Helix Nebula – The Science Cloud

Deliverable Title: HNSciCloud Overall and Individual Exploitation Plans

Partner Responsible: TRUST-IT

Work Package: WP7

Submission Due Date: 30 November 2018

Actual Submission Date: 31 January 2019

Distribution: Public

Nature: Report

Abstract: This document reports the exploitation plans for the innovative solutions resulting from the HNSciCloud pre-commercial procurement action. It outlines the future intentions of the buyers group and the two contractors of the pilot phase of the HNSciCloud project.

Document Information Summary	
Deliverable number:	*D7.3*
Deliverable title:	HNSciCloud Overall and Individual Exploitation Plans
Editor:	*Sara Pittonet (Trust-IT Services), Silvana Muscella (Trust-IT Services)*
Contributing Authors:	
Reviewer(s):	
Work package no.:	*WP7*
Work package title:	*Stakeholder Engagement, Dissemination, Communication & Exploitation*
Work package leader:	*TRUST-IT*
Work package participants:	*CERN, INFN, DESY, CNRS, KIT, SURFSARA BV, STFC, EGI.eu, EMBL, IFAE*
Distribution:	*Public*
Nature:	*Report*
Version/Revision:	
Draft/Final:	
Keywords:	*Stakeholder Exploitations Plans, Commercialisation plans*

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.

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

Log Table

Issue	Date	Description	Author/Partner
V0.1	15/01/2019	First draft version	Sara Pittonet (Trust-IT Services)
V0.2	25/01/2019	First draft version available for partners' input and specific contributions	Sara Pittonet, Silvana Muscella (Trust-IT Services)
V0.3	27/01/2019	Comments and first internal review	Bob Jones (CERN)
V0.4	29/01/2019	Trust-IT Feedback integration and second internal review	Sara Pittonet, Silvana Muscella (Trust-IT Services)
V0.5	31/01/2019	Final version for Buyers' review	Sara Pittonet (Trust-IT Services)
V.06	31/01/2019	Contributions from RHEA Group	Alastair Pidgeon (RHEA Group)
V0.7	01/02/2019	Contributions from T-Systems	Jurry de la Mar (T-Systems)
V0.8	02/02/2019	Integration and final internal review	Sara Pittonet (Trust-IT Services), Bob Jones (CERN)
V1.0	02/02/2019	Final Version	Sara Pittonet (Trust-IT Services)

Executive Summary

Since its creation in 2011, Helix Nebula has grown to become a leading **public-private partnership between public research actors and cloud service providers**. The initiative has undertaken its first joint Pre-Commercial Procurement (PCP) tender called Helix Nebula Science Cloud (HNSciCloud) to create a common hybrid science cloud platform for the European research community. The project has been referenced by NIST/IEEE Joint Cloud Federation Working Group as good example of a **multi-cloud environment**.
The R&D performed enabled the 10 research organisations (including 3 intergovernmental organisations) hosted in 7 countries to successfully pilot more than 30 diverse use cases ranging from high-energy physics to photon-neutron physics to life sciences. Each use case provided significant challenges regarding federation in the hybrid cloud, transparent large-scale data management and service management. A key success factor was that the required capacity to support the use-cases was provided on-demand at short-notice. During 2018 it exceeded 20,000 cores and 2 petabytes of storage with a network bandwidth of 80Gbs, all provided by European providers.
All the services are based on open source implementations that do not require licenses in order to be deployed on the in-house IT resources of research organisations connected to the hybrid platform.
The HNSciCloud platforms have been actively promoted to other potential users as part of the early adopter scheme, notably in the context of European Open Science Cloud (EOSC) to provide production quality and commercially supported cloud services that are tailored to the needs of the research communities. Recently, three ESFRI research infrastructure clusters have confirmed intentions to procure commercial cloud services as part of their work programmes.

HNSciCloud has permitted the contractors and buyers group to explore and address key challenges associated with the adoption of commercial cloud services in the public research sector.
Today, HNSciCloud hybrid cloud platforms link together commercial cloud service providers and research organisations' in-house IT resources via the GEANT network. The platforms offers data management and compute services accessible via eduGAIN and ELIXIR federated identity and access management systems with support services, account management facilities, documentation and training.

HNSciCloud has played an important role in evangelising commercial clouds by demonstrating specific and challenging use-cases. These use-cases are relevant to a wide range of fields of science in Europe. Furthermore, new cloud infrastructures such as the Copernicus Data and Information Access Services services may also be interfaced with this multi-cloud platform, facilitating access by users to earth observation data. Advanced user and quota management, coupled with unified authentication (e.g. eduGAIN, Elixir) and seamless data management solutions, complete the value proposition of HNSciCloud service. Finally, strong synergies could be built to important federating initiatives such as EOSC-Hub.

As procurement rules and practices evolve, more value will become available from a flexible and dynamic multi-cloud eco-system. Scientists and researchers will have easier access to more choice, more scalability, better performance and enhanced security. Greater flexibility in the procurement rules and practices would allow users, their institutes and funding

agencies to take advantage of the favourable dynamics that applies to the public cloud market with on-demand capacity and overall cost savings.

The lead contractors in the final pilot phase of the HNSciCloud pre-commercial procurement, T-Systems and RHEA Group, have different profiles and offer different approaches to the cloud solution. For these reasons, the commitment of the two contractors' towards the uptake of HNSciCloud hybrid cloud platform is different and their future exploitation plans are explored. Details on the innovations that the contractors were able to introduce to their commercial solutions thanks to the HNSciCloud project as well as information about their commercialisation plans. Finally, it includes specific references to future commercialisation opportunities that the two contractors will put in place, as a potential source of new revenue-streams from the HNSciCloud research activity. The OCRE and the EOSC initiative are the most relevant for this context and are highlighted here:

- The **EOSC initiative** launched by the EC in November 2018 is potentially the greatest opportunity to stimulate and aggregate the demand from the science community for commercial cloud resources. EOSC, and the related projects funded to support the implementation of EOSC [2019-2020] is seen as the major channel for engaging with the European publicly funded research market. That said, the true market scale of EOSC is not well understood and the support from the EOSC funding agencies to engage commercial cloud service providers needs to be expanded through policy (modernising procurement practices to take into account the dynamism of commercial cloud markets) and financial access (moving from traditional grant-based provisioning to open tendering processes).

- The **OCRE H2020 project** provides the main near-term procurement opportunity for the public research and education sector to exploit the results of HNSciCloud in the context of EOSC. RHEA & SixSq will work with GÉANT and CERN as beneficiaries in the OCRE project to develop a procurement framework for commodity commercial cloud services. Several of the HNSciCloud buyers' group members as well ESFRI clusters have expressed their intention to participate as procurers.

Table of Contents

1. Introduction

Many experts have reported that to solve new and more complex questions in science it is eminently important that the newest innovations in IT, e.g., HPC and GPU-processing and professional IT services, become available not only to select groups but to almost every scientist involved in the global research communities. To make commercial cloud services available to a large community, the HNSciCloud PCP has been addressing key challenges encountered in the public research sector:

- integration of commercial cloud services with publicly funded e-Infrastructures to create a hybrid cloud structure;
- providing transparent data access to large-scale research datasets in a hybrid cloud environment;
- establishing service agreements and payment models to match the requirements and constraints of the research sector and public procurement.

The hybrid cloud services developed within HNSciCloud enable research organisations to run very diverse applications and workflows on any combination of dedicated hardware, private cloud and public cloud infrastructure. The service includes functions for Infrastructure as a Service (IaaS), transparent data access, large-scale data management, container management, HPC-as-a-Service, batch-systems integration, GPU and FPGA computing, network connectivity, federated Identity and Access Management (IAM), service payment and service quality monitoring. This rich set of services are also valuable for commercial applications with dynamic compute and data storage requirements in a range of business sectors.

2. The ICT Market in the EU

The IT market in the EU had a strong year back in 2017. The total value of the IT market grew 4.0%, compared to 3.5% in 2016 and the business outlook remains positive. The top IT priorities named in an IDC survey in May 2018 are:

- Improving the quality of IT support for the business
- Enabling digital transformation of the business
- Making the business more profitable
- Supporting revenue growth

The migration to cloud continues to have a dramatic impact and is fuelled by organisations' needs to become more agile, as cloud accelerates the deployment of IT services. At the same time, cloud is increasingly preferred by software and service vendors, so new services are launched earlier or exclusively in the cloud. As organisations execute their digital transformation, they will need both private and public clouds to provide the agility required.

The need for scaling out and scaling up continues to drive cloud investments in Europe. The market is also propelled by organisations' focus on reducing and shifting the cost structure from CAPEX to OPEX.

Using cloud services means relinquishing control and allowing an outside party access to systems and data. Consequently, organisations need to move focus from IT perimeter security to the data – whether in the cloud, on-premises, on end-user devices, or in transit. It means that organisations need to comply with legislation e.g., EU GDPR, the geographic location of the data, and so forth. As such, a decision to migrate to cloud entails significant planning and a clear strategy that may slow down certain organisations, particularly those in public sector with a large legacy infrastructure.

