Helix Nebula – The Science Cloud

Deliverable Title: D6.3 Integrating commercial cloud services into the European Open Science Cloud

Partner Responsible: EGI.eu

Work Package: WP6

Submission Due Date: M36

Actual Submission Date: M36

Distribution: Public

Nature: Report

Abstract: The HNSciCloud hybrid cloud platform links together commercial cloud service providers and research organisations' in-house IT resources via the GEANT network, and provides innovative solutions for supporting the connection of the research infrastructures identified in the ESFRI Roadmap to the European Open Science Cloud (EOSC). These results provide an innovative vision of how to develop the capacity necessary to support the nascent EOSC intended to create a single digital research space for Europe's 1.8 million researchers. The document outlines the EOSC rules of participation and technical requirements for onboarding public and commercial services and provides a list of recommendations to encourage the engagement of commercial cloud providers in the EOSC.

Document Information Summary	
Deliverable number:	D6.3
Deliverable title:	Integrating commercial cloud services into the European Open Science Cloud
Editor:	Giuseppe La Rocca (EGI.eu)
Contributing Authors:	Bob Jones (CERN), Joao Fernandes (CERN)
Reviewer(s):	
Work package no.:	WP6
Work package title:	Pilots Evaluation and Recommendation
Work package leader:	EGI.eu
Work package participants:	CERN, CNRS, EMBL/EBI, ESRF, DESY, INFN, KIT, IFAE, SURFSara
Distribution:	Public
Nature:	Report
Version/Revision:	1.2
Draft/Final:	Final
Keywords:	HNSciCloud, EOSC

Disclaimer

Helix Nebula – The Science Cloud (HNSciCloud) with Grant Agreement number 687614 is a Pre-Commercial Procurement Action funded by the EU Framework Programme for Research and Innovation Horizon 2020.

This document contains information on the HNSciCloud core activities, findings, and outcomes, and it may also contain contributions from distinguished experts who contribute to HNSciCloud. Any reference to content in this document should clearly indicate the authors, source, organisation, and publication date. This document has been produced with co-funding from the European Commission. The content of this publication is the sole responsibility of the HNSciCloud consortium and cannot be considered to reflect the views of the European Commission.

Grant Agreement Number: 687614
Start Date: 01 January 2016
Duration: 36 Months

Log Table

Issue	Date	Description	Author/Partner
V0.1	06/12/2018	First draft version available for partners' input and specific contributions	Giuseppe La Rocca/EGI.eu
V0.2	13/12/2018	First internal review	Bob Jones/CERN
V0.3	21/12/2018	Second internal review	Giuseppe La Rocca/EGI.eu
V0.4	03/01/2019	Third internal review	Giuseppe La Rocca/EGI.eu
V0.5	07/01/2019	Additional inputs and internal review	Bob Jones/CERN
V0.6	17/01/2019	Incorporated feedback from RHEA Group and T-Systems	Giuseppe La Rocca/EGI.eu Alastair Pidgeon/RHEA Group Bob Jones/CERN Martin Canning/IDC
V1.0	29/01/2019	Final Version for review	
V1.2	27/02/2019	Published version follow final project review	Bob Jones/CERN

Executive Summary

The objective of the European Open Science Cloud (EOSC) is to give the European Union a global lead in research data management and ensure that European scientists reap the full benefits of data-driven science. The European Commission (EC) has recently approved the governance and the financial mechanisms to support the Implementation Roadmap for the EOSC.

A key aspect for the implementation of this vision is the definition, through the establishment of a compliance framework, of a set of rules for supporting the federation of European research data infrastructures wishing to join, and enable researchers to access services provided by commercial providers that comply with EOSC requirements. A preliminary set of rules of participation and the technical requirements needed for onboarding public and commercial services in the official EOSC services catalogue have been defined.

The initial EOSC eco-system will be composed of services delivered by publicly-funded providers across Europe and by existing research data repositories.

The HNSciCloud project has developed and piloted a hybrid cloud platform for data intensive science to support the requirements of 10 leading European research centres, and have made the services available to public sector institutions and private sector businesses to support a range of diverse use cases.
The HNSciCloud project has been highlighted by the EC High Level Expert Group as a concrete example of EOSC in practice, providing an innovative vision of how to develop functionality and capacity necessary to support the nascent EOSC digital research space for Europe's 1.7 million researchers.

The onboarding process for commercial cloud service providers described in this document builds on the technical architecture as documented by the EOSC-hub project.

The overall governance structure of EOSC consists of the following three-layer model:

- Strategic;
- Executive;
- Stakeholder.

According to this EOSC governance model, the only area where commercial cloud service providers can contribute directly to the set-up of EOSC is via the Stakeholder Forum working groups.

Today, commercial cloud services do not play a significant role in the production computing environment for the publicly funded research sector in Europe, even if they can:

- stimulate the development of new cloud-based solutions to target the requirements coming from cutting-edge researchers,
- offer users more choice of services,
- encourage cost-effective service delivery,
- offer certified services that can facilitate compliance with relevant standards and legislation notably in the domains of security, and data protection.

Building on the lessons learned during the HNSciCloud project, a set of recommendations are proposed to permit commercial cloud service providers to contribute to the nascent EOSC. Overall, commercial services need to be on an equal footing to the institutional services and be a central part of the EOSC vision. Currently, they appear as a bolt-on or after thought. Unless there is a significant commercial market prospective, commercial providers will be reluctant to engage and even less so to invest.

A lightweight process for onboarding commercial cloud providers in EOSC is needed to accelerate their engagement. The recommendations address the networking layer and federated identity management services as well as contractual aspects to ensure continued access to data. Beyond the technical aspects, it is important that the public sector procurement processes in the context of EOSC for European and national organisations are revised to simplify the provisioning of services and implement a value-for-money approach when comparing services with traditional procurements of hardware and software. The simplified procurement processes should consider the total cost of ownership that can compare reliably in-house services with services provisioned externally.

Table of Contents

1. Introduction

At the beginning of this year the European Commission (EC) has approved the governance and the financial mechanisms to support the Implementation Roadmap for the European Open Science Cloud (EOSC)[1]. The document gives an overview of six pillars for the implementation of the EOSC: architecture, data, services, access and interfaces, rules and governance, and describes the measures taken under Horizon 2020 Work Programmes to start implementing the European Open Science Cloud. In a nutshell, the objective of the EOSC is to give the European Union a global lead in research data management and ensure that European scientists reap the full benefits of data-driven science, by offering "1.7 million European researchers and 70 million professionals in science and technology a virtual environment with free at the point of use, open and seamless services for storage, management, analysis and re-use of research data, across borders and scientific disciplines". The document also states that: "the EOSC will be a fundamental enabler of open science and of the digital transformation of science, offering every European researcher the possibility to access and reuse all publically funded research data in Europe, across disciplines and borders", and [..] "widening the EOSC by gradually opening its user base to the public sector and industry".

Taking into account the experience and the lessons learned over the lifetime of the HNSciCloud Pre-Commercial Procurement project, this document highlights how the pilot services developed by the commercial cloud service providers can contribute to the nascent EOSC, and makes a set of recommendations that should be taken into account during the implementation.

The document is organized as follows:

- Section 2 presents the overall European Open Science Cloud (EOSC) vision.
- Sections 3 and 4 describe the first implementation of EOSC.
- Section 5 introduces the EOSC-hub Service Management System (SMS).
- Section 6 focuses on the establishing and operating EOSC.
- Section 7 highlights the overall structure of the EOSC governance.
- Sections 8 and 9 outline the rules of participation and the technical requirements needed for public and commercial services to be part of the official catalogue.
- Section 10 introduce the role of commercial cloud service providers in EOSC and how the HNSciCloud solutions can contribute to the nascent EOSC.
- Section 11 draws some conclusions.
- Appendix I – The HNSciCloud results.
- Appendix II – Service Maturity.

[1] https://ec.europa.eu/research/openscience/pdf/swd_2018_83_f1_staff_working_paper_en.pdf

In the first part, the EU vision to implement the European Open Science Cloud (EOSC) to support and develop open science and innovation in Europe and beyond is introduced. This includes also an overview of the EOSC-hub H2020 project which aims to create the integration and service management structure of EOSC.

