

HNSciCloud, a Hybrid Cloud for Science

João Fernandes^{1,*}, Bob Jones¹, Sergey Yakubov², Andrea Chierici³

¹CERN, IT Department, 1211 Geneva, Switzerland

²Desy, Notkestrasse 85, 22607 Hamburg, Germany

³INFN-CNAF, Viale Carlo Berti Pichat 6, 40127 Bologna, Italy

Abstract. Helix Nebula Science Cloud (HNSciCloud) has developed a hybrid cloud platform that links together commercial cloud service providers and research organizations' in-house IT resources via the GÉANT network. The platform offers data management capabilities with transparent data access where applications can be deployed with no modifications on both sides of the hybrid cloud and with compute services accessible via eduGAIN [1] and ELIXIR [2] federated identity and access management systems. In addition, it provides support services, account management facilities, full documentation and training. The cloud services are being tested by a group of 10 research organisations from across Europe [3], against the needs of use-cases from seven ESFRI infrastructures [4]. The capacity procured by ten research organisations from the commercial cloud service providers to support these use-cases during 2018 exceeds twenty thousand cores and two petabytes of storage with a network bandwidth of 40Gbps. 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. An early adopter scheme has been put in place so that more research organisations can connect to the platform and procure additional capacity to support their research programmes.

1 The HNSciCloud PCP

The Helix Nebula Science Cloud (HNSciCloud) is a €5.3 million Pre-Commercial Procurement (PCP) [5] tender for the establishment of a European hybrid cloud platform to support the deployment of high-performance computing and big data capabilities for scientific research. HNSciCloud is sponsored by ten of Europe's leading public research organisations and co-funded by the European Commission. The buyers group is formed by CERN (main procurer and project coordinator), CNRS, DESY, EMBL-EBI, ESRF, IFAE, INFN, KIT and SURFsara, sponsoring use-cases in several scientific domains.

The HNSciCloud PCP tender covers the procurement of R&D services for the design, prototype development and pilot use of innovative cloud services.



Fig. 1. Scientific domains represented in HNSciCloud.

* João Fernandes: joao.fernandes@cern.ch.

The innovative cloud services provided by the platform are designed and implemented to address a set of challenges that require a combination of services at the Infrastructure as a Service (IaaS) level. These challenges are:

- **Compute and Storage:** support a range of virtual machines and container configurations working with datasets up to the petabyte range.
- **Network Connectivity and Federated Identity Management:** provide high- end network capacity for the whole platform with common identity and access management.
- **Service Payment Models** - explore a range of purchasing options to determine the most appropriate ones for the scientific application workloads that will be deployed.

The procured cloud services integrated with the procurers’ in-house resources and publicly funded e-Infrastructures provide a hybrid platform that serve scientists and engineers working in high-energy physics, astronomy, life sciences including biomedical research and the photon/neutron sciences, as depicted in Figure 1.

2 Project Phases

HNSciCloud is organized in four different phases, as depicted in Figure 2. After a phase of requirement assessment that corresponds to the tender phase, the providers enter three R&D stages: Design, Prototype and Pilot, each of them competitive for a duration of three years.

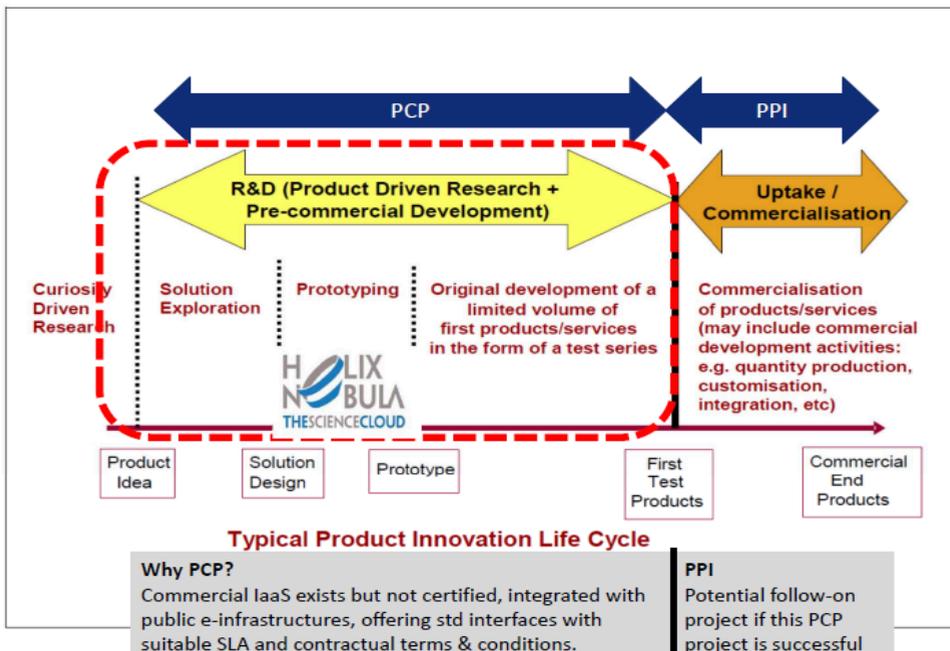


Fig. 2. HNSciCloud project phases.

