively narrower scope would result in a limited impact to the community. Even though, ideas exist how the production of best practises guidelines targeting the end-user perspective can be part of the business model to be developed.

2.) If the production of guidelines would be postponed until a specific request is posed, i.e. The production of the guidelines would be done on-demand, it would require too much time and due to that would most likely be superseded by one-on-one consulting activities.

In view of these two disadvantages, the implementation of the consulting category 1 from the end-user’s perspective will be postponed. Additionally, the analysis of the service requests posed from the end-user perspective gave on little differentiation between the service categories 2 and 3. Due to that, for the remainder of this section, we will disregard the difference between the first three consulting categories and only list requests targeting the fourth service category separately.

3.1.1 Service requests targeting one-on-one and one-on-n consulting

- Visualization methods for specific engineering work flows (C1U1, C1U3, C5U1).
- Data analytics as pure post processing and in-situ (C1U1, C1U3, C5U1).

Copyright © 2020 Members of the EXCELLERAT Consortium

- Visual analytics (C1U1, C1U3, C5U1).
- Data management for large data sets (C1U1, C1U3, C5U1).
- Meshing techniques/software (C1U1, C1U3, C5U1).
- Guidelines for scalable simulation workflows (C1U1).
- Optimal use of resources, when running uncertainty quantification (UQ) enhanced simulation (C1U1, C1U3, C5U1).
- Service to advise on runtime approach i.e. how to obtain best efficiency vs. elapsed time trade-off (C1U3, C5U1).
- Application of mesh adaption methods and for front tracking (C3U2).
- Application of multi-code coupling technologies (C2U2).

Information about and access to n-on-n consulting activities such as conferences, workshops and symposia were requested not that frequently from the end-user perspective compared to code developers and hardware manufacturers perspectives.

From the end-users' perspective especially the topic of a cross-competency experts panel with respect to multi-physics model configurations was requested.

3.1.2 Provisioning of tools

As already mentioned in Section 2.3.2, besides consulting services that a demand for provisioning of dedicated tools for the Exascale engineering cycle exists and that this confirmed the consortium in its decision to grant the enhanced services, developed in task 4.1 – task 4.5 their prominent position further on. From the perspective of the end-user especially requests targeting the work done in task 4.2 – task 4.5 were posed.

3.1.3 Dataset hosting

As already mentioned in Section 2.3.3, from the end-users' perspective service requests for dataset hosting were posed even though the specific purpose of the datasets to be hosted was not mentioned.

3.1.4 Services to be provided from the end-users' perspective

The summary of services that are envisaged to be provided from end-user-perspective are shown in Figure 4. Namely, these are:

- Management of requests targeting one-on-one consulting activities.
- Management of requests for special training courses or events.
- A training courses repository listing trainings.
- Management of requests about organization of special events like cross competency expert panels.
- An event repository listing community events like workshops and symposia targeting specific topics.
- A tools repository providing access to Exascale ready tools for the engineering cycle. If tools are developed within the framework of EXCELLERAT or can be made available directly by EXCELLERAT due to their licensing, their provision will be included in this service.