The second part of the document focuses more on the HNSciCloud project initiative, and how its solutions and experience can contribute to the implementation of a full-scale European Open Science Cloud. HNSciCloud has developed two hybrid cloud platforms for data intensive science to support the requirements of 10 leading European research centres with the intention to make the results available to public sector institutions and then private sector business in the longer term. Thanks to these outstanding results the HNSciCloud project has been highlighted by the EC High Level Expert Group as a concrete example of EOSC in practice, providing an innovative vision of how to develop capacity necessary to support the nascent EOSC digital research space for Europe's 1.7 million researchers.

The last part of the document reports the Rules of Participation (RoP) and the technical requirements needed for onboarding public and commercial services in the official EOSC services catalogue. These minimum set of requirements for commercial or non-commercial entities that wish to participate in the EOSC as service providers have been preliminary discussed in the "Principle of Engagement"[2] draft document elaborated by the EOSCpilot[3] project. These requirements included compute, storage and network capacity, accessibility, interfaces, identifiers and metadata, authentication and authorization, information assurance and a willingness to accept vouchers to access the proposed services.

[2] https://docs.google.com/document/d/1JBd30DMG_KhG8vIh1IrOIM4VpYFjOnk4sS59whjaIEI/edit
[3] https://eoscpilot.eu/

2. The European Open Science Cloud vision

European Open Science Cloud (EOSC) is a vision for a federated, globally accessible, multidisciplinary environment where researchers, innovators, companies and citizens can publish, find, use and reuse each other's data, tools, publications and other outputs for research, innovation and educational purposes. Making this vision a reality is essential to remove technical, policy and human barriers and facilitate the creation of new knowledge and economic prosperity in Europe. In terms of architecture, EOSC is expected to grow as an "eco-system" environment where different research data infrastructures, interconnected to operate as a seamless European research data infrastructure, are committed to provide services for helping research communities to use, provide and share scientific data and results.

A key aspect for the implementation of this vision is the definition, through the establishment of a compliance framework, of a set of rules for supporting the federation of other research data infrastructures wishing to join, and enable researchers to access services provided by commercial providers that comply with a defined set of requirements.

Broadly speaking, this compliance framework should outline:

- the rules for participation,
- how to meet the FAIR data principles,
- how services can be integrated,
- address security and privacy issues, and
- include governance practices and risks controls.

3. Implementation of EOSC

EOSC aims to support the following three main objectives:

1. increase the value of scientific data assets by making them easily available to a greater number of researchers, across disciplines and borders,
2. reduce the costs of scientific data management, and
3. ensure adequate protection of information/personal data according to applicable EU rules.

The initial EOSC eco-system will be composed of services delivered by publicly-funded providers across Europe such as: EGI, EUDAT and GÉANT, and by existing research data repositories. Work to integrate and federate such services has already began in Horizon

2020 Work Programme 2016-2017, with the EOSC-hub project and other related projects expected to deliver services under the EOSC. These projects will deliver the initial catalogue of services and data to be provided by EOSC and will define the delivery model(s) for the services. Those catalogues would be enriched periodically through the process of federation. From a technical point of view, the access to these services will be provided through the EOSC Portal[4] as a delivery channel connecting the demand-side and the supply-side.

Further support will be provided by H2020 projects funded via the following calls:

- INFRAEOSC-01-2018 - Access to commercial services through the EOSC hub.
- INFRAEOSC-02-2018 - Prototyping new innovative services.
- INFRAEOSC-03-2018 - Integration and consolidation of pan-European access mechanisms to public e-infrastructures and commercial services through the EOSC hub.

4. The EOSC-hub project

By mobilizing e-Infrastructures comprising more than 300 data centres worldwide and 18 pan-European infrastructures, this project is a ground-breaking milestone for the implementation of the European Open Science Cloud. Overall, the EOSC-hub project[5] creates the integration and management system (the hub) of the future European Open Science Cloud that delivers a catalogue of services, software and data from the EGI Federation[6], EUDAT CDI[7], INDIGO-DataCloud[8] and major research e-Infrastructures. Services and data will be provided ensuring confidentiality, integrity and accessibility in compliance with EU regulations by defining and enforcing community-defined policies and controls. Where personal data is concerned, controls will be introduced to comply with the European General Data Protection Regulation (GDPR) and its adaptations in the national regulations.
The hub builds on mature processes, policies and tools from the leading European federated e-Infrastructures to cover the whole life-cycle of services, from planning to delivery. The hub aggregates services from local, regional and national e-Infrastructures in Europe and worldwide, and acts as a single contact point for researchers and innovators to discover, access, use and reuse a broad spectrum of resources for advanced data-driven research. The services in the hub, which will be offered at highest Technology Readiness Levels (TRLs), address the need for interoperability by promoting the adoption of open standards and

[4] https://eosc-portal.eu/
[5] https://eosc-hub.eu/
[6] https://egi.eu/
[7] https://www.eudat.eu/eudat-cdi/
[8] https://www.indigo-datacloud.eu/

protocols (more details in Appendix II). Through the virtual access mechanism, more scientific communities and users have access to services supporting their scientific discovery and collaboration across disciplinary and geographical boundaries.

5. The EOSC-hub Service Management System (SMS)

The EOSC-hub project defines, creates and operates the integration and management system of EOSC. Through this Service Management System, the project delivers a catalogue of services, software and data from various internal and external service providers, including the EGI Federation, the EUDAT CDI, INDIGO-DataCloud and major research e-Infrastructures. From a technical perspective, EOSC-hub addresses service integration and management by:

1. allowing the end-to-end composition of services,
2. aligning scope, value, service catalogue entries and their specifications across EOSC providers,
3. managing relationship and collaboration between the providers, who will be represented in the project through a dedicated board – the Operations Advisory Board, and
4. defining standardization and modularization. EOSC-hub will manage the EOSC suppliers to provide flexible, innovative, standard and consistent services, and will be accountable for the integrated services that are delivered. In addition, the project will act as the central point of control between demand and supply of advanced services for data-driven research.

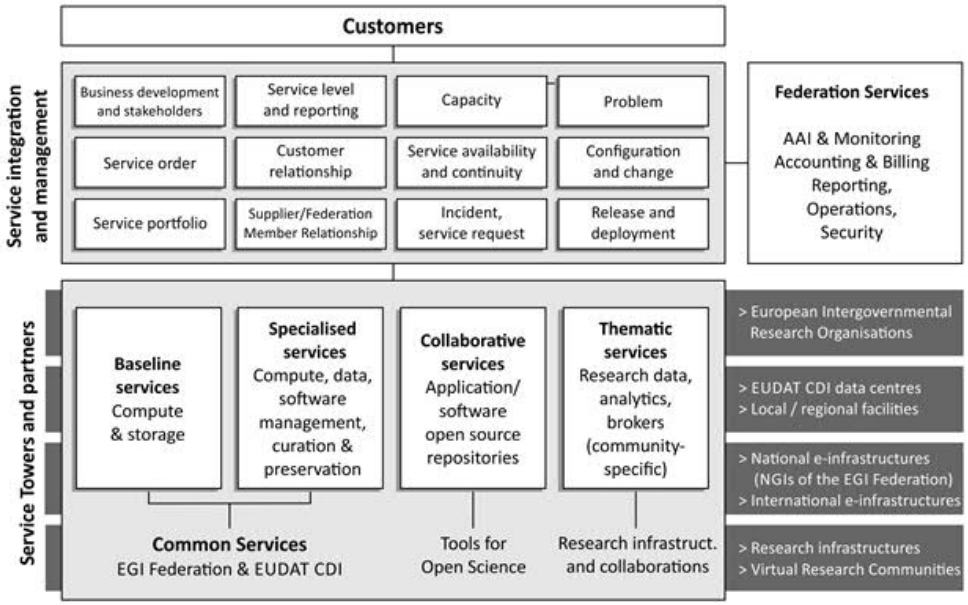

Figure 1. EOSC-hub provides service integration and management services and activities that ensure the delivery of services in a multi-supply environment.

EOSC-hub will define and implement the EOSC IT service management system (ITSM), i.e. the entirety of activities performed by service providers to plan, deliver, operate and control services offered to customers. These activities will be directed by policies and will be structured and organized by processes and procedures.
The EOSC-hub ITSM ensures all participants have a clear understanding of all concepts, terminology and activities to be carried out. EOSC-hub will implement best practices based on the lightweight standard, FitSM[9], regarding the service planning, delivery, operation and control of the services in the service catalogue.