The phases and awarded consortia are described in more detail in the next sections.

2.1 Tender phase

In March 2016, HNSciCloud organised an Open Market Consultation (OMC) event to analyse and discuss across more than seventy cloud provider representatives the requirements

and use-cases. The OMC conclusions and findings set the basis for the preparation of the tender that was launched in July 2016 and closed in September 2016.

A total of thirty multinational companies, SMEs and public research organisations from thirteen countries submitted bids during the period. In November 2016, four consortia were awarded contracts at the tender award ceremony at CNRS in Lyon, France.

Overall, the tender generated a high-level of interest in the cloud services market, where a significant number of FP7 [6] and H2020 [7] projects were cited in the bids as sources of innovation contributing to the proposed solutions. Four consortia were awarded contracts for the Design phase, as depicted in Figure 3.

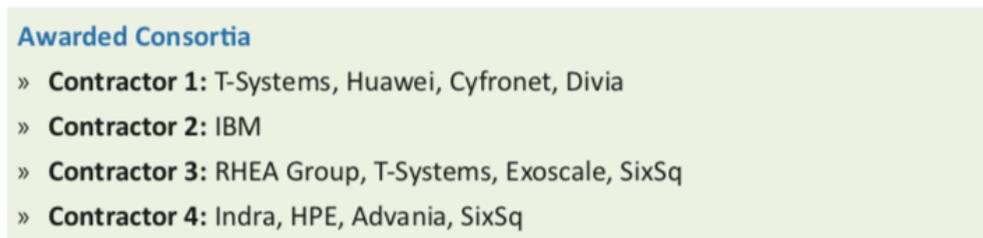


Fig. 3. Awarded Consortia of HNSciCloud for the Design phase.

2.2 Design phase

The Design phase of HNSciCloud kicked-off during the tender award ceremony. This phase had 15% of the project budget allocated. In February 2017, the four consortia delivered their designs (including architecture, technical design of components together with unit costs). The eligible contractors submitted their Prototype phase bids in early March and the HNSciCloud evaluation committee assessed the bids and selected the three most promising. The three winning Consortia moving to the HNSciCloud Prototype phase are shown in Figure 4. The announcement took place on 3rd April 2017 during the awards ceremony held at CERN in Geneva, Switzerland.

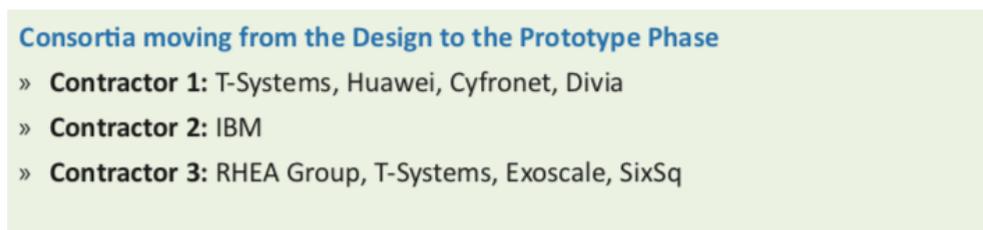


Fig. 4. Awarded Consortia of HNSciCloud for the Prototype phase.

There were a number of key aspects that scientific organisations need to take into account for the design phase of a PCP project, in the cloud domain:

- Prepare detailed information on each use-case.
- Start the dialogue between the scientific organisations and the cloud providers as early as possible and continue discussions during the whole phase.
- Set intermediate milestones for each contractor to present their progress.
- Revise the effort planning when the Request for Tender is published.
- Establish a dedicated technical team to collect input from all procurers, interact with the contractors and actively perform tests.

2.3 Prototype phase

During the Prototype development phase, which kicked-off in April 2017, the cloud providers developed prototypes and made them accessible to experts from the scientific organisations for testing.

In order to ensure a successful outcome of this phase, the buyers group closely monitored the providers activities and performed extensive testing of the prototype solutions. In order to facilitate the information flow between the supply and demand sides, progress review events were held between putting together contractors, researchers and IT experts.