D-PIL-3.12 T-Systems Commercialisation Plan

The HNSciCloud activity is predicting that the Buyers Group (and ultimately the wider European science community) will need to augment their existing computing resources with commercial cloud resources to meet the massively increased demand in compute and storage – the hybrid cloud.
Therefore, there is potentially a huge demand for commercial cloud services in the science community. However, this will only work if there is a real commercial, sustainable market. The HNSciCloud initiative sets out to answer questions such as:

- Is the total cost of ownership compelling?
- Can it provide better services?
- Is the science community willing to embrace?
- How can it co-exist with existing research infrastructures?
- Are scientists and their applications cloud ready?
- Can existing procurement practices support it?
- Are there strong enough incentives from funders (e.g. Local/Regional/National governments, EU, etc.)?

It is unlikely, and arguably undesirable, that a single cloud provider will be able to deliver such a hybrid cloud that meets the expected needs of the European Open Science Cloud (EOSC). Therefore, in our view, the most likely outcome is that the EOSC will need to make available a pooled set of cloud resources combining Institutional Clouds, Commercial Clouds (e.g. European Clouds – Exoscale, OTC, Cloud28+, Orange and, inevitably, the "big three" – AWS, Google and Microsoft) and public European cloud infrastructures (e.g. provided by EGI or GÉANT).

D-PIL-3.12 RHEA Commercialisation Plan

3. The TCO Study

A Total Cost of Ownership (TCO) study for selected use cases was introduced in the pilot phase to help the buyers' group understand the impact of the contractors commercialisation plans for their organisations. The buyers' group have selected relevant use cases and provided the high-level input requirements and workflow information to perform the TCO study. Based on this, HNSciCloud contractors of the pilot phase, T-Systems and RHEA Group, have evaluated alternative approaches to support the use cases most effectively using commercial cloud services and derived the TCO per use case.

The PANCANCER and ALICE use-cases each have their own specific requirements for cloud deployment. Use cases are complex and there were many factors to consider in the TCO (e.g. performance and scaling, job efficiency, real world data through-puts, CPU contention and so-on). This required close working between the science community and the contractors' cloud experts to identify the most cost-effective solutions.

The PANCANCER & ALICE use-cases demonstrate that important savings can be made by shopping around and choosing a cloud that best matches a specific use-case. The ALICE use case is sufficiently large enough that cloud providers could be willing to invest in specific commercialisation plans.

The results of the TCO study led the two contractors towards the definition of commercialisation plans for the services implemented as part of their effort in the HNSciCloud initiative and that are described in the chapters below.

4. Commercialisation Plan of the RHEA-led consortium for HNSciCloud

4.1. Innovation of RHEA Solution: Exoscale and Nuvla

RHEA consortium includes **Exoscale**. The technical platform has already demonstrated interfacing with the Open Telekom Cloud (OTC), Amazon Web Service (AWS), EGI Clouds (e.g. CESNET and Cloudferro) and HPE Cloud28+ (e.g. Advania). In collaboration with the Institut Français de Bioinformatique (IFB), SixSq has already interfaced a number of its community clouds. It would also be possible to interface to Microsoft Azure, Google Cloud and IBM's cloud infrastructure (aka SoftLayer or Bluemix).
Fundamentally, integration of a new cloud service provider (CSP) requires that:

- RHEA supports the CSP's API;
- the CSP has an appropriate GÉANT connection;
- the CSP supports the network connectivity requirements;
- the CSP provides transparent resource pricing and;
- the CSP supports the minimum SLA requirements.

Further, the buyers own cloud infrastructures can also be connected to the RHEA platform, as illustrated with the case of the IFB - Institut Français de Bioinformatique and has also been demonstrated by interfacing with EMBL's Embassy cloud service[1].

RHEA's commercial offering is based on the **Nuvla**[2] platform offering seamless access to multi-cloud resources, where the user can choose the cloud provider(s) that best meets their needs at the lowest price. RHEA offers a *managed service* approach for users wanting maximum simplicity and a *self-service* approach to allow users greater flexibility and control.

[1] https://www.embassycloud.org/
[2] https://nuv.la/

Note: A current constraint of RHEA's offering is the GÉANT policy which prevents direct traffic between commercial clouds across the GÉANT network. For most use-cases this is not a problem, but where, for example, two public research institutes are using different commercial clouds (e.g. OTC and Exoscale) and want to share data, that data would have to be transferred via one of the public research institutes and not directly between commercial clouds, which introduces additional complexity.

The Buyers' Group use cases (and by extension those of the wider European science community) have different operational models. There is no one-size-fits-all solution. The flexibility and choice offered by the multi-cloud approach allows users to choose the CSP(s) that best meet their needs for compute and storage resources. In addition, different pricing models allow users to minimise the costs of the resource usages for their own use cases.

> "HNSciCloud has showed us that monitoring usage, managing quotas and metering resources, especially across clouds, is a challenge for any user community and customer institution, and can turn out to be time consuming and effort heavy. This is why we believe a management service, such as Nuvla, can significantly increase visibility over usage and costs, while providing tools to better manage and optimise resources consumption."
>
> D-PIL-3.12 RHEA Commercialisation Plan

Figure 1 RHEA's SWOT analysis of their commercial offering based on Nuvla.

A critical issue for the viability of this business plan is that there must be simplified mechanisms in place by which multi-cloud services can be easily and quickly procured by the Buyers Group members and the wider community. The HNSciCloud hybrid-cloud concept is a fundamental change from the traditional procurement of hardware to the provision of cloud services, including Infrastructure as a Service (IaaS). While most publicly funded bodies have well established mechanisms for procuring hardware and support services, which deliver services from a fixed consortium for a fixed period of time, these practices are not necessarily well suited to a multi-cloud approach where the IaaS services are, to some extent, commoditised and the market changes rapidly. This means the cloud provider that delivers the best price-performance today may not do so next year.

Therefore, RHEA's commercial offer must be matched by new (or modified) procurement practices to allow the Buyers Group and wider scientific community to get the maximum benefit from what the hybrid cloud service can offer.

The brokerage and management of the Nuvla service are included within the discounts compared to public pricing. The benefits of the Nuvla service include the following for multi-cloud and hybrid-cloud:

- Frictionless cloud provider switch
- Federated identity authentication integration
- Service catalogue for pricing prediction
- Dashboard and standard API (based on DMTF CIMI)
- Rich and flexible usage metering
- Quota monitoring
- Application deployment automation and management
- Initial voucher support
- First-line support

Exoscale will introduce the concept of pre-emptive virtual-machine (VM) scheduling and CloudFerro also has an interest in offering a similar service. AWS has a similar system for academic users. The pre-emptive approach could benefit use-cases such as ALICE with low-priority jobs that can benefit from low-pricing by exploiting the spare capacity of a CSP. If no single CSP has sufficient spare capacity, jobs could be deployed on several clouds. Pre-emptive (headroom or opportunistic) scheduling is simpler to understand and use than spot pricing and allows costs to be determined in advance and can be useful for sporadic use-cases or to handle certain peak loads.

While benchmarking has shown Exoscale to be very competitive compared to AWS, no CSP is consistently cheapest for all use cases, meaning the ability to easily access and deploy to multiple CSPs should be an option when striving to get the best value for money for the Buyers' Group and, ultimately, the European tax payer. Nuvla makes it commercially and contractually easier to achieve this kind of flexibility.

4.2. Minimum Viable Market: the EOSC "freemium" model

A recent paper[3] from the Science|Business Network's Cloud Consultation Group suggests there may be a market of up to €2.25B per annum, based on 5% of €45B per annum research spending. While this is a very large figure, it appears to assume that all IT spending would be in the cloud from day one, whereas it may take a number of years to reach this level, if at all. The paper proposes the "freemium" model for the EOSC to make commercial IaaS & PaaS, as well as tools and data available to the science community, free at the point of use[4].

[3] "How the Science Cloud could pay its way, Business models for a sustainable European Open Science Cloud", A report of the Science|Business Network's Cloud Consultation Group, July 2018, https://sciencebusiness.net/science-cloud/news/how-science-cloud-could-pay-its-way

[4] In chapter 4 "EOSC Business Model: Financing the EOSC" of the Final report and recommendations of the Commission's 2nd High Level Expert Group on the European Open Science Cloud[4], three main business models for provisioning access to Research Infrastructures are envisaged: A. Excellence-Driven Access; B. Market-Driven Access and C. Wide Access mode. The document also provides a comparison among possible funding/revenue models to ensure that data is open (opened up) and shared (and reused), as well as payment models to envisage how money is transferred to ultimate data and service providers: Direct Support, where elements of EOSC receive direct

The EU plans to spend €400M on EOSC related activities in the next 2-3 years. RHEA have analysed this spend to assess what is the likely spend on commercial cloud services. This is summarised in the table below. As can be seen, the actual spend on commercial cloud services could be less than 10% of this over three years and very small in the next two years.