6. Costs for establishing and maintaining the EOSC

As stated before, even if the initial implementation of the EOSC federation will be based by services delivered by providers across Europe such as: EGI, EUDAT CDI, GÉANT[10], and by existing research data repositories, it will require additional investments. The seven main cost categories to establish and maintain EOSC are described below.

- Employing cloud-computing services

[9] www.ftism.eu
[10] https://www.geant.org/

The cost of getting data into the cloud and storing some of it for decades, and the cost of using cloud computing resources to access and analyse scientific data, including the necessary connectivity.
Although data within the scope of the EOSC is intended to be free for registered users to access at the point of use, researchers will need to pay to use the cloud computing resources required to analyse and manipulate the data. As this marketplace will be demand-driven, the June 2017 Declaration envisions that both public research data Infrastructures and commercial operators will develop and provide services based on user needs, and discontinue provision when not justified by the level of adoption.

- Opening up scientific data

The implementation of data management plans to make research data findable, accessible, interoperable and re-usable (FAIR principles). Academics, industry and public services will have to be persuaded and incentivised to share their data, and improve their data management training, literacy and stewardship skills. Research projects covered by the H2020 programme will have to implement data management plans to make research data findable, accessible, interoperable and re-usable (FAIR principles).

- Federation of existing scientific data infrastructures

This federation will be the critical building block for EOSC. The vision is that national and disciplinary nodes will be connected at a pan-European level creating federated data infrastructures for specific thematic areas. The movement of existing data is not cost-free.

- Development of specifications for application interoperability

To enable data to be shared across disciplines and infrastructures, more standardization of meta-data and, perhaps, the actual data itself will be needed.

- Creation of search tools

The development of new software tools to enable scientists to search, browse and access research data will be required. These new tools will make use of the meta-data that annotates and identifies the underlying data. This meta-data will need to conform with specifications that enable it to be processed through common, open source data analysis tools. Ideally, all research data will be available programmatically, through web APIs, so that it can be identified and accessed by search engines and automated systems.

- Creation and maintenance of a security environment

The European Commission envisions a suitable certification scheme to guarantee security, data portability, and interoperability in compliance with the new General Data Protection Regulation (GDPR). Such a scheme will need to be flexible enough to enable the EOSC to keep pace with the evolution of scientific research.

To manage access to research data and tools, the EOSC will need a robust authentication system, which combines a single sign-on process, resulting in a federated identity and credentials for all users of the EOSC. Such a system should build on the existing systems employed by the research infrastructures to authenticate members of their respective user communities.

7. The governance of the EOSC process

According to the final report and the recommendations of the 2nd High Level Expert Group on the European Open Science Cloud, the overall governance structure of EOSC consists of a three-layer model as depicted in Fig. 2. This governance structure will be funded by the European Commission for a 2 year period, following which a revised governance model will be put in place taking into account the experience gathered and with the goal of being more sustainable.

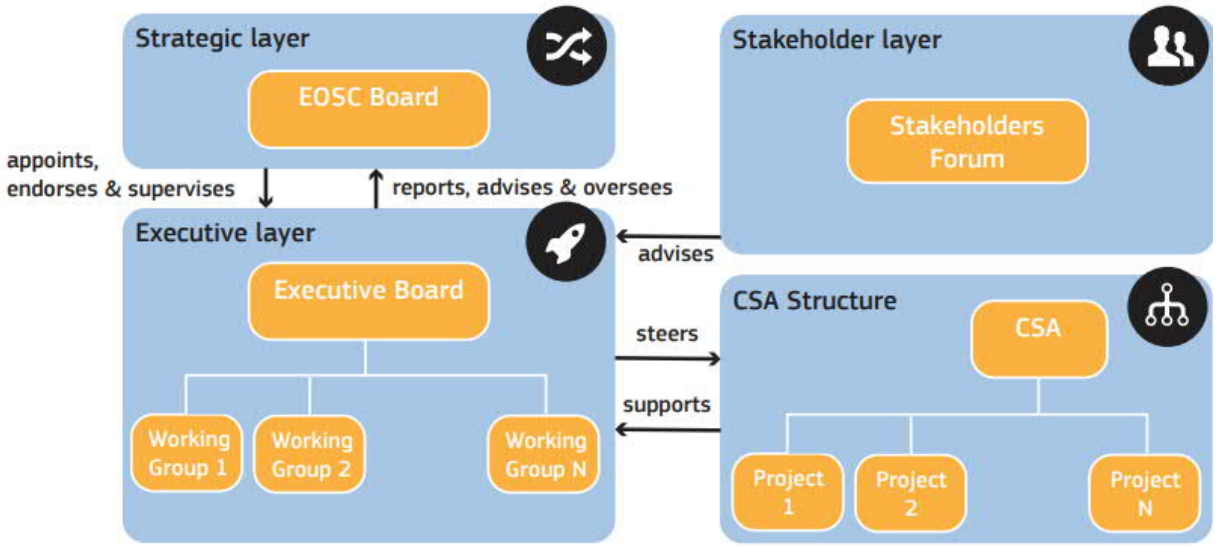

Figure 2. The EOSC three-layer governance model

To summarise, the three layers are:

- Strategic Layer in the form of an EOSC board to combine state-of-the-art expertise on scientific cloud infrastructures with the Funders and Policy Makers. The EOSC Board will therefore include EU Member States and Associated Countries representatives. The EOSC Board will involve Member State, Associated Countries and Commission representatives. The EOSC board is responsible to define the strategic aim of the EOSC, review, agree and prioritize the EOSC proposals and requirements from the strategic vision point of view. The EOSC Board has to ensure

the coordination of Member States and Commission initiatives, as well as monitor and assess the achievements of EOSC, approve the list of the executive board members and executive board work plan.

- Executive Layer in the form of an Executive Board to manage day-to-day operation of the EOSC and procurers designing and planning work-related future developments. The list of individual experts selected as members of Executive Board has been announced during the official launch of EOSC in November 2018[11].

- Stakeholder layer organized in the form of a Stakeholders Forum to provide a medium for stakeholders: Users (Consumers), Providers and Intermediaries of EOSC Resources. This would have the main role to discuss, supervise and channel communication between the EOSC and the communities across all three layers. The Stakeholders forum must grant all stakeholder groups the possibility to determine the requirements, policies and principles of participation. The forum must also have an organizational structure that enables consumers, providers and intermediaries of all sizes to participate and interact with the boards.

The Executive Layer is complemented by a number of different Working Groups (WGs), created with a time limited duration (e.g. up to 2-year duration), working on specific areas identified by the Stakeholder forum.

A Coordination Support Action (CSA) Structure will help the Executive Board coordinate all relevant European Commission-funded projects and support the implementation of the work plan while also acting as a secretariat to deliver the EOSC main functionalities.

According to the EOSC governance model described in Fig. 2, the only area where commercial cloud service providers can contribute directly to the set-up of EOSC is via the Stakeholder Forum.

8. The Rules of Participation

The EOSC Rules of Participation (RoP) set out, in a transparent and inclusive manner, the rights, obligations and accountability of the different stakeholders taking part in the initiative (e.g. data producers, service providers, data and service users). The initial set of

[11] https://ec.europa.eu/info/news/results-call-applications-selection-members-expert-group-members-executive-board-eosc-2018-nov-23_en

requirements for the service and data providers have been preliminary discussed in the "Principle of Engagement"[12] draft document elaborated by the EOSCpilot[13] project and shall be further specified in the upcoming months by other EOSC initiatives. These requirements included compute, storage and network capacity, accessibility, interfaces, identifiers and metadata, authentication and authorization, information assurance and a willingness to accept vouchers which can be used to purchase services from the cloud service providers. The resulting requirements, elaborated in the context of the EOSC-hub project, are described below.

- The Rules of Participation for data and service providers

The rules of participation for data and service providers include principles that drive the participation of service providers and users in EOSC. The main rule of participation is that EOSC services shall be registered in an EOSC compliant or compatible service catalogue visible to the global EOSC gateway. In the EOSC compatible catalogues, the registered databases and services must be described according to the appropriate EOSC guidelines, which could include information such as the service availability, functionalities, operations, maturity, user support, interoperability (metadata schemata supported), openness (licenses), privacy (GDPR compliance), terms of use and contractual framework.