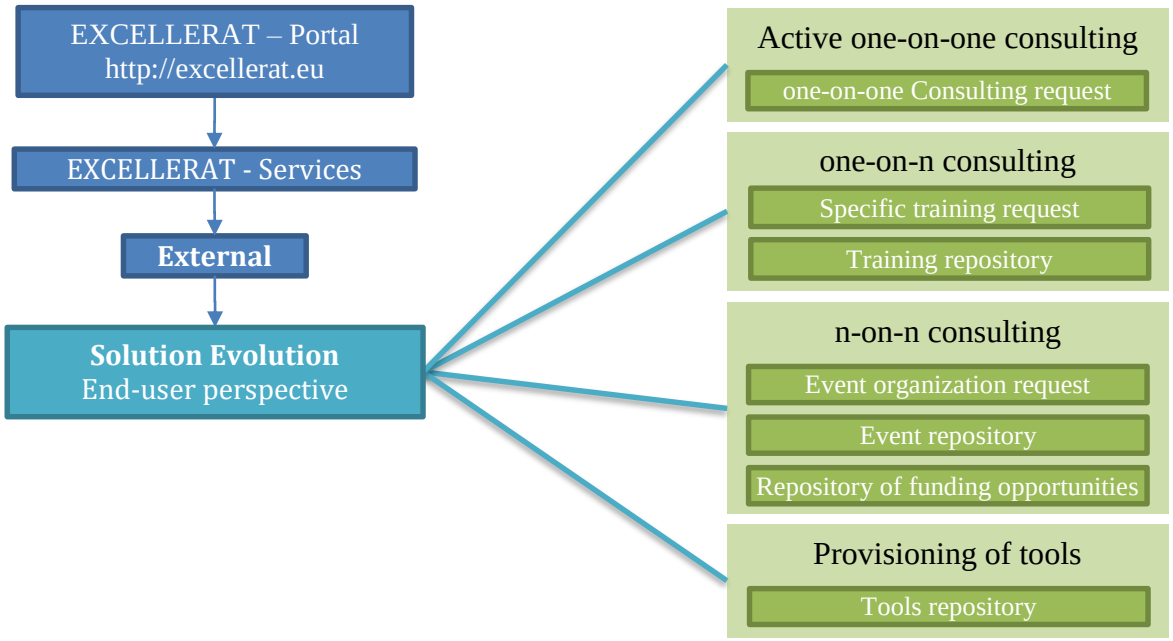

Figure 4: Currently envisaged list of services from the perspective of the end-user

3.2 Application / Code evolution

As stated above, from the code developers' perspective in the first place best practise guidelines with rather generic topics targeting methodologies were requested. Discussions within the consortium on this topic showed that the focus on basic methods is due to the fact that highly scalable applications have very specific implementation strategies. For instance, a mesh partitioning scheme can be transferred from one code to another if both codes feature the same or a similar mesh topology description even though this does not apply to its explicit implementation. In most cases, the direct transfer of an implementation will fail due to the programming language and code-specific data structures used. For this reason, code developers are apparently not interested in explicit guidelines tailored to a code, but in descriptions of methodological procedures. However, it should be noted that a clear difference between scientific publications and the requested best practice guidelines could be recognized during the discussions. In this case, methodological description does not mean the mathematically correct formulation of a numerical method, but rather the technical methodological description of a procedure of the type: "The network partitioning library was used in the following way in connection with a CFD finite volume code and the following results were achieved with regard to efficiency, etc.".

3.2.1 Service requests targeting passive one-on-one consulting

Explicitly the following list of topics for best practice guidelines was extracted from the requests formulated in the user-stories. For reference the core-code along with the use-case to which the request applies is given in brackets behind the respective topic:

- Porting of legacy applications to modern hardware (C5U1).
- Know-how about efficient data transfer/management (C1U1, C1U3).
- In-situ analysis workflows (C5U1).
- Development of automated process chains (C6U1).
- Implementation for In-situ methods (C1U1, C1U3, C5U1).

Copyright © 2020 Members of the EXCELLERAT Consortium

- Error indicators and estimators (C1U1, C1U3, C5U1).
- Using accelerators/heterogeneous systems (C1U1, C5U1).
- Efficient data redistribution (C1U1, C1U3, C5U1).
- Testing and validations procedures - a debug procedure (C3U2).
- Handling of adaptive mesh refinement with physical constraints (C2U1).
- Multi-code coupling technologies and implementation strategies (C2U2).
- Two-layer wall models (C2U2).
- Optimal use of resources, when running UQ enhanced simulation (C1U1, C1U3, C5U1).

3.2.2 Service requests targeting active one-on-one and one-on-n consulting

In this section, the overview of the requested consulting topics which were categorised in the active one-on-one and also the one-on-n consulting categories from the code-developer perspective is given.