Topic	Title	Action	Budget (€M)	Accessible to Commercial Clouds	Comment
INFRAEOSC-01-2018	Access to commercial services through the EOSC hub	RIA	12	4.5	OCRE over 3 years
INFRAEOSC-02-2019	Prototyping new innovative services	RIA	28.5	0	
INFRAEOSC-03-2020	Integration and consolidation of pan-European access mechanisms to public e-infrastructures and commercial services through the EOSC hub	RIA	79	35	OCRE follow-on, but assume 50% goes to e-infrastructures
INFRAEOSC-04-2018	Connecting ESFRI Infrastructures through cluster projects	RIA	95	0	
INFRAEOSC-05-2018-2019	Support to the EOSC governance (a) Setup of an EOSC coordination structure	CSA	10	0	
INFRAEOSC-05-2018-2020	(b) Coordination of EOSC-relevant national initiatives across Europe and support to prospective EOSC service providers	RIA	30	0	
INFRAEOSC-05-2018-2021	(c) FAIR data uptake and compliance in all scientific communities	CSA	10	0	
INFRAEOSC-06-2020:	Enhancing the EOSC portal and connecting thematic clouds (a) Support to the EOSC portal	RIA	5	0	
INFRAEOSC-06-2020:	(b) Connecting thematic clouds into the EOSC	RIA	5	0	
INFRASUPP-01-2018-2019	Policy and international cooperation measures for research infrastructures (b5) Support to the eInfrastructure Reflection Group (e-IRG) CSA	CSA	0.6	0	

payments from funding agencies; Pure, with: Researchers use EOSC vouchers ('cloud coins') to support EOSC services/data; or a Hybrid model: EOSC vouchers and direct support.

Other actions GÉANT Partnership projects	Two actions covering the support to the pan-European data network for research and education, including increase of the backbone capacity, which underpin EOSC SGA Q3 2018 Q4 2018 €128m	SGA	128	0	
			403.1	**39.5**	

Table 1 EU plans to spend on EOSC related activities in the next 2-3 years

As can be seen from the table, the eco-system surrounding the EOSC is in its preliminary stages. Therefore, any marketing and outreach strategy will have to take this eco-system into account, by either contributing to it, or supplementing it. It is not yet clear, for example, the role of existing research infrastructures in this ecosystem and whether they will also be competing with the commercial providers for a share of the EOSC market. One of the prime features for 2019 is that Research Infrastructures should commit, with KPIs on their practical involvement to the EOSC ecosystem, together with the newly funded CSA on supporting EOSC as well as with the EOSC Executive Board, distinct practical measures will be put into place to foresee the engagement.

Due to the uncertainties in the market, RHEA constructed a model to assess what would be the minimum viable market for RHEA's approach to be profitable. The model tries to estimate what they believe is the real addressable market (taking account of EOSC and a fraction of the €2.25B spend, increasing slowly each year). They then assume a maximum of 20% of this is available to RHEA (i.e. RHEA acts as the broker for it). This is probably a high figure. Of this, they assume 10% revenue. This provides the annual expected income against which are the costs of sales, operations, leaving a profit (or loss).

The simple analysis shows that for the RHEA business model to be profitable the annual volume of the addressable market has to be significantly greater than it is today.

RHEA's model assumes that a small proportion of the science spend on infrastructure will start to be spent on commercial cloud services. This is very conservative but even so is ten times the foreseen EOSC spend over the next few years. The model is clearly very sensitive to whether that exists or not. Without it, the volume would not be sufficient for RHEA at all, nor for Exoscale until ~2022.

4.3. RHEA Commercial Offering for the NuvlaScience

RHEA's HNSciCloud offering is built on the existing commercial Nuvla service to deliver large-scale IaaS services, via access to a wide range of IaaS providers, exemplified by OTC, Exoscale and Advania (all demonstrated during the project), as well as specialist services such as managed data solution based on Onedata, by Cyfronet.

RHEA's offering is open to other suitable cloud providers to give the buyers group the widest possible choice, near unlimited scalability and to avoid lock-in.

As part of HNSciCloud, the Nuvla service was upgraded to fulfil buyers' requirements to provide unified authentication leveraging federated identity management such as eduGAIN and Elixir. This feature also allows organisation owners to manage users via groups and roles. Further, unified features such as usage monitoring, metering, costs estimates, access to quotas, benchmarking, and reporting were deployed to facilitate management of the multi- and hybrid-clouds. Advanced usage records, with flexible query capabilities and fined grained support (e.g. user, group, organisation) provides users with a unified view over their usage, over a user defined period. Costs estimate is also available (e.g. based on per minute measurements) based on service offers and available pricing.
Importantly, these above features work whether or not Nuvla is used to provision cloud resources. This reduces constraints on users, while increasing the reach of benefits from Nuvla.

Finally, the Nuvla API now provides a powerful tool for users to query the system and build simple, yet sophisticated, queries to generate a range of reports for downstream use.

Benefiting from this development and experience, RHEA can offer these services under commercial conditions to the Buyer's Group members, their communities and other European academic organisations.
The proposed commercial service offering recognizes the wide variety of needs and types of users in the Buyers Group and can respond to scale:

1. **Self-Service**: customers have their own contracts with cloud, using Nuvla to simplify management and provide a unified web, CLI, API and tools to all clouds. This model is flexible and allows users to manage their cloud contracts themselves, at the cost of having to manage each contractual relationship.
2. **Brokerage**: RHEA offers a single contract to each customer, providing frictionless access to all the cloud services it has reseller or equivalent agreements with. This allows customers to take advantage of a range of purchase options, including Pay-As-You-Go, Bulk Subscription, Reserved Instances and/or Storage, or a combination of these.
3. **Managed Services**: these services are compatible with both above Self-Service and Brokerage and can include: Managed Data Management, VPN as a Service, container support, HPCaaS, SLA support, Consultancy Support, etc.

All of the above commercial offerings share a common important point: **the services are free at the point of delivery.** In other words, users are given quotas by organisation or group owners, from which they can consume cloud services.

The charging and billing granularity are dictated by cloud providers and RHEA simply propagates these. For example, Exoscale charges per minute, while others might charge per hour. CSPs might invoice monthly, but not at the same date. Nuvla's accounting system adapts to these schedule differences with a daily aggregation, based on a per minute accuracy. Billing statements are issued monthly, following the cloud providers' schedule.

For those users wishing to use Microsoft Windows or other licensed products, RHEA's Bring-Your-Own- License (BYOL) approach does not impose any restrictions, unless coming from cloud

providers. This is currently not an issue with available cloud providers on Nuvla (e.g. OTC and Exoscale).

4.4. Relationship with the Cloud Providers: reseller agreements

Some cloud service providers require dedicated contractual relationships with a cloud broker in order for the cloud broker to be allowed to resell their cloud services.
Different cloud providers operate differently, with their own terms and conditions on users and resellers, and have different organisational structures. In general, only resellers will have a clear contractual context through which they can resell cloud services. With complex rules such as data privacy issues, reselling cloud services outside such a contractual agreement is risky, especially in case of dispute and litigation. This relationship also generally allows the reseller to bulk purchase or be rewarded for sales, such that these rewards can then be shared with its customers in order to make its value proposition more attractive.

Some organisations operate at a national level and could allow resellers to contract directly with the national entity of their choosing, but this is not the norm.

4.4.1. Relationship with Cyfronet (Onedata)

Transparent data access is a key challenge of HNSciCloud and is supported by Onedata[5]. Although the Onedata software is open source and free of charge, support for the product would be offered as a managed service.

For a Onedata based managed service, RHEA would have to setup a contractual arrangement with the commercial organization Cyfronet is setting up to offer professional services on their open source product (in a Red Hat type model). RHEA/SixSq would continue offering 1st line support (single entry point) and this commercial entity would provide 2nd and 3rd line support. This would offer best user support to customers. And since SixSq's technical knowledge in Onedata is already significant, the quality of the 1st line support should remain high, even for non-SixSq software-based managed services.

4.4.2. Integrating a New Cloud Provider

RHEA's multi-cloud service and its architecture allows the addition of cloud service providers. However, for a cloud to be eligible to be offered as part of the HNSciCloud offering, it must meet a set of minimum requirements of the buyers' group, including:

- a high-speed GEANT connection (40 Gbps)
- minimum performance requirements
- minimum SLA requirements, provide public and open pricing
- be within an EU member state or other eligible state.

[5] https://onedata.org/

- IaaS API
- provide pricing information

If these criteria are satisfied, the next step it to perform the technical integration. Experience has shown that if the cloud provider does not have a GÉANT connection or requires additional bandwidth, this can take weeks or even months to set-up, depending on the NREN.

In the case where a connector would not exist, a commercial discussion must take place to understand if and how to proceed with development of the missing connector.

The software used by the RHEA team for the HNSciCloud data management is entirely open source, for which no royalties are required. Further, complete documentation is provided to assist on-site software installations. Furthermore, if customers private cloud (if available) are connected to Nuvla, it can be used to automate software installation of several applications.