The following basic requirements for service and data providers that aim to become EOSC providers, have been identified by the EOSC-hub project:

- Enable Machine-Readable Metadata, and include Persistent Identifiers along with a machine-readable description.
- Have detailed Terms of Use and Policies.
- Ensure Accessibility and Interoperability using standards and APIs.
- Allow the portability of data and services.
- Provide information on any costs associated to use of the service and the charging model.
- Declare the minimum quality guidelines and fulfil them.
- State how data would be stored and processed, for sake of transparency.

- The Rules of Participation for users

Users' Rules of Participation may be provided within a framework such as the Terms of Use presented by the Service Providers, including at least the two following principles:

Data sharing and sustainability

[12] https://docs.google.com/document/d/1JBd30DMG_KhG8vIh1IrOIM4VpYFjOnk4sS59whjaIEI/edit
[13] https://eoscpilot.eu/

- In the domains where community-recognised data deposition repositories exist, users should be encouraged to deposit data in them as a first instance.
- Users are recommended to publish their data in a FAIR format: and as open as possible and as closed as necessary.

Acknowledgement of use of EOSC services

- If requested by the Terms of Use of the data and service provider, users should acknowledge through citation in their publications, or by other means, the specific database, service/services accessed through the EOSC that have enabled their research.

9. Procedure for onboarding new services in EOSC

This section outlines the operational requirements that will be applied for existing and potential future service providers wishing to participate in the European Open Science Cloud (EOSC). This work has been carried out as part of the Rules of Participation task force within the EOSC-hub project, which aims to define the conditions for service providers to offer services through the EOSC-hub.
From a technical perspective, service onboarding within EOSC-hub is the process whereby a new service joins the EOSC-hub service catalogue and EOSC Marketplace. This provides services with all the benefits offered by the catalogue and marketplace - promotion of the service to users outside their local community domain with a single gateway for users to discover and use services, regardless of their nature and the scientific discipline of the user, and potential integration with other services in the catalogue.

Within EOSC-hub, two service catalogues are available:

1) The internal catalogue, containing access enabling services developed as part of the project and necessary for the operation of the EOSC-hub (e.g. helpdesk and AAI).

2) The external catalogue, containing the common services, on which many services depend (data, compute, orchestrators), and research-enabling services offering services to the end user, typically building on the common services. Services in these catalogues may attain varying levels of service integration:

- Low level has the minimum set of SMS requirements. Nevertheless, along with all other services in the external catalogue, services entering at this level benefit from the advantages of the EOSC-hub as listed above. Services which enter the catalogue at

this level may either have a less mature SMS which they plan to develop, or a mature SMS but would like to join the EOSC-hub initially without committing additional resources for integration until a later stage.

- Medium level is aimed at services in the external catalogue that are being delivered as part of an existing and mature SMS complying with the majority of requirements of FitSM or other recognized Service Management Framework. Services achieving a Medium level of integration are encouraged to participate closely with the EOSC-hub by being represented at the Service Management Board meetings.

- High level is for services in the internal catalogue delivered as part of the EOSC-hub SMS. Services from the external catalogue may also achieve this highest level of integration. Services with a High level of integration are expected to participate closely with the EOSC-hub and be represented at the Operations Management Board meetings.

It is anticipated that new services will join the external catalogue at either the Low or Medium level of integration. The benefits for joining as a service provider in EOSC, and the technical requirements requested have been documented[14]. The onboarding process for each case is described in Fig. 3.

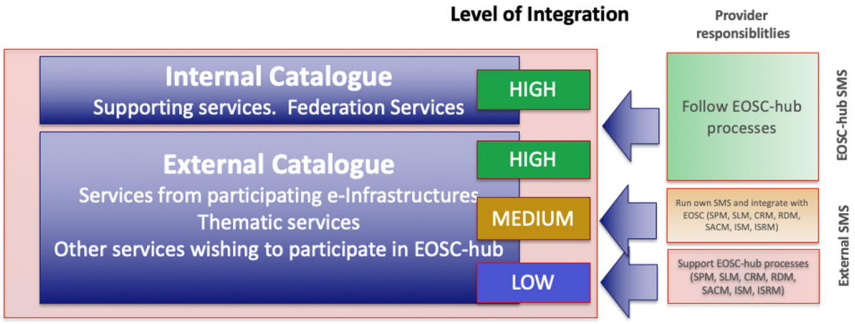

Figure 3. The levels of integration for EOSC services

- The onboarding process of new service providers

The simplified workflow to on-board new service providers in EOSC is depicted in Fig. 4. The different steps of this workflow are described in the next sections.

[14] https://www.eosc-portal.eu/for-providers

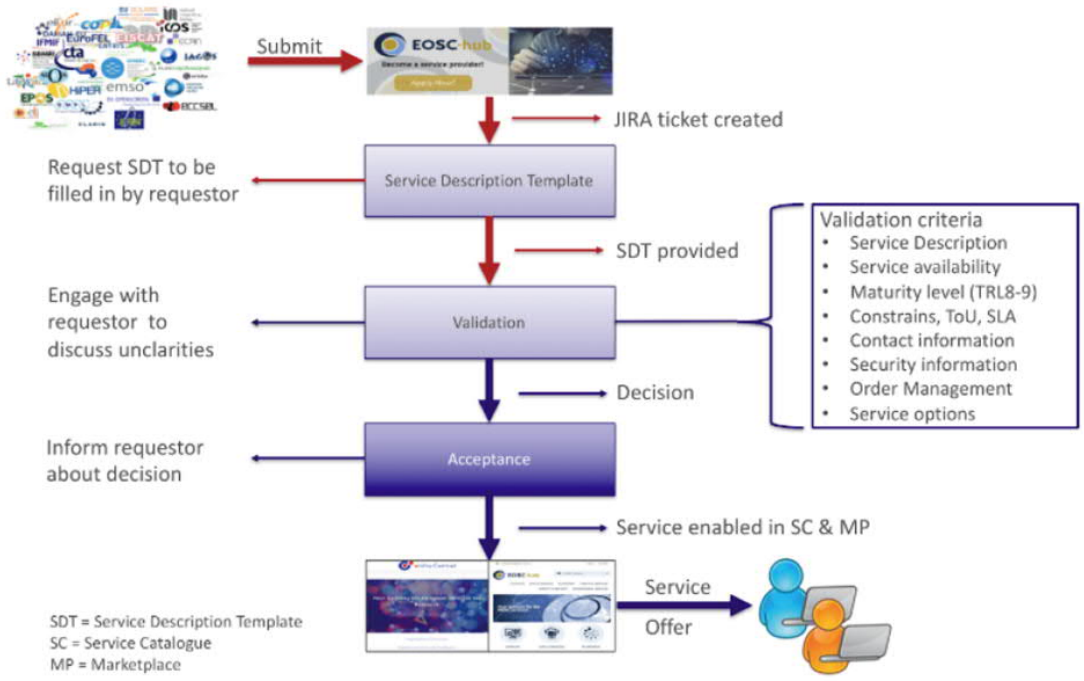

Figure 4 – Simplified workflow for onboarding new services in EOSC

- Initial contact with EOSC-hub

The service provider completes the Join as a Provider form[15] on the EOSC-hub website, providing high-level information about the service and motivation for joining. The submitted information becomes a ticket in Jira to track the request.

- Initial evaluation of request

EOSC-hub staff working with Service Portfolio Management (SPM) evaluate the request, contacting the submitter for supplementary information to initially gauge the validity of the request and to confirm that the service is appropriate for inclusion into the service catalogue. If so, the prospective service provider fills in the Service Portfolio Entry Template[16] (at least all mandatory fields) and service options are discussed.

- Determination of the initial level of integration and verification of requirements

EOSC-hub staff working with Operations review the operations validation checklist to evaluate the maturity of the service (i.e. TRL level of at least 8. For more details, please refer to Appendix II) and the maturity of the service management system of the service provider (this may be done as an online meeting with the service provider, and the checklist may be

[15] https://www.eosc-hub.eu/join-as-service-provider

[16] https://docs.google.com/document/d/1gB30A_1UBqXEGjRpqNEI7_DnZVM8ggUb_TdqPXK0Dt0/edit

sent in advance to aid preparation for the meeting). The initial level of integration is agreed with the service provider.