First of all, the analysis of the requested topics showed the expected demand for expertise in the direction of performance engineering, code efficiency improvements as well as topics connected to those like programming models and knowledge about hardware. Explicitly the following topics were named:

- Detailed performance analysis for identification of code sections with high potential for runtime improvements (C6U1).
- A benchmark ready to test the load balancing performances (C3U1).
- Know-how of using accelerators/heterogeneous systems (C5U1).
- Expertise to select and implement suitable programming models for enhanced code efficiency (C6U1).
- Service to advise on runtime approach (how to obtain the best efficiency vs elapsed time trade-off) **(**C1U3, C5U1).

Beside the requests for performance engineering, which targeted directly the codes, there were some requests identified, which were collected under the topic of “system-application-interaction”. This means, these requests cannot be handled by one performance engineering of a single application but could be tackled with an integrated approach based on the integration of the system’s and the application’s performance monitoring and error handling. Explicitly, the following topics were named:

- A detailed performance assessment that gives pertinent figures at the end of a year of production (C3U1).
- Error handling strategy that goes beyond codes, to relay precise error messages (C2U1, C3U2).

The third area in which topics for one-on-one consulting and training were requested from the code developer perspective was concerned with numerical and algorithmic methods.

- Smart load balancing (C3U1, C1U3, C1U1).
- Multi-code coupling technologies (C2U2).

Copyright © 2020 Members of the EXCELLERAT Consortium

- Mesh adaption techniques for front tracking (C2U1, C3U2).

Furthermore, an interesting request was made for temporary technical project management, which takes over the coordination and organisation of complex code development tasks e.g. like the ones originating from the implementation of end-user solutions. Even though currently no final strategy is developed to implement this as a dedicated service it will for now be continued under consulting and also being kept on the agenda of the business development task.

3.2.3 Service requests targeting n-on-n consulting

From the code-developers' perspective several requests directly targeted access to n-on-n consulting events such as conferences, workshops and symposia even though most of the topics additionally appeared also in requests towards the other consulting categories. This was not that much of a surprise since many of the code-developers of highly scalable engineering applications are affiliated with academic research institutes in which the know-how exchange via such events is more present than within industry. Explicitly the requests targeted the following topics

- New architectures (C1U3, C5U1).
- Optimization method (C3U2).
- HPC resources (C3U2).
- Numerical methods (C3U2).
- Programming models for enhanced code efficiency (C6U1).
- Co-Design.

Even though not directly named within the user stories, the topic of Co-Design was also grouped under the code-developers' perspective since based on ongoing discussions within the consortium about the topic in addition to the dedicated task, a transversal working group was established. This group is comprised of members from different partners, tasks and WPs to discuss their approach and findings with hardware vendors, system integrators etc. This means for this group a service which provides dedicated workshops and connection to the system and hardware developers would be of great help.

3.2.4 Dataset hosting

As stated in Sections 2.3.3 and 3.1.3 the service to host and provide datasets for validation and verification was extracted from the requests posed in the user-stories from the end-users' perspective but was also derived from the ongoing work carried out within the use-cases.

3.2.5 Services to be provided from the code-developers' perspective

The summary of services that are envisaged to be implemented from the code-developers' perspective are shown in Figure 5:. Namely, these are:

- Management of requests targeting one-on-one consulting activities.
- A repository of best practise guidelines.
- Management of requests for special training courses or events.
- A training courses repository listing trainings.
- Management of requests about organization of special events like cross competency expert panels.

Copyright © 2020 Members of the EXCELLERAT Consortium

- An event repository listing community events like workshops and symposia targeting specific topics.