4.5. Nuvla Fees

The Nuvla platform is funded via volume fees and/or user subscription. The RHEA objective behind this model is to find the right balance between minimum operational revenues to cover operational costs, while ensuring a minimum upfront cost is required from customers. To this effect, the brokerage service is included in the nominal IaaS costs for public clouds or negotiated for private clouds. The self-service offer can include a combination of a reduced revenue fee and per user subscription.

These models are currently evolving, taking customer feedback into consideration. **RHEA's ultimate objective is that the value of the Nuvla service far outweighs the service fees** (both revenue share and user subscription). This delivered value includes:

1. Unified usage metering and costs monitoring, across all available cloud services.
2. Support for federated authentication (e.g. eduGain, Elixir) to simplify access to cloud services
3. Ability to implement Most-Economically Advantageous Tender (MEAT) purchase strategy and provisioning mechanism by allowing users to express their requirements in the form of policies to rank providers
4. Benchmarking feature, allowing users to benchmark their application on all cloud candidates, such that they can compare cloud services based on performance directly relevant to their workload
5. Differentiated pricing, providing the ability for providers to make special offers to customers, or for customers with an existing contract including a discount, to be registered as such in the Nuvla pricing service.
6. Platform ability to provide unified interfaces (i.e. web, CLI, API, tools) to all clouds, simplifying operational costs while providing large services choices.
7. Saving from bulk purchases by RHEA shared with the customers, providing prices lower than standard market conditions.
8. Automation, custom reporting, user management and monitoring benefits from the Nuvla service
9. Benefiting from regular addition of new cloud providers to the catalogue to enrich choices.

10. Access to voucher scheme, such that new users can more easily try the hybrid cloud, or used by community managers, encourage its users to join the services.

Through HNSciCloud, these features are being adapted to ensure maximum benefits to the Buyers Group.

Managed services, including managed Onedata, cluster, continuous monitoring and benchmarking and container management would be charged separately.

One of the key value propositions of the broker (Nuvla) is its ability to compare pricing and monitor the costs of the deployed VMs. This means that customers can continuously ensure that they get the best deal the market can offer and deploy their applications to their chosen cloud. Further, it allows them to integrate in Nuvla private clouds they might have, thus building a hybrid and interoperable model at low cost. And finally, it allows them to experiment with other clouds, at low cost and risk, such that by the time their cloud contract is due for renewal, they get the best possible deal.

4.6. Payment Models

Model	Description	Applicability in HNSciCloud & envisaged benefits (if applicable)
Pay Per Use	Pay-per-use offers the user (consumer) great flexibility; they only pay for what they need, when they need it. For the supplier, it means they cannot predict demand and therefore have to have excess capacity to handle un-expected demand or peaks of demand. For on-demand use, the payment granularity could be an important choice depending on the users' operational model – e.g. short-term VM usage (fractions of an hour) would benefit from a pay-per minute provider, compared to pay-per hour provider.	The policy engine of Nuvla includes this type of functionality, such that the user is be able to enter its expected resource profile and be presented with the cloud or clouds which provide the best price. One of the advantages of the multi-cloud model offered by the RHEA consortium is that users can choose the cloud to match their needs rather than a one-size fits all approach.

Model	Description	Applicability in HNSciCloud & envisaged benefits (if applicable)
Reserved instances	Reserved instances mean the user commits to using a certain volume of resource (processing and storage) for a period of time – typically 1, 2 or even 3 years. This guarantees the provider with fixed revenue and therefore can offer substantial discounts. As well as the discounts, the user can reliably forecast its costs for the duration of the reservation. However, the user is then committing to a fixed minimum expense for a fixed time period.	One mechanism that could be explored is to have an adaptive approach to reserved instances. For each period (nominally a year) the usage of VMs would be analysed and the broker would forward buy a certain proportion of VMs to cover those which are always active. This is a conservative approach but filters-out peaks of high usage.
Payment in Advance	Cloud providers can offer attractive discounts if users pay in advance. Combined with reserved instances, this gives both a positive cash-flow and clear forward loading of resources, which can assist with planning of hardware purchases, upgrades and expansions.	
Volume Discounts	Some cloud providers will give volume discounts for very large consumers of processing and storage, but will typically want this to be combined with the use of reserved instances to guarantee demand, rather than high-levels of sporadic use	

Model	Description	Applicability in HNSciCloud & envisaged benefits (if applicable)
Pre-Emptive and Spot Pricing	Pre-emptive and spot pricing is when a cloud provider offers its spare capacity at potentially high discount – e.g. on a non-priority basis (where a user can be kicked-out if a full paying customer needs resources) or at a time of lower demand (e.g. nights and weekends). With spot pricing the capacity is provided to the organisation who offers the highest price for a given time window but this is complex to administer. Pre-emptive pricing offers a fixed discount on the understanding that the VM can be pre-empted (shut down) at any time – but the user then does not pay. Exoscale and Amazon offer such a scheme and is simpler for users and providers than spot pricing but the latter could be attractive for some use-cases.	While significant discounts could be achieved via this mode, the provisioning and procurement strategies need to support this dynamic provisioning strategy e.g. using HTCondor, Slurm, Nuvla or some other mechanism.
Pooling & Sharing	The HNSciCloud is an example of a group of buyers coming together to pool their demand and, potentially, share resources (i.e. processing, storage and networking).	The largest economic benefits and maximum flexibility for the Buyers Group could be gained by pooling the demand for a certain base level of reserved instances, storage and networking, then using on-demand resources until another defined base-level is determined by the Buyers Group or Group member.

4.7. Consulting & Support

From time to time a member of the Buyers Group may wish to obtain additional consulting or support for their own specific needs. This could include assistance for "on-boarding" their applications, training for their team, custom extensions to meet specific requirements, feasibility studies, security assessments and so-on. Customer organisations might also want to access expertise in setting up their cloud purchasing strategy.

4.8. Procurement Practices

In this section we explore a number of current (and potential) procurement models and the influence they could have on the hybrid multi-cloud approach.

Practice	Description	Influence on HNSciCloud
CAPEX vs. OPEX	The classic result of the introduction of IaaS is that there is a transfer of spending mode from capital expenditure (CAPEX) to operational expenditure (OPEX). Many public organisations have annual budgets for capital expenditure and use these budgets to purchase computing infrastructures plus operational funding to maintain and operate the infrastructures. In some cases, additional funding may come from collaborative projects (e.g. H2020) which also have budgets for infrastructure procurement.	In times of austerity, even though OPEX is being cut, CAPEX may nevertheless be available as it is often regarded as an investment in the future. This need to be aligned with new services, such as IaaS, since they potentially can provide much greater flexibility and actually be cheaper over time. Therefore, a mechanism needs to be found by which IaaS procurement can be seen as equivalent to infrastructure procurement, operations & maintenance, perhaps aided by pricing approaches including partial advanced payment to secure maximum discounts.
Most-Economically Advantageous Tender (MEAT)	The Most-Economically Advantageous Tender (MEAT) is a procurement model that allows public bodies to evaluate different bids against a clearly defined set of criteria (e.g. compared to lowest priced fully compliant bid). The procurer must clearly set-out which criteria must be satisfied and the weighting that apply to each of them in the evaluation, and the evaluation model that will be used. The advantage to the procurer is that they can evaluate different bids against a common set of criteria and get the best balance of price and quality rather than, say, a cheap but minimally compliant bid.	MEAT can be applied to the hybrid and multi-cloud approach, where the buyers can define its own criteria, such that they choose systematically the most appropriate cloud offer(s). This selection is based on a clear and specific set of criteria to best suit its use-case/operational model. The Distributed Management Task Force (DMFT) Cloud Infrastructure Management Interface (CIMI) standard based policy placement feature of Nuvla enables this type of rich policy placement logic. This powerful tool is available to buyers to experiment cloud ranking along the lines of the MEAT principals.
Service Contract	Service contracts are increasingly common in public procurement whereby the procurer defines an initial set of services to be provided over a specific period and the bidders offer a price for those services and maximum daily/hourly rates for additional unspecified services during the period	With the hybrid multi-cloud approach, the procurers could request a service (e.g. compute, storage, networking) which must meet a certain SLA, be within a maximum price and subject to certain constraints usually defined in the contract or statement of work.

Practice	Description	Influence on HNSciCloud
	of the service contract. During the service contract, the procurer can request additional services based-on the rates. The service contract is governed by a Service Level Agreement (SLA).	Then, it is up to the provider how to deliver that service with one or more clouds, choosing the clouds on the most economically advantageous basis, including deciding whether to bringing-in new cloud providers and not.
Fixed Frame Contracts	In a frame contract the procurers select one or more consortia usually led by a prime contractor based on a set of technical, financial and other criteria. The procurement normally covers a period of 3-5 years. The bidders may be asked to quote maximum prices for the duration of the contract. Under the frame contract the procurer issues Requests for Proposals (RFP) to the selected consortia who then submit bids to compete for that work in a restricted competition. The consortia can be “fixed” or “flexible”	In the HNSciCloud context the “fixed” frame contract would potentially lock-in the procurer to the same consortia for a long period. In a highly competitive and dynamic market, such as ours, it could mean that cheaper and/or more innovative cloud providers are “locked-out” for the duration of the contract. This could be addressed by creating a “flexible” frame contract, which would allow consortia to bring- in additional provider(s) where it can be justified, based on agreed criteria (e.g. pricing, capacity).