- EOSC-hub SMS process integration

At this point where all requirements are fulfilled at the agreed level of integration, the following steps are required. In all cases, the step is tracked in Jira.

- Adding the new service to the catalogue

Information required by SPM about the service is added to the catalogue.

- Preparation for SLA creation

The SLA template is populated for the new service in readiness for SLAs with end users, and the service contact is made familiar with its format.
More details on how to formalize the preparation of the agreement between the end users and the service providers will be defined in the Rules of Participation.

- Security

The nominated contact from the service provider (or another nominated security contact) is added to the appropriate database (e.g. GOCDB) for security-related issues.

- Support

For services at a Low and Medium integration level, the appropriate contact for providing support is added to the EOSC-hub catalogue. For services at a High integration level, a new support unit is created in XGUS and associated with the appropriate contact for providing support.

- Marketplace integration

The prospective service provider completes the Marketplace Service Description Template[17] providing information relevant to the Marketplace.

- Operational integration

A nominated contact from the service provider is invited to engage with the Service Provider Forum by joining its mailing list and attending quarterly meetings. The Service Provider Forum is designed to facilitate communication between service providers and the EOSC-hub project, in addition to requirements gathering. If the level of integration is Medium, the nominated contact from the service provider is additionally invited to attend Service Management Board meetings.

[17] https://docs.google.com/document/d/1OWsWBChJ9fh8EIK8pIzIwIXByT1NE_tnmA3X0gxWuks/edit

The EOSC-hub Service Management Board (SMS) is a lightweight management/governance body for all Service Providers of EOSC-hub for the purpose of:

- Reporting of major operational news regarding services in the EOSC-hub Marketplace: security, upgrades, operational problems etc. involving access-enabling services.
- Discussion of potential or actual problems involving operational aspects of services in the EOSC-hub Marketplace that are not being addressed with existing processes and procedures. Agreeing on solutions to such problems, or escalating them to the EOSC-hub Activity Management Board (AMB) if appropriate.
- Definition and approval of non-strategic changes to policies affecting Operations, including operational aspects of service onboarding and Rules of Participation.

10. The role of commercial cloud service providers in EOSC

The European Commission's vision is that the EOSC will ensure every research centre, every research project and every researcher in Europe has access to the computing, data storage and analysis capacity they need to conduct the research they want to do. By putting cutting-edge computing resources at the fingertips of researchers, EOSC could catalyse broader adoption of hyper-scale cloud computing by European science, enable scientific breakthroughs, and generate financial savings that could be used to pay for the costs outlined earlier, such as compliance with the FAIR principles and the creation of a secure environment.

There are significant benefits the EOSC could gain if commercial clouds are introduced as part of the provisioning strategy. A major benefit of commercial clouds is the ability for public research organizations to access services on a variety of payment models, such as a pay-as-you go, thereby potentially reducing the overall costs and enabling applications to scale rapidly. Additional models such as: reserved instances, opportunistic, paying upfront, etc. have been identified in HNSciCloud as having the potential to reduce costs for specific use cases.

For individual researchers, this translates into less time waiting in a queue for the necessary IT resources to become available, leading to potentially substantial productivity gains. With commercial cloud services, complex workflows and complete environments can be set-up quickly and shared widely, enabling rapid prototyping and development of new tools.

Today, commercial cloud services do not play a significant role in the production computing environment for the publicly funded research sector in Europe, even if they can:

- stimulate the development of new cloud-based solutions to target the requirements coming from cutting-edge researchers,
- offer users more choice of services,
- encourage cost-effective service delivery,
- offer certified services that can facilitate compliance with relevant standards and legislation (e.g. in the domains of security, data protection).

Concerning service compliance, commercial cloud providers offer services that are compliant with international security standards, such as ISO/IEC 27000 series, to establish controls for implementing an information security management systems and minimize the consequences of information security incidents. In particular, ISO/IEC 27018:2014 specifies guidelines for the protection of Personally Identifiable Information (PII) which might be applicable within the context of the information security risk environments of a provider of public cloud services. ISO/IEC 27018:2014 is applicable to all types and sizes of organizations, including public and private companies, government entities, and not-for-profit organizations, which provide information processing services as PII processors via cloud computing under contract to other organizations.
For what concerns the data privacy, the General Data Protection Regulation (GDPR) came into effect during the HNSciCloud project execution, replacing the Data Protection Directive 95/46/EC and imposing new obligations on organizations that process the personal data of European Union residents. According to this new regulation, research organizations that process personal data for research purposes must have a lawful basis for any processing activity. In order to process personal data research organizations must have the data subject's consent and this consent must be "unambiguous" and specific to the processing operation. This poses new challenges for research organizations because it is often not possible to fully identify the purpose of personal data processing for scientific research purposes at the time of collection. In practice, this requires that research organizations have to invest in modifying their in-house IT services so they become GDPR compliant.

While commercial cloud services do not play a significant role in the production computing environment for the publicly funded research sector in Europe, they do play a role in the Copernicus' Data and Information Access Service (DIAS)[18] flagship space programme. The European Space Agency (ESA) has opened up the data generated by the Copernicus earth observation programme, one of the largest data providers in the world. The DIAS flagship programme has been designed to kick-start the development of data access and cloud processing services, which can be used by entrepreneurs, developers and the general public to build Copernicus-based services and applications. The overarching goal is to boost user

[18] https://www.copernicus.eu/en

uptake, stimulate innovation and the creation of new business models based on Earth observation data and information. Through public procurement the EC in collaboration with ESA has established five alternative DIAS service providers to give maximum choice and support user and market diversity. Three of the five service providers choose to make use of existing European commercial cloud services (Orange, OVH and T-Systems) and two are relying on build-for-purpose cloud systems also provided by industry. Some of the services include hybrid cloud functions and are linked to public infrastructure e.g. T-Systems has fully integrated a Long Term Archiving (LTA) service provided by The German Aerospace Centre (DLR), In summary, users will have full and free access to Copernicus data and services, and will, at commercial conditions provided by the DIAS providers, be able to process the data and information with alternative service options.

The hybrid cloud platforms developed and piloted by the HNSciCloud project have been highlighted, by the EC High Level Expert Group, as a concrete example of EOSC in practice since it provides an innovative vision of how commercial cloud service providers and research organizations' in-house IT resources, connected via the GÉANT network, can be used to link the research infrastructures identified in the ESFRI Roadmap to the European Open Science Cloud.

The commercial cloud provider services in these platforms are compliant with the new data protection directives and the commercial solutions developed within the project are already compliant with international standards, including: 1.) the ISO/IEC 19086 standard for Service Level Agreements, 2.) the ISO/IEC 27000 series of standards for information security of the resulting services, and 3.) the IEEE 1016-2009 system design document template for the design of the hybrid cloud architecture.

All the pilot services developed by the commercial contractors in the project are based on open source implementations that do not require licenses in order to be deployed on the in-house IT resources of research organizations connected to the hybrid platform, and can be accessible via eduGAIN[19] and ELIXIR[20] federated identity and access management systems. For more details, please refer to Annex I where the pilot services are explained in more detail.

Overall, the HNSciCloud PCP project has showcased how, by adopting a hybrid cloud model, an organisation can quickly and economically add resources as needed by bursting out of its own IT infrastructure to public cloud processing and storage capacity.

[19] https://edugain.org/

[20] https://www.elixir-europe.org/

Broadly speaking, the adoption of the hybrid cloud model allows the organisations to:

- Rely on the public clouds for dynamic workloads, state-of-the-art functionality e.g. for AI and operations that require vast scalability and on the private clouds for well-predictable and legacy workloads or operations that may be subject to specific (legal) constraints, providing enhanced agility to move applications easily between the in-house and off-site resources taking into account aspects of policy, cost, security and availability.
- Leverage the richness in the diversity of European suppliers and to match it with the expertise available in production e-Infrastructures, demonstrating the technical feasibility of interoperability between these players.