The Prototype phase has produced results that address the major challenges defined for the project. Those challenges included compute and storage supported by a range of virtual machines and container configurations for PB range datasets, network connectivity with high-end network capacity via GÉANT [8], and Federated Identity Management using the SAML2 protocol supporting both eduGAIN and ELIXIR AAI.

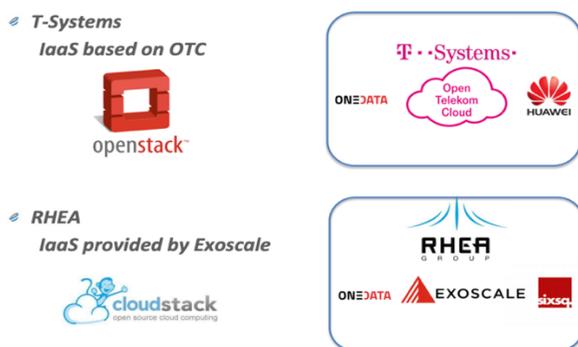


Fig. 5. Awarded Consortia of HNSciCloud for the pilot phase.

2.4 Pilot Phase

During the Pilot phase, which started in February 2018, contractors expanded their prototypes. The selected consortia are shown in Figure 5. These prototypes were made available first to IT experts who performed scalability tests and then, to end-users so that they could deploy their applications. The assessment of each of the pilots provided valuable feedback to the contractors.

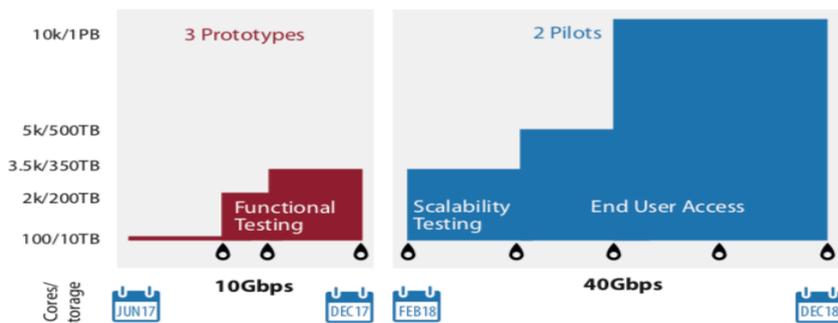


Fig. 6. Cloud capacity during prototype and pilot phases of HNSciCloud.

The Pilot phase corresponded to the real deployment of scientific use-cases foreseen in the beginning of the project, after a scalability testing period. This phased deployment is illustrated in Figure 6.

3 HNSciCloud Use-Cases

The group of research organisations participating in HNSciCloud deployed several use-cases in the pilot services. The IaaS services were used across four scientific domains, comparing the performance and costs of diverse applications. The following sections provide details about a number of flagship use-cases.

3.1 Astronomy - LOFAR ASTRON

The LOFAR (Low Frequency Array) is the first of a new kind of telescope that uses an array of simple omnidirectional antennas as opposed to a dish antenna for mechanical signal processing. The electronic signals obtained from the antennas are digitized and transported to a central digital processor. The antennas are simple enough, but there is a large number of them – about 7,000 in the full LOFAR design. The main goal of the LOFAR use case is to test and subsequently turn into production, scattered storage and computing resources distributed across Europe. The Helix Nebula Science Cloud grants this opportunity due to the availability of the highspeed network connection of the GÉANT Cloud VRF infrastructure, and the pricing model based on compute and storage alone.

3.2 Photon/Neutron Sciences - CrystFEL Serial Femtosecond Crystallography

The CrystFEL framework is used for the technique of Serial Femtosecond Crystallography (SFX) and comprises programmes for data processing, simulation and visualization. CrystFEL is part of a complex, non-distributable software stack, which is free for use by academia and non-profit organizations. The CrystFEL framework is increasingly used at various synchrotrons to analyze data from serial (femto-second) x-ray crystallography. The nature of these experiments makes a cloud-based distributed pipeline particularly appealing, as the framework can fully exploit large computational resources with tunable demands. The objective of the CrystFEL use case is to run ‘medium’ data intensive data analysis tasks.

3.3 Life Sciences - Large scale genomics analyses in cancer studies

The Pan-Cancer initiative has the objective of comparing 12 tumor types profiled in the context of The Cancer Genome Atlas [9] research network. Cancer can take hundreds of different forms, depending on external factors such as localization and cell type. Pan-Cancer currently represents the most comprehensive computational study dealing with cancer genomics, with roughly 1 PB of data to be processed. This has forced researchers to implement new pipelines able to cope with the massive quantity of data, with a focus on leveraging cloud resources provided by public and commercial clouds, including the EMBL-EBI Embassy Cloud [10]. In the context of HNSciCloud, the Pan-Cancer project was able to determine genetic variation for more than 5000 tumor samples with more coming in on a monthly basis.