5. Commercialisation Plan of the T-Systems-led consortium for HNSciCloud

The HNSciCloud Project results have significantly contributed to improve and extend the Public Cloud services portfolio of T-Systems. The new functionality enables T-Systems to address more application, more potential user communities and more business.

The integration of the new functionality already started during the PCP project. The OTC roadmap maintains two major release updates per year. As a result, newly developed solutions from the PCP are commercialised in a short time frame, and several new functions have already been included in the 2017 and 2018 releases.

A remaining technical concern is the significant number and diversity in eduGAIN implementations in Science and the lack of user entitlements in eduGAIN systems. Even when commercial cloud services can be federated with eduGAIN, significant work may continue to be needed to onboard new users and manage authorisations, thereby risking slow take-up. A consideration could be the evaluation of new ideas and approaches for a complete revamp of identity management based on newer technologies e.g., blockchain while continuing to improve the existing landscape.

5.1. Innovation of the T-Systems Solution: OTC, Onedata, HPC, DODAS Service

The T-Systems solution combines innovation that resulted from the PCP, creating value for science users and organisation with innovations driven by the global OpenStack community. The advantages are that the value created benefits both science and industry and thus sets the **Open Telekom Cloud** apart as a viable alternative to other global public cloud services. Being part of the OpenStack community has enabled T-Systems to introduce new powerful functions during the prototype and pilot phase in the **OTC** production environment. These include new VMs flavours; new high-performance compute flavours with InfiniBand; new network and storage functions; scalable file services; relational database and big data services; the multi- project function and the Tag Management Service to allocate resources and consumption to e.g., cost centres.

5.1.1. Data Management

The PCP has provided the start to directly integrate large-scale data management with IaaS services. User requirements, however, are diverse. By selecting **Onedata**, it was possible to implement various use cases with different needs.

5.1.2. HPC-as-a-Service

The T-Systems HPC-as-a-Service represents an optimal technical and commercial solution, whereby workloads can be treated differently with regards to throughput and capability, making best use of the large variety of the infrastructure resource type options. It could be demonstrated with both the leading workflow managers SGE and SLURM, that job scheduling and management can be highly automated, thereby enabling the less-experienced scientist to work with the service directly and reducing the workload on experienced IT managers to support them. Other workload

managers may be added in future according to user demand, due to the broad support in the OpenStack community.

A new innovation that could be added stems from the work with CERN in the prototype phase. Scientist have shown a strong interest to evaluate machine and deep learning frameworks for simulation e.g. to save large amounts of resources being used for Monte Carlo calculations. Due to the flexibility of the T-Systems solution, it was possible to run various learning scenarios making use of P100 and V100 GPUs, providing valuable feedback to the science team for future production.

D.PIL-3.12 T-Systems Commercialisation Plan

5.1.3. Container Management

The container management enabled a full dynamic and on-demand analytics service for the complex CERN CMS use case. The so-called **DODAS service** is unique in functionality and performance, it can initiate a full datacentre within 2 hours with all the required hard and software services and following thousands of concurrent analytics jobs can be scaled up and down within minutes. DODAS provides a sustainable and stable operation for jobs running several hours or days and includes innovative self-healing capabilities when micro services need to be restarted.
CERN scientist also successfully demonstrated how container clusters can be federated in a multi-cloud, in order to scale up and out across federated pools of containers.
A trend in industry being observed is to enable fully automated workload distribution in multi-cloud solutions through container management and container federation as an alternative to cloud management platforms or brokers.

5.1.4. Identity and Access Management

T-Systems solution enables federated identity and access management for eduGAIN and ELIXIR AAI through easy-to-use menu functions. Federation includes conversion rules to assign specific and admin rights in OTC tenants and sub-tenants. Federation also supports Web UI, CLI and API access. Metadata updates between the external IdP and T-Systems solution can be synchronised automatically. Both for OTC and Onedata as a service provider, entries in the eduGAIN repository https://technical.eduGAIN.org/entities have been created and maintained.

5.2. Payment & Service management models

5.2.1. Service Management

T-Systems solution supports an ISO/IEC-type Service Level Agreement, branded as “Enterprise Agreement”. With Enterprise Agreements, OTC offers extended support for commercial services in a hybrid cloud for an additional fee – such as faster response times, architecture support, and dedicated service and delivery management. On individual tenant level, users can select a Silver, Gold and Platinum service level, with a minimum commitment of 12 months. The services are backed by service commitments and users are entitled to service credits when commitments would

not be met. Prices are published and range from 2,500 Euro (Silver) to 5,500 Euro (Gold) to 17,500 Euro (Platinum) per month per tenant.

5.2.2. Service payment

The new **Financial Dashboard** interface is considered very valuable for users, since it enables the customizable, highly granular and daily distribution and evaluation of cloud budgets and consumption. It especially facilitated the unique requirement of the Buyers Group, with a central fixed budget of resources needed to be distributed in a flexible manner over more than 30 use cases in 9 organisations with weekly updates and changes. The dashboard, combined with the quota functionality managed by T Systems avoided unexpected consumption overruns and usage could be managed according to plan.

A state-of-the-art and extensive TCO study was performed for two selected use cases to demonstrate how service payments can be combined in an optimal way to achieve lowest TCO for the user. And the importance of cloud resource benchmarking to be able to select the best resource in terms of price/performance for an application could be demonstrated.

The new innovative functionality will be made available to all OTC customers in early 2019 and is considered a key outcome and impact of the HNSciCloud project.

D-PIL-3.12 T-Systems Commercialisation Plan

5.3. Cloud Credit Model

The Cloud Credit Model was based on the concepts outlined in the EOSC studies[6]. It should be noted, that the implementation of Cloud Credit Models is a functionality that is deeply embedded in business and operational support systems. Therefore, the design and implementation are time- consuming and shall be done with care to avoid iterations. Since during the pilot there was not sufficient time for a design process, assumptions had to be made regarding the user functionality. Although the basic functionality provided could be used, the result shows that in a future scenario sufficient interaction between Buyers Organisations and potential Service Providers shall be considered to discuss and design the user scenarios and avoid multiple iterations.

5.4. Data Protection and Security - Concepts and Policy

During the pilot phase the Buyers Group evaluated the Data Processor Agreements provided by T-Systems. The agreements were found to be compliant with EU GDPR and user requirements. An open issue identified was the need to better reflect the “diplomatic status” of some of the science organization’s e.g., CERN and EMBL. This will be part of ongoing work in T-Systems to be included in updated Agreements in early 2019.

[6] The National Institutes of Health in the United States has conducted a *cloud credit* experiment. Eight service providers (e.g. Infrastructure, Platform, Software), became conformant providers. Approximately $3m were distributed as cloud credits.

5.5. Innovative Competencies Resulting from the HNSciCloud PCP

Innovation and a very short response time to new customer demand are key success factors for commercial cloud services, which are expected to face fierce global competition over the next decade. **Proof of the innovation in OTC has been demonstrated with a large set of new functions and performance options that have been introduced since the market launch in March 2016, resulting from the PCP to address the needs of science and research organisations.**

- HPCaaS has been fully integrated into the OTC UI and API and includes specific HPC resource types for high-performance and high-I/O processing (in clusters), MPI-library and InfiniBand fabric.
- GPU- and FPGA-optimized compute resources have been introduced and will be extended in 2019 with further higher-performance options.
- Large-memory compute resources have been introduced with up to 940 GB RAM.
- Dedicated server and bare metal services have been introduced with up to 96 cores per
- physical server.
- New public server images are being provided, e.g., for HPC resource types.
- ONEDATA Data Management solution has been provided on a project basis, including block and object storage access for servers and docker containers, the import of existing data collections and a federated authorisation management.
- Tag Management Service to add metadata to cloud resources for cost allocation.
- Multi-project feature to support isolation, quota management and organisation models within atenant.
- High-bandwidth connectivity to GÉANT Cloud VRF (currently up to 40G).
- A new SDN-based Private Line Access Service to easily connect customer data centres to OTC through a layer 2 network
- perfSONAR network performance monitoring.
- Detailed usage reporting and accounting through a new Enterprise dashboard
- Enterprise service level agreements based on ISO / IEC 19086.

5.6. Data Management Competencies

Distributed data management for (large) datasets is a common pain point within research activities. In many organisations, however, little experience may be available on how to migrate to new scenarios and use existing infrastructures more efficiently. While the Onedata solution can address these types of challenges, such diverse scenarios are still complex to implement.