Based on the experiences gained during the three competitive phases of the HNSciCloud PCP project, the following recommendations are proposed for the implementation of EOSC:

- There are technical aspects that need to be taken into account for enabling eduGAIN, the inter-federation service. Even if eduGAIN provides the basic specification to configure and create a federated identity environment, additional agreements and specifications are needed so that service providers can determine which identity attributes are to be considered and matched in order to provide granular and manageable access and authorisation to data and resources. ELIXIR has already formalised such agreements and specifications in the context of the Life Sciences community. An enhanced management of attributes in the home identity systems would reduce complexity. Alternatively, the integration of a proxy service in EOSC would simplify attributes management, but would shift the management burden to EOSC and require additional resources.
- High-performance I/O applications require dedicated network connectivity and data traffic management. From a technical perspective GÉANT offers policies to support the traffic exchange between commercial providers and the NRENs, and provide data traffic management, but the process is not straightforward since it involves different stakeholders. Streamlining the process to connect commercial providers to the GÉANT network is another critical point that EOSC should try to address. Also, since one of the objectives of EOSC is to promote the development of data-intensive research using solutions and resources distributed across Europe, the monitoring of the network performance is another critical aspect that needs to be taken into account. The adoption of a solution, such as the perfSONAR networking monitoring service, to ensure end-to-end performance is desirable.
- For all the public and commercial services available in EOSC, a contractual agreement should permit users to repatriate their datasets and configuration information. From a technical standpoint, this means that the contractual

agreements between the cloud service providers and customers have to include a concrete cloud-service "exit plan", and a "meaningful and effective" redress in case the portability requirements defined in the contract are not met. To overcome the technical, economic and legal challenges linked to data repatriation, EOSC should follow the recommendations outlined by the report "Switching of Cloud Services Providers"[21], that IDC and Arthur's Legal developed for DG Connect and delivered in 2018. According to this report the most promising way to facilitate the portability of data and applications in cloud environments would be the introduction of a right to data and application portability obligation that applies to cloud service providers. The Under the free flow of non-personal data Regulation, the European Commission has encouraged the development on a self-regulatory basis of a Code of Conduct for Cloud Switching/Data Porting to reduce the risk of "vendor lock-in". In anticipation of this regulation, SWIPO has been formed as the collective name of a group of stakeholders that are working towards a code of conduct in the domain of Cloud Computing. A first version of the PaaS code of conduct has been opened for public consultation[22]. These recommendations and best-practices should also be taken into account in EOSC.

- A lightweight process for onboarding commercial cloud providers in EOSC has to be set-up. This process should take into consideration how the commercial support structures can be integrated with the EOSC helpdesk. The current rules of participation and the procedure for onboarding new services providers in ESOC are documented in Sections 8 and 9.
- The cost prediction, usage monitoring and billing, across different providers, are critical points that the HNSciCloud project has partially addressed and should be taken into consideration by the nascent EOSC.
- A voucher provides pre-paid access for users to a cloud providers services for a specified monetary value. Vouchers can allow also researchers to "shop around" and choose the service provider that gives them the best value for money for their needs. The introduction of vouchers offer a convenient means to finance the use of EOSC services and increase their adoption by more users. A preliminary implementation of voucher schemes have been provided by the contractors involved in the pilot phase of the HNSciCloud project.

Additional non-technical points to be considered are the following:

- The uptake of commercial public cloud services in the public sector is hampered by procurement regulations that continue to favour "hardware and software" provisioning based on CAPEX rather than service provisioning based on OPEX.

[21] https://publications.europa.eu/en/publication-detail/-/publication/898aeca7-647e-11e8-ab9c-01aa75ed71a1/language-en

[22] https://docs.google.com/forms/d/e/1FAIpQLSejA2MGW_UCkfjgK-Yxo5q6lzhD1yh1a0LhDrrcPWYEoxVE1A/viewform

- Many public sector research organizations have a limited understanding of the full costs of provisioning services in-house which distorts total costs of ownership comparisons with commercial services.
- The choice of services often depends on a complex decision making process involving intermediaries between the user and the funder that hinders innovation since the user may not get access to newest service options available in the market.
- There is the need to stimulate an attractive market where the demand and the supply sides can meet each other and understand their requirements. Currently, there are no initiatives aiming to develop and stimulate this interaction between the two parties. OCRE, the three-years H2020 project started in January 2019, is the first initiative aiming to bridge this gap.
- More focus needs to be placed on stimulating the market.
- A channel is required by which commercial cloud service providers can provide their input to the EOSC governance structure.

11. Conclusions

The Final report and recommendations of the Commission 2nd High Level Expert Group on the EOSC, entitled 'Prompting an EOSC in practice'[23] concluded with a list of long-term challenges for EOSC that included the following:

- Private sector involvement in the EOSC.
- Further development of procurement models for EOSC delivery.

This document contributes to addressing these challenges taking into account the experience and lessons learned from the HNSciCloud PCP project. It highlights some crucial topics for the implementation of a full-scale European Open Science Cloud, such as the definition of the main Rules of Participation (RoP) and the light-weight process for onboarding new services providers wishing to join the EOSC catalogue. By the time of writing, the RoP and the onboarding process described in this document are in line with what the ESOC-hub project has reported in the D10.3 – EOSC-hub Technical Architecture and standard roadmap v1.

To contribute to the nascent EOSC, the experience and the lessons learned by the project during the three competitive phases have been reported in this document. These aspects should be taken into account during the implementation of EOSC.

A summary of these recommendations for integrating commercial cloud services into the EOSC are reported below:

- The attribute management of eduGAIN providers is desirable or the use of the foreseen integration of a proxy service in EOSC, to simplify attributes management.
- High-performance I/O applications require dedicate network connectivity. The adoption of solution, such as the perfSONAR networking monitoring service, to ensure end-to-end performance is also desirable.
- For all the public and commercial services available in EOSC, a contractual agreement should permit users to repatriate their datasets and configuration information. Recommendations outlined by the report "Switching of Cloud Services Providers" and the best practices being defined by the SWIPO working group to reduce the risk of "vendor lock-in" should be taken into account in EOSC.
- A lightweight process for on-boarding commercial cloud providers in EOSC has to be set-up. This process should also take into consideration how the commercial support structures can be integrated with the EOSC helpdesk.

[23] ISBN 978-92-79-94836-7 doi:10.2777/112658

- The cost prediction, usage monitoring and billing, across different providers, are critical points that have to be take into consideration by EOSC.
- Vouchers offer a convenient means to finance the use of EOSC services and increase their adoption by more users.

Additional non-technical points to be considered are the following:

- The uptake of commercial public cloud services in the public sector is hampered by procurement regulations that continue to favour "hardware and software" provisioning based on CAPEX rather than service provisioning based on OPEX.
- Many public sector research organizations have a limited understanding of the full costs of provisioning services in-house which distorts total costs of ownership comparisons with commercial services.
- The choice of services often depends on a complex decision making process involving intermediaries between the user and the funder that hinders innovation since the user may not get access to newest service options available in the market.
- There is the need to stimulate an attractive market where the demand and the supply sides can meet each other and understand their requirements. Currently, there are no initiatives aiming to develop and stimulate this interaction between the two parties. OCRE, the three-years H2020 project started in January 2019, is the first initiative aiming to bridge this gap.
- More focus needs to be placed on stimulating the market.
- A channel is required by which commercial cloud service providers can provide their input to the EOSC governance structure.

Appendix I – The HNSciCloud platforms

In Annex I we focus on the pilot services developed, by the two contractors involved in the project, that support research organisations to deploy very diverse applications and workflows on any combination of dedicated hardware, private cloud and public cloud infrastructure. These services could be adopted for supporting the implementation of the EOSC vision.

The hybrid cloud service based on the Open Telekom Cloud

Open Telekom Cloud (OTC)[24] is an international large-scale public cloud from Deutsche Telekom, based on OpenStack, supported and operated by Deutsche Telekom's entity T-Systems out of Europe. During the HNSciCloud project T-Systems has developed a hybrid cloud service, that enables science organizations to run very diverse applications and workflows on any combination of dedicated hardware, private cloud and public cloud infrastructure.

The resulting hybrid cloud service includes integrated access to all functions through a single (federated) identity, that can be based on existing Identity Providers in science e.g., eduGAIN and Elixir AAI. The existing IaaS-functionality has been extended to provide High-Performance, GPU and FPGA compute resources in a cloud environment including access to advanced functionality e.g., InfiniBand networking and MPI library functions. Scripts and images enable users to use HPC-as-a-Service, organizing and scheduling jobs with established HPC tools and batch systems. The CPU flavours available in the Open Telekom Cloud platform are reported in Table 1.