3.4 High Energy Physics - On Demand Analysis Services for CMS

CERN does not have the computing or financial resources to crunch all of the data on site.

Recent developments in cloud computing have come to public attention due to the promise of providing as much computing power as users need, simplifying management and reducing Total Cost of Ownership (TCO). CMS is one of the LHC experiments in the Worldwide LHC Computing Grid (WLCG) collaboration, linking up 170 computing centres in 42 countries. Demonstrating that HNSciCloud can satisfy the requirements of LHC experiments has an enormous potential impact on cloud adoption in the particle physics communities. DODAS [11] has been a successful deployment in HNSciCloud, allowing CMS to utilize “any cloud provider” to generate sites on demand, with almost zero effort.

3.5 Relevance of the resulting services

The assessment of the relevance of three categories of services offered in the context of the project that was done by the buyer organisations users is shown in Figure 7. The service categories in the left are Federated AAI, Storage and Compute. The deployed use-cases are FDMNES [12], CrystFEL, DODAS, LOFAR, PanCancer and WeNMR/HADDOCK. The relevance was assessed with three levels: useful (one dot), relevant (two dots) and fundamental (three dots).

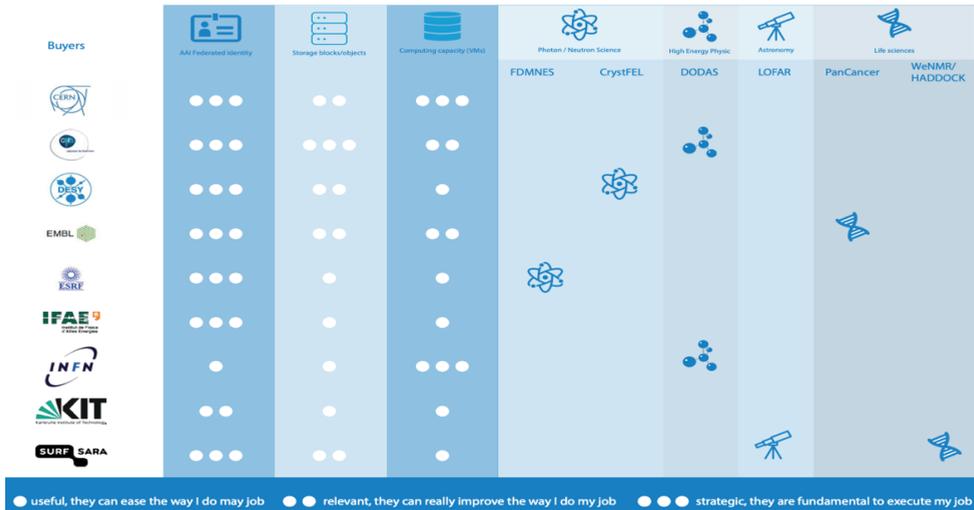


Fig. 7. Assessment of relevance of HNSciCloud for the scientific use-cases deployed.

In addition, HNSciCloud has proven to offer strategic opportunities to rapidly deploy use-cases across architectures not available on-premises at scale such as GPU accelerators, avoiding the burden of internal hardware procurements for R&D purposes. A very successful example has been the deployment of Deep Learning (DL) workloads for fast detector simulation [13]. The deployment of cutting-edge DL workloads in HNSciCloud allowed to obtain and publish results in a period of less than six months. Such an aggressive timeline would not be possible when recurring to the traditional hardware procurement processes currently available for public organisations.

4 Conclusions and future directions

There are a number of lessons learned based on the experience gathered in the several phases of HNSciCloud. The main findings are summarized in the following sections.

4.1 Preparation and execution phases

During the tender preparation phase, corresponding to a nine to twelve months period, efforts are needed to precisely define the R&D challenges, the objectives and expected outcome of the PCP project. This includes aspects such as an in-depth needs assessment and an open market consultation evolutionary process during the tender preparation phase. Concretely, organize a series of events if possible, where the procurers and potential tenderers can progressively refine the focus under the guidance of the scientific experts. Another aspect to be considered is to launch a survey among the known market players to allow the procurers to detect the capabilities and the willingness of the market to participate in the tender.

Concerning the tender process, it was found that nominating a lead procurer that already had longstanding relationships with all members of the buyers group proved to be a successful approach. The close cooperation between the members of the buyers group was essential to the success of the project.