Typical **new service modes of operation** that have been tested by T-Systems and are the basis for commercialisation include:

- Onedata as a high-performance service, with a preconfigured, validated and automated setup that can be run as a self-service by an individual tenant.
- Onedata as a shared service for generic purpose applications in diverse use cases, for which an account is only used to obtain (federated) access. An example is the new Copernicus DIAS service, where collaboration and data sharing are a key requirement. Onedata-as-a-Service will provide users with easy access to diverse datasets that may include remuneration options for usage.

5.7. Onedata.io S.A spin-off

To address the demand for the Onedata data management solution and user support, a spin-off company named **Onedata.io S.A.** has been created. The new company will be in charge of support and evolution of the software.

The company will offer various support models with regards to performance and service response time. In addition, the company will provide development of custom-made features for specific needs to tailor the solution to individual demands and solution designs for particular needs.
Contracting will be preferably through value-added resellers, e.g., public cloud providers, that would provide Onedata with an infrastructure service. T-Systems is planning to become a registered VAR in order to to support its customer base on Open Telekom Cloud.

5.8. Price Models

One of the main arguments for using cloud computing is its dynamic availability – with costs based on actual use. This means it is quick and easy to ascertain the prices for the hourly usage of virtual machines, pricing issues become less clear, however, when additional components are also used. This is important to note, as hardly anyone uses only VMs. Network connection, storage, licenses for software (operating systems, middleware, databases and applications) – increasingly, complete IT environments comprise many components. Self-service always means that users themselves have to keep track of costs.

With the **Open Telekom Cloud**, bills are issued every calendar month. In these bills, the services used are offset against any credit balance and free contingents (from reserved contracts).

For the core offering of the **Elastic Cloud Server**, in addition to hourly **pay-as-you-go models**, there are also reserved models available that offer users discounts of up to 60 percent on regular orders of resources over 12, 24 or 36 months. **Two discount models** are available: monthly advance payments and pre-financing of the entire period (upfront). The upfront model offers the largest discounts.

6. HNSciCloud legacy & sustainability path: the OCRE project and the Voucher scheme

In science the user of an IT service is usually not the funder. The funder may even be another organisation with another governance. This may significantly slow-down innovation since the user may not get access to newest service options available in the market and a service provider may face a time-consuming complex decision-making process before a new service is procured.

The OCRE H2020 project [INFRAEOSC-01- 2018] begins January 2019. It is led by GÉANT but involving CERN, RHEA and SixSq will provide a means for the science community to access commercial cloud services. Approximately €5M will be available in the form of vouchers and other innovative mechanisms that can be spent with cloud providers over a period of 5 years. It builds on the work already carried out in HNSciCloud and provide a mechanism by which cloud resources can be procured from multiple vendors. INFRAEOSC-03-2020 will build on this with a substantial budget.

> Projects such as OCRE will be a valuable companion to provide insights into the appetite of different communities in consuming commercial cloud services. One of the desires of the consortium is to get users (or at least their institutes and projects) used to the real costs for the IT infrastructures they require, and therefore factor in their strategy the need to pay for these.
>
> D-PIL-3.12 RHEA Commercialisation Plan

OCRE will also introduce the **Cloud Credit Model as described by the 2nd EOSC High-Level Expert-Group Report[7].**

HNSciCloud's results have encouraged the ESFRI cluster projects ESCAPE[8]: European Science Cluster of Astronomy & Particle physics ESFRI research infrastructures; EOSC-Life: Providing an open collaborative space for digital biology in Europe and PaNOSC: The Photon and Neutron Open Science Cloud, to include the procurement of commercial cloud services as part of their work programme. OCRE will also liaise with these ESFRI cluster projects in order for them to all use the same procurement framework.

6.1. Voucher Scheme

First successfully experimented by SixSq and Exoscale for the Human Brain Project, a voucher scheme is an effective mechanism to lower the barrier to entry to cloud service for new users. This was since generalised during the last phase of HNSciCloud.

[7] https://publications.europa.eu/en/web/eu-law-and-publications/publication-detail/-/publication/5253a1af-ee10-11e8-b690-01aa75ed71a1
[8] https://twitter.com/ESCAPE_EU

How does the Voucher Scheme work?

How the transaction is carried out, e.g. the payment model proper, is largely up to the entity funding the EOSC vouchers. Options could include pre-purchase of services, escrow of funds with investigators given withdrawal rights up to certain amounts, or distribution of credits via a pre-paid debit system. In addition, in this model, funding agencies or participating states could choose to directly fund the costs of maintaining/archiving key scientific data sets or other electronic resources in environments that scientists regularly select for research purposes. Moreover, this model foresees a dispute resolution scheme. For example, if a user pays with EOSC vouchers but does not obtain the quality of service foreseen (i.e. the service is 'down' for a period of time in breach of the Service Level Agreement) a mitigation is introduced.

The scheme enables community managers to distribute voucher codes to users, such that they can more easily access cloud resources. This type of feature is still evolving, but feedback is generally positive. With a limited quota associated to a voucher, users can easily deploy cloud resources, associated to a community or customer account, such that the procedure to access cloud funds is simplified.

The voucher scheme will be further refined and applied in the OCRE project.

7. Onboarding of commercial cloud providers in the EOSC: HNSciCloud recommendations

There are significant benefits the European Open Science Cloud (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.

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.

D6.3 Roadmap for the implementation of a full-scale European Open Science Cloud

Both contractors consider the European Open Science Cloud (EOSC) as a potential route to exploit the results of the HNSciCloud. In their commercialisation plans they both agree that several points need to be taken into consideration in the process of drafting the governance and the financial

mechanisms to support the Implementation Roadmap for the EOSC in order to ensure the onboarding of commercial cloud providers in the EOSC service catalogue. HNSciCloud Roadmap for the implementation of a full-scale European Open Science Cloud[9] provides a list of recommendations to encourage the engagement of commercial cloud providers in the EOSC. We report below few additional non-technical points that have been considered relevant also by the two Contractors RHEA and T-Systems:

- 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.
- A channel is required by which commercial cloud service providers can provide their input to the EOSC governance structure.

8. Conclusions

T-Systems and RHEA Group have different profiles, T-Systems being a very large international company while the RHEA Group less than 500 people, and offer different approaches to the cloud solution (T-Systems high-value cloud service provider vs RHEA Group's multi-cloud brokerage model). For these reasons, the commitment of the two contractors' towards the uptake of HNSciCloud hybrid cloud platform is different.

8.1. Conclusions from RHEA

The RHEA-led HNSciCloud consortium is committed to commercially exploit the service resulting from the project. The unique selling proposition of RHEA's service is based on a multi-cloud solution, able to scale to the future needs of the European scientific community. Further, RHEA service recognises and embraces the diversity this community represents. To this effect, RHEA setup a set of flexible offers and corresponding pricing models. These include for example: unified monitoring and metering (with flexible user interface and API), brokerage (with a unique contract to access all clouds), self- provisioning (bring your own cloud contract and benefit from

[9] Deliverable D6.3 Roadmap for the implementation of a full-scale European Open Science Cloud

the Nuvla service to better manage users and resources) and managed-services (e.g. data management service).

However, the analysis has shown that, on its own, it will not provide the volume to commercial cloud providers to provide a sufficiently attractive business; other existing funds need to be redirected to the use of commercial providers.

Despite the initiatives by the EC with HNSciCloud, OCRE and EOSC, the real market for science usage for commercial clouds it still to be established. A market needs both sellers and buyers who are both willing to trade. Further market research efforts are needed in this respect to cement the work already done and demonstrate to the commercial sector that, under the right conditions, the science community will buy, while convincing the science community that it is advantageous for them to do so.

Success tends to breed success. A special effort should therefore be made to attract early adopters, which could spearhead the next generation of hybrid-cloud and EOSC users, and in return provide these communities with new means in terms of capacity and access to compute, storage, networking resources, and much more, to produce new science. **To this effect, the RHEA consortium has introduced a voucher scheme, such that community managers can more easily enrol future cloud users, with limited process complexity.**

Since all HNSciCloud stakeholders will benefit from the success of the project and its exploitation, we should ensure we pull together all of our resources to disseminate and evangelise the current and future potential of the service, such that we together reap the fruits of collective success.

Both OCRE and EOSC are the most natural routes to exploitation of the results of the HNSciCloud PCP project. We are committed to pursue these avenues.

D-PIL-3.12 RHEA Commercialisation Plans

8.2. Conclusions from T-Systems

T-Systems has outlined the key objectives to be achieved and the relevant technical, functional and service aspects to focus on. Success can only be facilitated, however, if demand is also managed appropriately.

A key challenge to growing demand is to enable public procurement to easily use commercial cloud services as a supplementary or alternative infrastructure. The exchange with the Buyers Group on these aspects has been highly appreciated. **The resulting recommendation is to focus on implementing the appropriate procurement models and to support these with the required technical, functional and service options.**

T-Systems will invest a significant effort to extend the Hybrid Cloud solution with new functionality for procurement and service payment, to provide maximum incentive for the take-up of state-of-the-art services in research activities, e.g., through:

- New scientific Open Elastic Price List, eligible to science users and organisations only
- New volume discount scheme to enable attractive procurements based on financial volume, while maintaining maximum flexibility with pay-per-use billing
- New "Budget" function, for funding entities to self-service and manage financial quota, that can be linked to technical quota
- New "Cloud Credit" function, for funding entities to self-service and manage distribution of quota and provide users with services free.at-the-point-of-use.