OTC flavour	CPU	Cores/ Chips/ Threads	CPU2006 RATE relative performance	vCPU/Core Cost (Euro/h.)	PPR 2006 Ratio per Core[1]
S1/C1/C2/M1[2]	Xeon E5-2658A v3	24/2/2	1.00	0.02100	23.8
S1/C1/C2/M1[2]	Xeon E5-2658A v4	28/2/2	1.03	0.02100	24.6
S2[2]	Xeon 6161	44/2/2	1.08	0.02100	25.8
C3[3]	Xeon 6151	36/2/2	1.37	0.06180	22.2
M2[3]	Xeon E5-2690 v4	28/2/2	1.17	0.07094	16.5
M3[3]	Xeon 6151	36/2/2	1.37	0.07980	17.2
HP I[3]	Xeon E5-2690 v3	24/2/2	1.11	0.07875	14,1
HP II[3]	Xeon E5-2667 v4	16/2/2	1.45	0.12875	11.3

Table 1 – OTC CPU flavours and price-performance ratios

[24] https://cloud.telekom.de/en/infrastructure/open-telekom-cloud/products-services

The GPU flavours available in the Open Telekom Cloud platform are reported in the Table 2 (based on virtual servers) and Table 3 (based on bare-metal servers) below:

Flavor	vCPU	GPU	Memory	NVME	Network	Description
p2.2xlarge.8	8	1*V100	64GB	800G	10GE	1:1,1:2,1:4 V100 are provided by ECS P2 o High Performance (V100 GPU with GPU Direct/P2P) o Local NVME SSD o Agility of ECS
p2.4xlarge.8	16	2*V100	128GB	800G*2		
p2.8xlarge.8	32	4*V100	256GB	800G*4		

Table 2 – Virtualized GPU V100 flavours

Flavor	CPU	GPU	Memory	NVME	Network	Description
physical.p3.large	2*2690 v4	8*V100	512GB	4.8T	IB-EDR	1:8 V100 is provided by BMS o Ultra high performance o Local NVME SSD o IB-EDR

Table 3 – BMS for GPU V100 flavours

Starting from July 2018, 1 FPGA hardware accelerator is also included in the release. The technical requirement of the new hardware is described in Table 4.

INSTANCE	VCPU	MEMORY	FPGA	DISK	NETWORK
fp1.2xlarge.11	8	88G	1	EVS	10Gb

Table 4 – Accelerated FPGA flavour

Cloud storage functions have been significantly enhanced with the state-of-the-art Onedata global namespace data management solution. Onedata[25], the distributed virtual filesystem developed by Cyfronet, offers high-performance and unified data access across globally distributed environments and multiple types of underlying file systems, enabling users to share, collaborate and perform computations on the stored data. Both servers and containers can access the stored data in a seamless way.

The hybrid cloud service is complemented by an Enterprise agreement, that provides users a choice of SLA quality for each cloud service and commitments are backed by service credits. The service is accessible through GÉANT with n*10 Gbps throughput based on full diverse fibre infrastructure and monitored by a perfSONAR system. The network monitoring server with external 40 Gbps capability, includes a daily, very detailed and granular consumption

[25] https://onedata.org

monitoring and enables users and organizations to configure and provide authorized access to their own metrics and dashboards.

Users can be provided with cloud credits through vouchers to get access to the service free at the point of use. In 24 hours, at the latest, the tenant on the Helix Nebula Science Cloud will be configured and the user will start to use cloud computing and storage resources in the amount of the voucher budget. For the project vouchers could be registered at: https://www.divia.de/hnscicloud, and their budgets could be monitored accessing the financial dashboard based on Grafana[26]. Additional, a financial dashboard alert functionality is available to notify the user when his/her voucher budget reaches the 50% and the 80% of the total credit. For future EOSC and other projects, T-Systems has committed to further develop the cloud credit function into a fully integrated part of the commercial cloud service.

From a technical perspective, the hybrid cloud service functionality is at different levels of maturity (see Annex II for more details). Most services are commercially available based on Open Telekom Cloud including e.g. IaaS, Container Management, HPC-, GPU- and FPGA-computing, Networking, Identity and Access Management and Service Quality (TRL-9). Other functions range in TRL levels 6 to 8, e.g. Data Management and extended features required for IAM and Service Payment.

From the end-users perspectives, the hybrid cloud service is simple to use through a web or command line interface, it can be fully automated through APIs, and it provides extensive dashboards for consumption monitoring. Resulting, large-scale on-demand analytics with thousands of containers or large memory machines with up to 1TB RAM can be run within minutes as has been demonstrated e.g. for the CERN CMS experiment and the Marine Metagenomics at EMBL-EBI. In the past 2 years and half of operation the Open Telekom Cloud service has been able to receive significant recognition in the market. In 2018 Open Telekom Cloud has been also selected to run the Copernicus DIAS services – providing access to all Sentinel Satellite data for direct processing – for the European Union under contract with ESA, it has been certified as the first public cloud service in Germany to be compliant to the GDPR.

[26] https://grafana.com/

The RHEA Group hybrid cloud service

The RHEA hybrid cloud is based on a solid foundation of engineering expertise, existing open source software and commercial services:

- RHEA System Engineering & Cyber Security expertise[27].
- SixSq's Nuvla[28], a SlipStream-based hybrid-cloud management platform.
- Cyfronet's OneData for Data Management.
- Exoscale[29] IaaS Cloud Service (CloudStack[30]).

The RHEA Group HNSciCloud Nuvla multi-cloud solution provides the European science community a means to automate the deployment and the dynamic management of applications and minimise the costs of operations (see Figure 5), as well as monitoring resources consumption and quota management (see Figure 6).

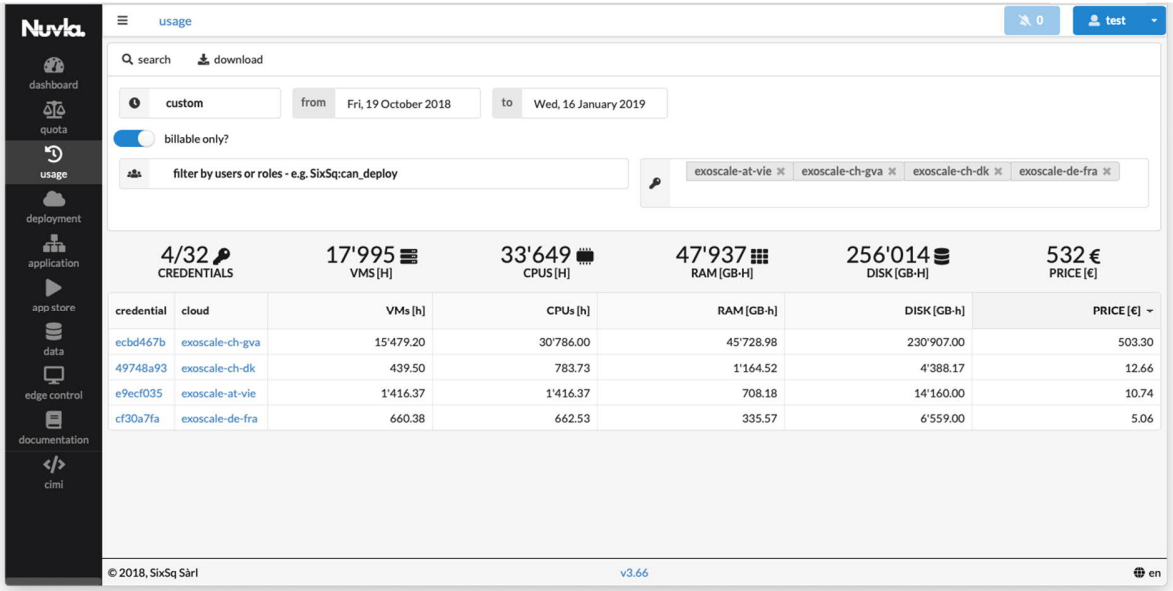

Figure 5 – Nuvla usage report

[27] https://www.rheagroup.com/
[28] https://sixsq.com/products-and-services/nuvla/overview
[29] https://www.exoscale.com/
[30] https://cloudstack.apache.org/