In order to minimize potential conflicts of interest during the project execution, the tender text should have a restriction forbidding a company from being both a lead contractor and a participant in other bidding consortia. In addition, no more than one tender should be allowed to be submitted by a natural or legal entity. Provisions in the selection criteria must be included to ensure the proposed solutions by competing consortia are sufficiently distinct and multiple solutions for each PCP challenge are developed. Future projects should strive to contract a higher number of consortia during each phase to ensure both sufficient competition and an increased likelihood of a successful completion.

During the execution phase, it was found to be important to continuously monitor R&D priorities taking into account the project challenges, communicate clearly these priorities to the contractors and take them into account when allocating resources during execution.

The concept of an early adopter programme should be promoted from the very beginning of the PCP, during the planning to allow sufficient time for other organisations to benefit from the resulting services.

The testing process can clearly benefit from the development and execution of a test suite as a separate task with its own milestones, deliverables and funded effort, complemented with a centralised repository to be shared across the buyers group and contractors.

The call-off for the pilot phase must include progressive ramping-up of capacity of IaaS resources in order to contain costs of IaaS provisioning. The management teams need to regularly monitor the consumption of resources and make adjustments accordingly, in order to control costs and ensure all the testing and deployment objectives can be achieved. A formal compensation scheme must also be in place in the tender in case of service outages during the project lifetime.

The schedule of intermediate reviews and partial financial payments are important checkpoints for the buyers group and contractors, to ensure that all parties remain active and engaged during the several phase cycles.

4.2 Cost assessment of the resulting services

The Total Cost of Ownership (TCO) study [14] for selected use-cases was introduced in the pilot phase to help the buyers group understand the impact of the commercialisation plans for their organisations. In future projects, the inclusion of a TCO study should be considered as an award criterion for each execution phase (design, prototype and pilot) to have costs of resulting services intimately linked to the architecture designs and R&D performed. Specialists in reviewing the tender specifications should be considered to ensure compatibility with market offerings, to estimate the current market costs of the capacity requested and to determine if the resulting services are competitive when compared to market prices. In order to turn TCO studies in a useful tool for service provisioning decisions, cost

transparency needs to be increased for in-house provisioned services by research procuring organisations in the public sector.

4.3 Access and wider adoption

HNSciCloud adopted a voucher/credit scheme to lower the entry barriers to cloud services for new users. The scheme enabled community managers to distribute voucher codes to users, such that they could more easily access cloud resources.

A limited quota associated to a voucher, allowing users to easily deploy cloud resources, associated to a community account. The buyer organisations oversighted the usage of the funded services. The scheme is being expanded and developed further in future initiatives.

The innovative services have been registered in the European Open Science Cloud (EOSC) [15] catalog, promoting adoption by other user groups. The process developed in HNSciCloud is considered a working example of the EOSC in engaging commercial cloud services for the scientific community.

References

- [1] eduGAIN: <https://edugain.org/>
- [2] ELIXIR AAI: <https://www.elixir-europe.org/services/compute/aa1>
- [3] HNSciCloud Buyers Group: <http://www.hnscicloud.eu/partner-type/buyers>
- [4] ESFRI infrastructures: <http://www.hnscicloud.eu/hnscicloud-user-groups>
- [5] PCP process: <https://ec.europa.eu/digital-single-market/en/pre-commercial-procurement>
- [6] 7th EC Framework Programme: https://ec.europa.eu/research/fp7/index_en.cfm
- [7] Horizon2020: <https://ec.europa.eu/programmes/horizon2020/en>
- [8] GÉANT network: <https://www.geant.org/Networks>
- [9] The Cancer Genome ATLAS: <https://cancergenome.nih.gov/>
- [10] EMBL-EBI Embassy cloud: <https://www.embassycloud.org/>
- [11] D. Spiga, *CMS Dynamic On Demand Analysis Service*, Proceedings of this conference (CHEP 2018)
- [12] FDMNES: <http://neel.cnrs.fr/spip.php?rubrique1007>
- [13] S. Vallecorsa, *A Machine Learning tool for fast simulation*, Proceedings of this conference (CHEP 2018)
- [14] M. Devouassoux, *Method to Calculate the Total Cost of Ownership of Infrastructure as a Service*?, DOI: [10.5281/zenodo.1463981](https://doi.org/10.5281/zenodo.1463981) (2018)
- [15] EC Directorate-General of Research and Innovation, *Prompting an EOSC in practice*, DOI: [10.2777/112658](https://doi.org/10.2777/112658) (2018)