8.3. Conclusions from the Buyers' Group

In November 2018 the 10 members of the HNSciCloud Buyers' Group (BG) completed a survey to gather their feedback on what would be required to convince them to pay for the services developed in the project. Specifically, considering the high-priority R&D activities and the experience with the use-cases that were deployed during the pilot phase, the survey addressed the Buyers' opinion about Compute and Storage services, Network Connectivity and Federated Identity Management services and service payment models. The primary aim of the survey was to understand whether they considered the hybrid cloud model as appropriate for their future needs.

The majority of the Buyers consider the hybrid cloud model to be appropriate for their future needs. That said, the buyers have very high expectations for commercial cloud service providers that plan to enter the research market. In particular:

- For what concerns compute and storage services, minor improvements are considered necessary by the Buyers' Group in the provision of VM configuration, GPUs and support to batch systems, while significant work still needs to simplify the provisioning of HPC resources. In terms of support provided by the contractors to the BG, documentation, operational security and compliance to GDPR were considered overall satisfying, while training should be improved in the future, with hands-on tutorials as the preferred training style.
- While major to significant improvements are perceived as needed to the Federated Identity Management services, in terms of network connectivity, half of the buyers would continue to connect via GEANT, but another 30% would prefer to procure their own connection.
- Finally, in terms of service payment models, the buyers would suggest significant improvements to the financial dashboards available to manage billing and accounting, while the use of vouchers is seen as positive by the BG.

An interesting outcome applies to the perception of how the cost of the contractors' services compare to in-house provisioned services. The contractors' services are considered more expensive than in-house costs by half of the respondents, while 40% of them don't know the costs that their institutes are facing for in-house services. This matches with the perception of the contractors who believe that despite the fact that many scientists, for the first time, were able to access significant IaaS resources on-demand in such a short time span, the members of the BG did not fully appreciate the value of this transformation due to the engagement of commercial cloud services via HNSciCloud.

Annex: Assessment of what is needed for the future: HNSciCloud Buyers' Group feedback on the Contractors' services

The results of the survey performed in November 2018 with 10 members of the Buyers' Group are provided below.

Compute and Storage

- CPUs: any VM configurations missing (e.g. large memory etc.)
- HPC: OTC – do they need to support 'hybrid model' so that one can launch a cluster from a client machine?
- GPUs: testing and development using new architectures is seen as a benefit
- Support for batch systems (HTCondor, SLURM)
- Storage: is OneData sufficiently stable to be used without the intervention/support of Cyfronet?
- Support for the services: interfaces (email vs. tickets, phone etc.), coverage (24h or 8hr, weekends), response time
- Training: What sort of training would be needed in a production environment? What style suits best your needs –
- Hands-on tutorials, presentations etc?
- Documentation: what is still missing?
- Operational security challenge – are the results sufficient?
- Provisions for GDPR sufficiently demonstrated?
- Do you see a need for a multi-cloud provider brokering service (technical brokering - interfaces, financial brokering – contract management etc)?

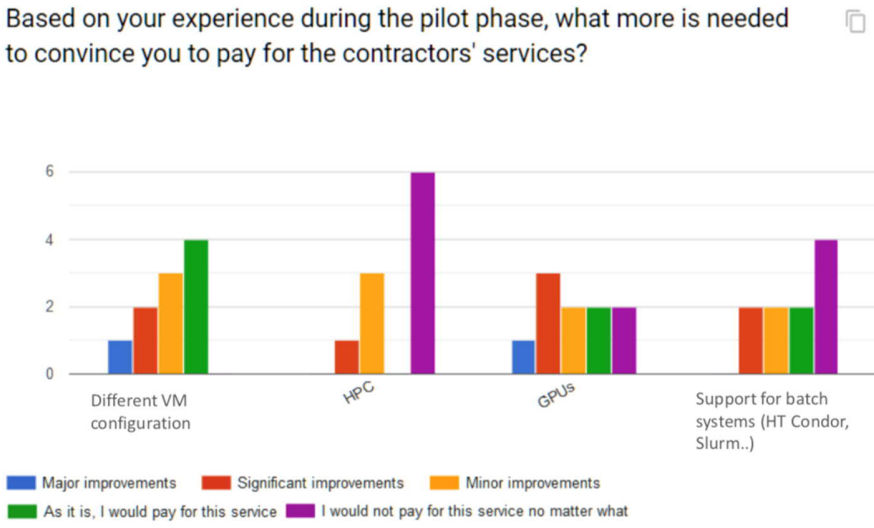

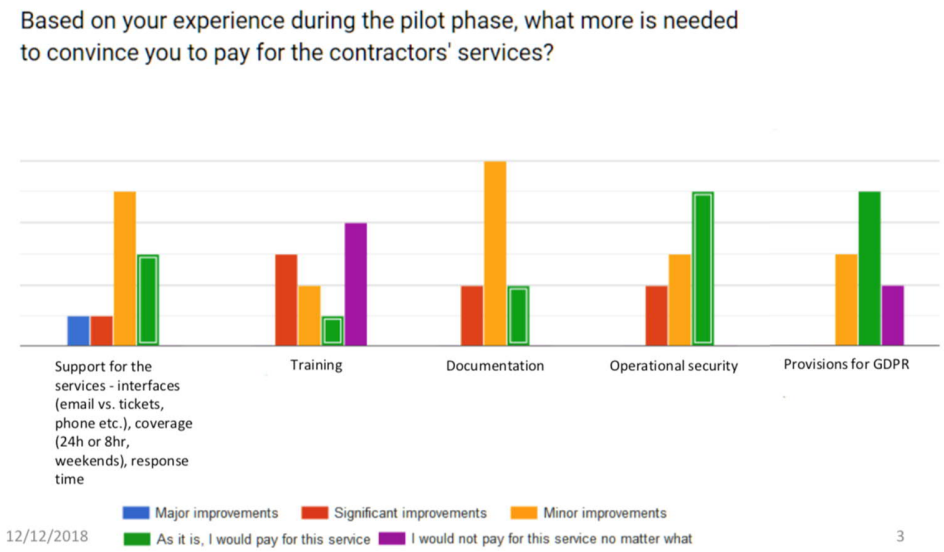

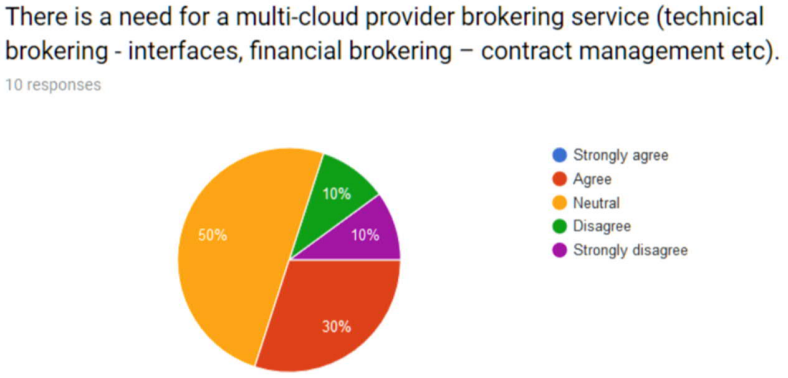

Network Connectivity and Federated Identity Management

Is the basic eduGAIN compatibility sufficient or are additional features, such as token translators (e.g. keycloak) necessary?

- Is the AAI sufficiently integrated across all the services (compute, data, Onedata, dashboard etc)?
- Support for non-web applications?
- Network connectivity: taking into account the pricing of connections given in the TCO reports and Quality of Service observed, would you continue to connect via GEANT or procure your own connection?