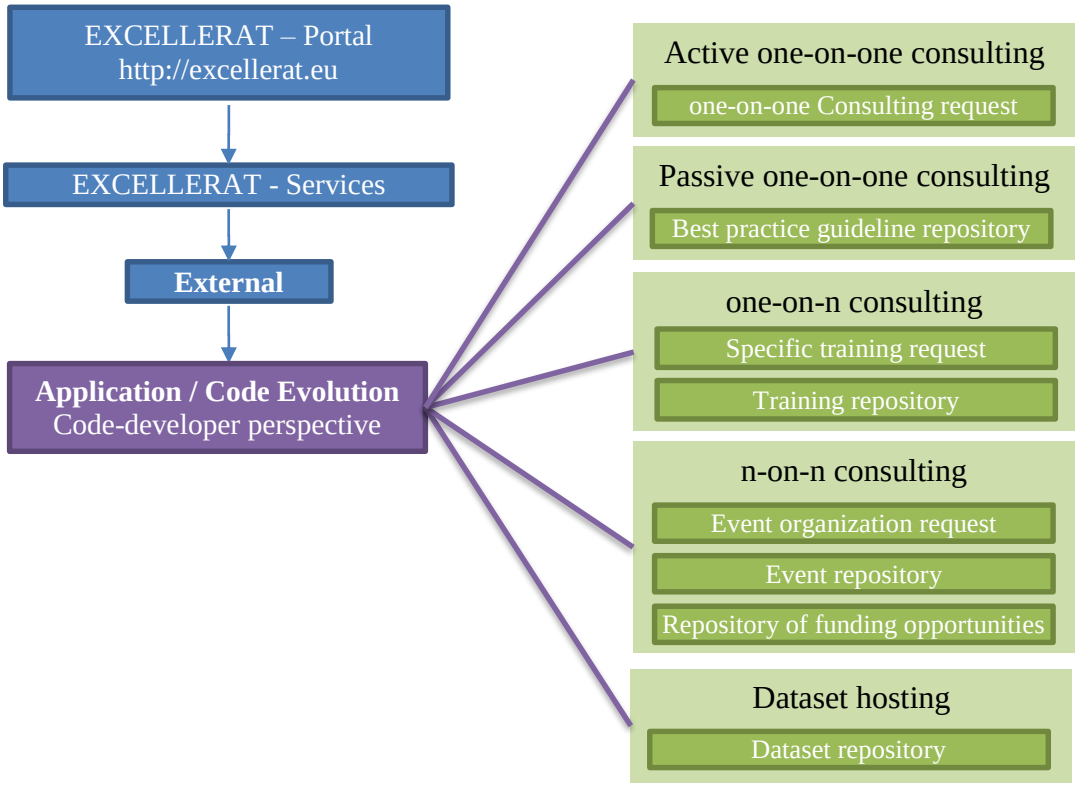

Figure 5: Currently envisaged list of services from the perspective of the code-developers

3.3 System evolution

Since all use cases are driven by the requirements of the user perspective, which in turn induce the requirements from the point of view of the code developers, the requests arising from the perspective of the system integrator, hardware developer, system-software developer or hardware vendor are at the moment relatively underrepresented. Currently the three topics directly related to system and hardware engineering are:

- Co-design.
- Detailed performance assessment that gives pertinent figures at the end of a year of production.
- Error handling strategy that goes beyond codes, to relay precise error messages.

Since these topics are already dealt with from the user and code developer perspective, the implementation of the vendors' service perspective is postponed for the time being. Once task 4.1 and the Co-Design working group are producing meaningful output, the implementation of the respective consulting and community services will be done.

3.4 Training

Since training is the most advanced and elaborated form of know-how transfer in all of the areas mentioned so far, this was taken into account during the proposal writing phase by creating a special task for it and now by implementing a dedicated option to approach the training services from the community perspective. Special topics for training courses requested within the user-stories covered so far:

Copyright © 2020 Members of the EXCELLERAT Consortium

- Using accelerators/heterogeneous systems (C5U1).
- Porting legacy F77 codes to modern hardware (C5U1).
- External aerodynamics and adjoint-based optimization (C5U1).
- Implementation of scalable simulation workflows (C1U1).
- Multicode coupling technologies (C2U2).
- Load balancing libraries (C3U1).
- Mesh adaption libraries (C2U1, C3U2).
- Adaptive mesh refinement techniques (C2U1).
- Two layer wall models (C2U2).
- Simulation of rotating parts and meshing (C1U3).