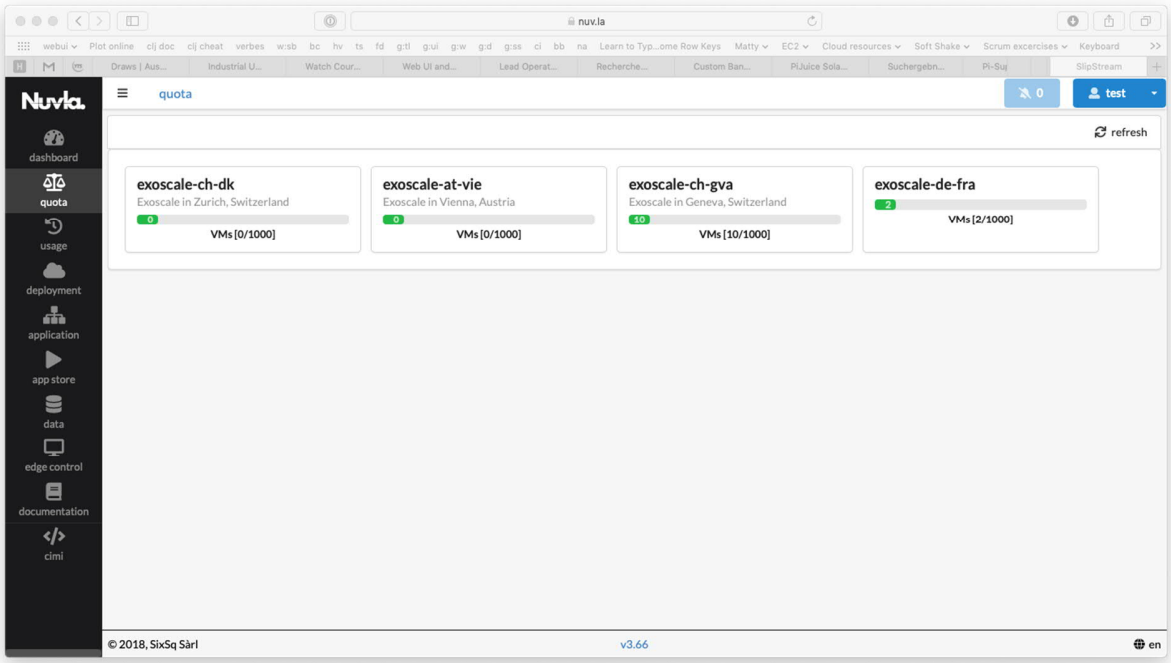

Figure 6 – Nuvla quota management

Further, in the context of the project, the RHEA hybrid cloud platform (i.e. Nuvla) has fully integrated the eduGAIN and Elixir AAI identity federations. KeyCloak[31] based implementation allows management of users, groups, and attributes defined by the federations' identity providers. This allows users to authenticate with their academic on-line credentials to deploy Virtual Machines (VMs) and/or containers to multiple clouds. The Nuvla's multi cloud platform allows users to choose which cloud to use based on price, performance, location or other factors which are important to them, and manage the full lifecycle of cloud applications with a high degree of automation.

Furthermore, in collaboration with the European Space Agency, this last feature was extended by SixSq to include policy placement based on datasets location, a potential key enabler for communities wanting to share data across organisations and scientific domains, using the cloud.

To support peta-scale data management solution, and deliver to the end-user a transparent data access and uniform view of all shared data sets no matter where they are physically stored, the RHEA Group hybrid cloud has been extended with Onedata, the open-source data management solution built around shared "spaces" developed by Cyfronet and GlusterFS/S3. This solution, which is integrated in the Nuvla cloud application management platform, allows users to access and to manage large datasets hosted in hybrid cloud infrastructures and/or at a Buyers Group organisation with minimum intrusion.

[31] https://www.keycloak.org/

To run High-Performance Computing application (HPC) on Nuvla two approaches are available: 1. direct use of virtual machines on the hybrid cloud platform to run a single MPI-based HPC job on a "raw cluster", or 2. deploy a dedicated batch system within the hybrid cloud platform to manage multiple HPC jobs. The HPCaaS solution offered by the consortium allow the deployment of batch clusters using high-performance hardware. These batch clusters can be configured with VM with up to 16 vCPUs and 128GB of RAM. The CPU flavours and the pricing model adopted by the Exoscale commercial cloud provider are reported in the tables below.

	CPU	RAM
Micro	1x 2198Mhz	512 MB RAM
Tiny	1x 2198Mhz	1 GB RAM
Small	2x 2198Mhz	2 GB RAM
Medium	2x 2198Mhz	4 GB RAM
Large	4x 2198Mhz	8 GB RAM
Extra-large	4x 2198Mhz	16 GB RAM
Huge	8x 2198Mhz	32 GB RAM

Table 5 – Exoscale's instance type

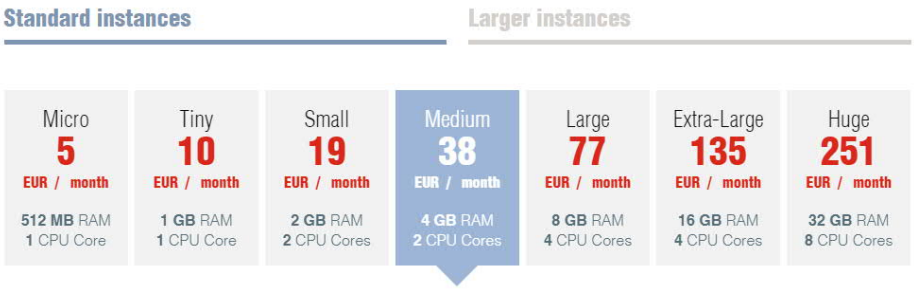

Table 6 – Exoscale's pricing model

From July 2018 onwards, Exoscale expanded the IaaS architecture integrating a limited number of recent GPU configurations. The GPU flavours available in the RHEA Group hybrid cloud platform are reported in the tables below:

GPU Profiles

vCPU	12	48
RAM	56 GB	225 GB
SSD (max.)	2 TB	8 TB
GPU Cores	3584	14336
GPU NVIDIA Tesla P100	1	4

Table 7 – GPU profiles available in Exoscale

All the services of the RHEA Group hybrid cloud platform are connected to GÉANT with 40 Gbps connectivity. Measurements of the reliability of the connections and the delivered bandwidth have been done using perfSONAR deployments.

RHEA has proposed a voucher scheme by which access to the pilot platform services can be granted to researchers sponsored by the Buyers Group for limited scale usage. The available documentation[32] describes how the user can either create a new account using the voucher provided by Exoscale, or add credits to a voucher in case of an existing account. This work benefited from shared effort with EPFL and the Human Brain Project to support the dynamic deployment of Massive Open Online Courses (MOOC), where both communities share similar goals to facilitate the on boarding of users.

Last but not least, a financial dashboard integrated in the platform also offers the possibility to monitor current activities, resource utilisation and costs (see Figure 5).

[32] http://hn-docs.readthedocs.io/en/latest/getting-started/exoscale.html?highlight=vouchers

Appendix II - Service Maturity

The following guide includes proposed characteristics to help assessing the Technology Readiness Level (TRL) of a service - either one entering EOSC or one within EOSC looking for a TRL re-classification. These definitions are results of the EOSC-hub RoP Task Force:

- TRL 6 (System/subsystem model or prototype demonstration in a relevant environment)
  - Representative model or prototype system is tested in a relevant environment and represents a major step up in a technology's demonstrated readiness.
- TRL 7 (System prototype, demonstration in operational environment)
  - Service has passed through development and is an advanced stage of pre-production. The software is stable, reliable and has been deployed in an operational environment.
  - Functionality as required by the target users is documented, understood, validated with target sample users and accepted by them. Internal documentation exists regarding preliminary validation tests.
  - An assessment has been made of the required load of the system once the transition into production is complete and a plan has been made to service this load. This assessment has been documented.
  - An SLA is optional.
- TRL 8 (System complete and qualified)
  - There are real users who are making use of the service and relying on it for their work.
  - Service documentation exists and is made available for end-users.
  - An acceptable use policy/terms of use/SLA is in place.
  - Evidence that the service is being delivered in a way consistent with user expectations.
  - Provision is made for user support, with response to incident and problem management.
- TRL 9 (Actual system proven in operational environment)
  - All requirements of TRL 8 are met.
  - Customer feedback is gathered and documented. The service has been in a production state and relied upon by users for at least 1 year and evidence is provided to show this.
  - There are quantitative outputs as a direct result of the service usage.