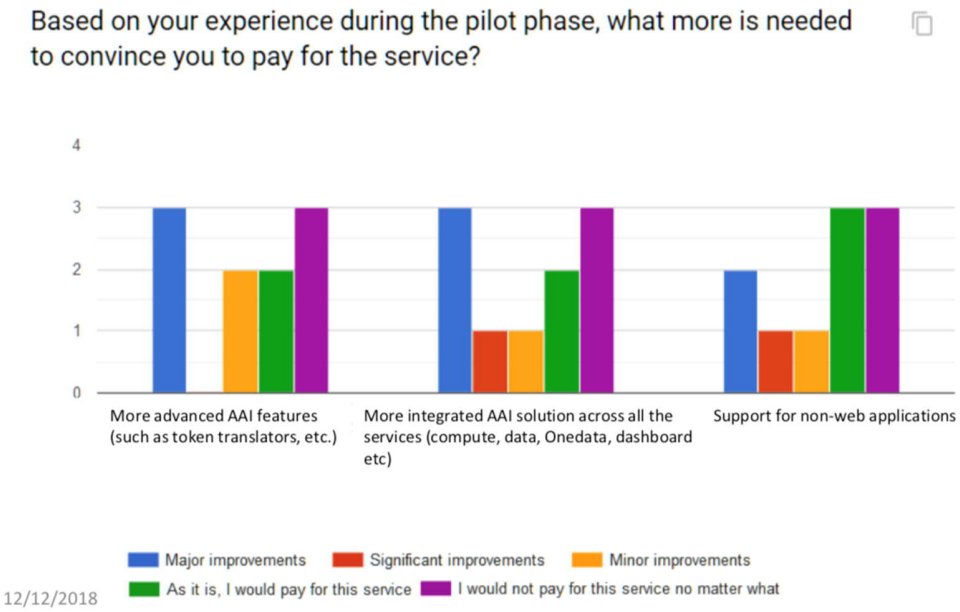

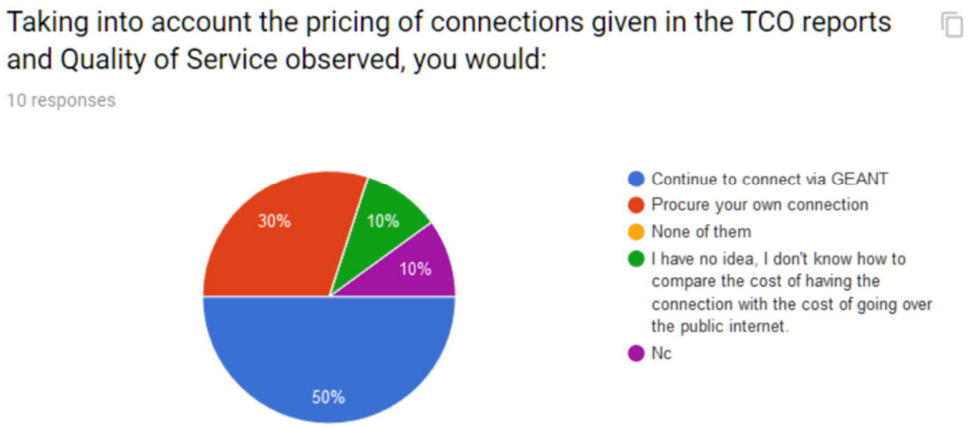

Service Payment Models

- Billing and accounting
  - Reporting and accounting (i.e. financial dashboards)
    - Does the financial dashboard to cover all types of billed services in the pilot platform?
    - Is the usage information Up to date and accurate?
    - Is the accounting history accessible?
  - Does quota management for sub-tenants/projects work?
  - Reserved instances vs pay as you go (on-demand) or pre-emptible (spot-instance)?
  - Duration of contracts(1 or 3 years)?
  - How many IaaS providers would you work with (0, 1, 2, more?)

- Use of vouchers – under what conditions would you use them and for which types of users and use-cases?
- TCO: How does the cost of the contractors' services compare to your in-house provisioned services?

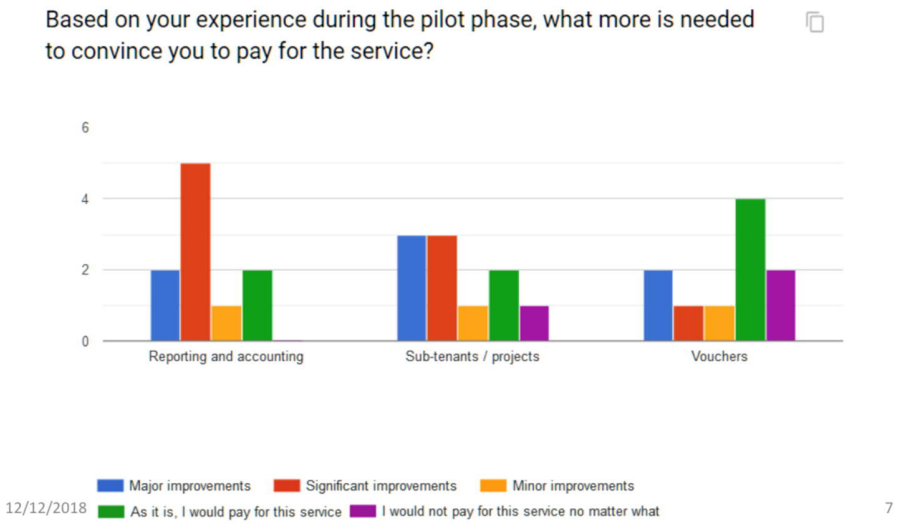

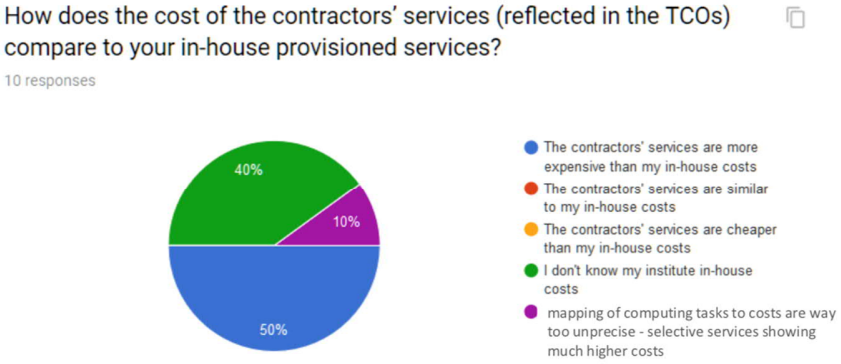

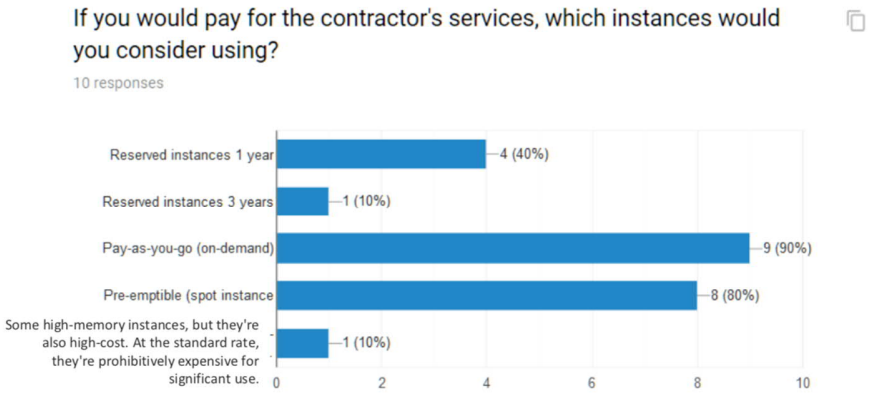

And finally, do you consider the hybrid cloud model as appropriate for your future needs?

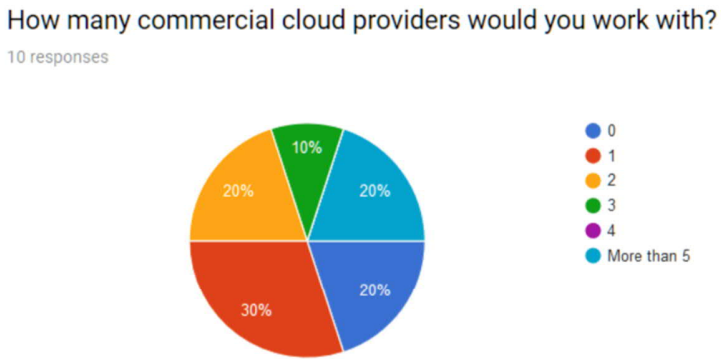

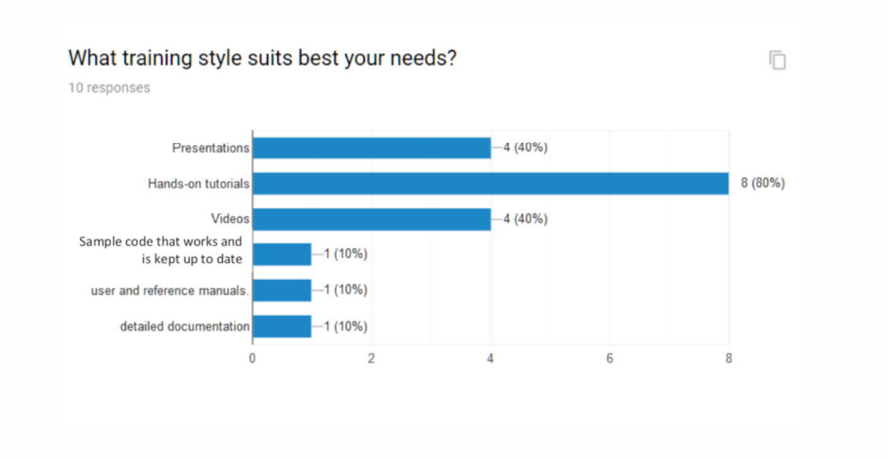

Do you consider the hybrid cloud model as appropriate for your future needs?

10 responses