3.4.1 Services provided within the training option of the community perspective

The summary of services that are envisaged to be provided under the training category are shown in Figure 6. The services to be implemented are:

- Management of requests for special training courses or events.
- A training courses repository listing trainings.

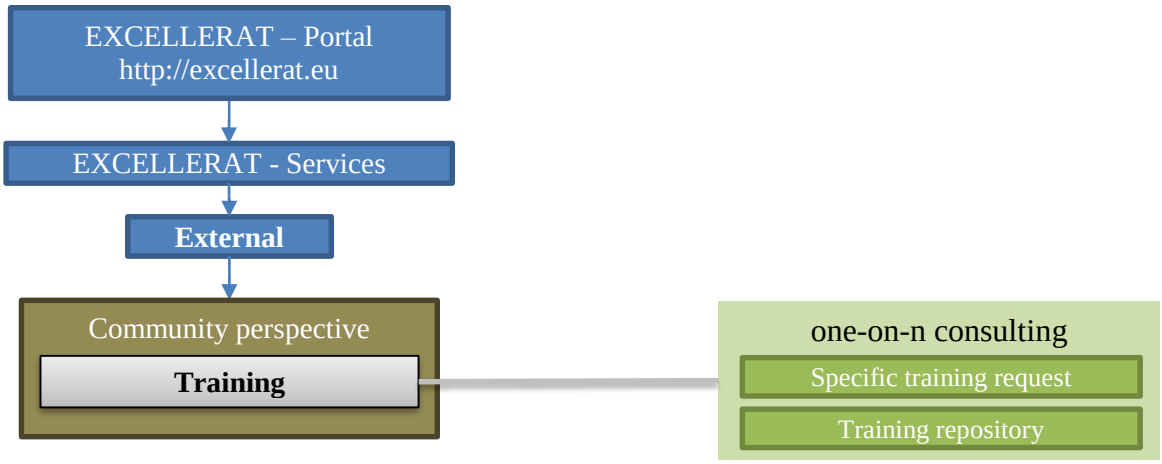

Figure 6: Currently envisaged list of services to be provided within the training category

3.5 Events

Since from all three perspectives end-user, code-developer and hardware / system provider know-how transfer in form of n-on-n consulting was requested, this category of services will, like training, also be covered separately. Specific requests for networking events and access to expert networks were posed with respect to the following topics:

- Accelerators/heterogeneous systems.
- Modern and future hardware.
- Optimization methods.
- HPC resources.
- Numerical methods and algorithms.

- New programming models.
- Physical modelling.

Additionally, besides the requests extracted from the user-stories the two service topics envisaged in the original proposal:

- Support access to industry funded research opportunities;
- Promoting outstanding applications i.e. exascale demonstration runs as well as high capacity runs

will be implemented as soon as respective opportunities can be linked and results of outstanding applications within the framework of EXCELLERAT are produced.

3.5.1 Services provided within the events option of the community category

The summary of services that are envisaged to be provided under the community category are shown in **Error! Reference source not found.**. Namely, these are:

- Management of requests about organization of special events like cross competency expert panels.
- An event repository listing community events like workshops and symposia targeting specific topics.
- A repository with funding opportunities targeting HPC engineering.

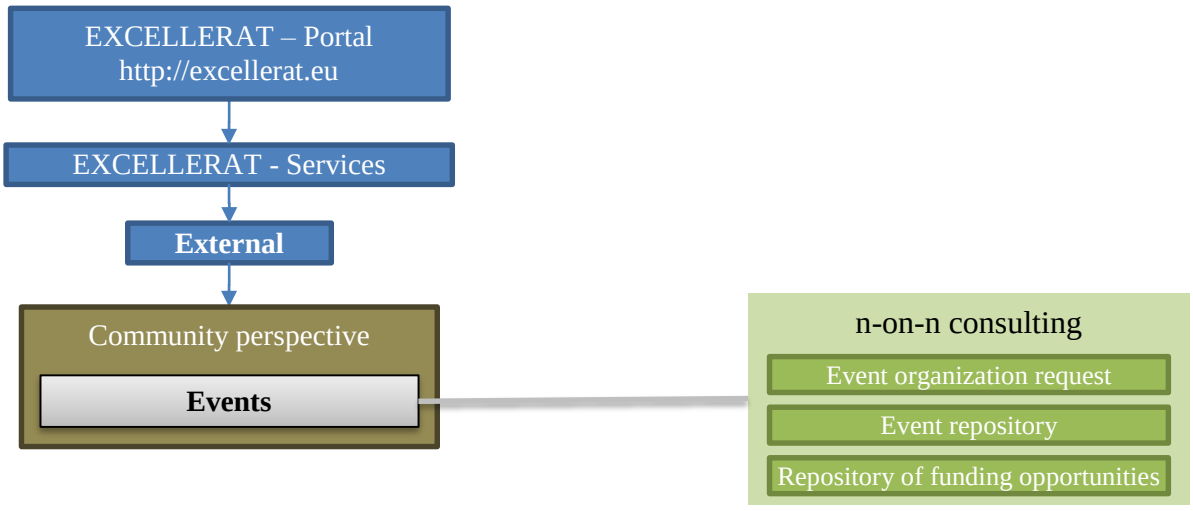

Figure 7: Currently envisaged list of services within the community category

Copyright © 2020 Members of the EXCELLERAT Consortium

4 Conclusion

The Evolution of EXCELLERAT's service portfolio is driven by two main objectives:

(1) To derive a service portfolio that covers to the largest possible extend all requirements and especially all user-groups targeting exascale engineering applications as represented by EXCELLERAT's use-cases.

(2) To identify the services with the largest impact on the engineering community by proper analysis, categorization, market assessment and community building activities.

This will help to further on refine the prioritisation for the implementation of the services by which the centre's sustainability can be achieved after its funding period.

In this deliverable the methodology applied to extend the originally envisaged service catalogue by service requests posed from different perspectives was described in section 2 along with the transfer towards EXCELLERAT's initial service layout which will be implemented via the EXCELLERAT portal. In summary nine services were identified in six service categories which will be accessible via four different service perspectives. Moreover, the relation between services, service categories and service perspectives have been carefully analysed (Figure 3).

Copyright © 2020 Members of the EXCELLERAT Consortium

5 References

[1] The EXCELLERAT consortium, “EXCELLERAT grant agreement,” EC, 2018.

[2] NASA, “NASA Common Research Model,” 10 9 2019. [Online]. Available: https://commonresearchmodel.larc.nasa.gov/. [Accessed 29 11 2019].

[3] Y. Yokokawa, M. Murayama, T. Ito and K. Yamamoto, “Aerodynamics of High-Lift Configuration Civil Aircraft Model in JAXA,” *Journal of the Japan Society for Aeoronautical and Space Sciences,* vol. 55647, pp. 563-571, 2007.

[4] T. Ménarda, S. Tanguyb and A. Berlemonta, “Coupling level set/VOF/ghost fluid methods: Validation and application to 3D simulation of the primary break-up of a liquid jet,” *International Journal of Multiphase Flow,* vol. 33, no. 5, pp. 510-524, 5 2007.

[5] K. P. Geigle, R. Hadef and W. Meier, “Soot Formation and Flame Characterization of an Aero-Engine Model Combustor Burning Ethylene at Elevated Pressure,” *J. Eng. Gas Turbines Power,* vol. 136, no. 2, 2 2014.

[6] U. A. A. Group, “UIUC Airfoil Coordinates Database,” 2019. [Online]. Available: https://m-selig.ae.illinois.edu/ads/coord_database.html. [Accessed 29 11 2019].

[7] G. I. Taylor and A. E. Green, “Mechanism of the Production of Small Eddies from Large Ones,” *Proceedings of the Royal Society A,* no. 158, p. 499–521, 1937.