# HNSciCloud - Overview and technical Challenges

IOP Conf. Series: Journal of Physics: Conf. Series **898** (2017) 052040 doi:10.1088/1742-6596/898/5/052040

# HNSciCloud - Overview and technical Challenges

**Martin Gasthuber**[1], **Helge Meinhard**[2] **and Robert Jones**[2]

[1] DESY, Hamburg/Germany

[2] CERN, Geneva/Switzerland

**Abstract.** HEP is only one of many sciences with sharply increasing compute requirements that cannot be met by profiting from Moore's law alone. Commercial clouds potentially allow for realising larger economies of scale. While some small-scale experience requiring dedicated effort has been collected, public cloud resources have not been integrated yet with the standard workflows of science organisations in their private data centres; in addition, European science has not ramped up to significant scale yet. The HELIX NEBULA Science Cloud project - HNSciCloud, partly funded by the European Commission, addresses these points. Ten organisations under CERN's leadership, covering particle physics, bioinformatics, photon science and other sciences, have joined to procure public cloud resources as well as dedicated development efforts towards this integration. The HNSciCloud project faces the challenge to accelerate developments performed by the selected commercial providers. In order to guarantee cost efficient usage of IaaS resources across a wide range of scientific communities, the technical requirements had to be carefully constructed. With respect to current IaaS offerings, data-intensive science is the biggest challenge; other points that need to be addressed concern identity federations, network connectivity and how to match business practices of large IaaS providers with those of public research organisations.

In the first section, this paper will give an overview of the project and explain the findings so far. The last section will explain the key points of the technical requirements and present first results of the experience of the procurers with the services in comparison to their 'on-premise' infrastructure.

## 1. Introduction

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 brought together cloud infrastructure providers and large scientific user communities to overcome barriers to adoption of Infrastructure as a Service (IaaS) for scientific use. In addition to offering the production Helix Nebula marketplace platform (HNX), the initiative has undertaken the first joint Pre Commercial Procurement (PCP) tender called Helix Nebula Science Cloud (HNSciCloud) [1]. This 5.3 million (Euro) joint tender, led by CERN, will establish a European hybrid cloud platform that will support high-performance, data-intensive scientific use-cases sponsored by 10 of Europe's leading public research organisations (CERN, CNRS, DESY, EMBL-EBI, ESRF, IFAE, INFN, KIT, STFC, SURFSara) and co-funded by the European Commission. Figure 1 shows the overall timeline with the three phases separating the fundamental different goals - design, functionality, scaling.

IOP Conf. Series: Journal of Physics: Conf. Series **898** (2017) 052040 doi:10.1088/1742-6596/898/5/052040

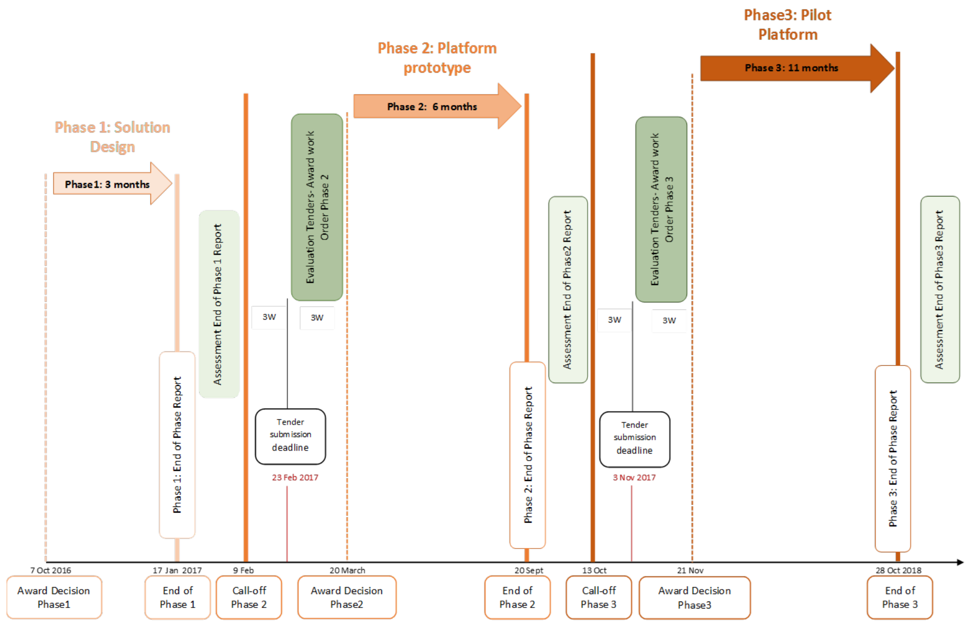


**Figure 1.** HNSciCloud - the fundamental 3 phases of this PCP project

The procured cloud services are being integrated with the procurers in-house resources and publicly funded e-Infrastructures to provide a hybrid platform for end-users from a wide range of scientific fields including high energy physics, life sciences, astronomy, neutron/photon sciences and the long tail of science. The set of use-cases that will be supported by the procurement include those directly connected to 7 Research Infrastructures that appear in the ESFRI 2016 roadmap, namely

- Euro-BioImaging: European Research Infrastructure for Imaging Technologies in Biological and Biomedical Sciences
- CTA: Cherenkov Telescope Array
- BBMRI: Biobanking and BioMolecular resources Research Infrastructure
- ELIXIR: A distributed infrastructure for life-science information
- ESRF Upgrades: Extremely Brilliant Source
- European-XFEL: European X-Ray Free-Electron Laser Facility
- HL-LHC: High-Luminosity Large Hadron Collider

The HNSciCloud tender was published during 2016 and the adjudication process has completed with contracts awarded for the design phase to four consortia:

- T-Systems, Huawei, Cyfronet, Divia
- IBM
- RHEA Group, T-Systems, exoscale, SixSq
- INDRA Systems, HP, Advania, SixSq

The procured services will be competitively evaluated during prototype and pilot phases with the end-users through 2017 and 2018 after which the most successful services will be commercialised and made generally available.

Unlike 'classical' EU projects, the representatives of the scientific communities are not carrying out the R&D work - this work is part of the work carried out by the contractor during

IOP Conf. Series: Journal of Physics: Conf. Series **898** (2017) 052040 doi:10.1088/1742-6596/898/5/052040

the three phases of the project. This construct should lead to industry standard service offerings with a competing market.

## 2. Motivation

All current and planned large scale scientific projects have demands for huge data storage and analysis resources. Upcoming projects like the HL-LHC and in the domains of Genome Analysis, Astrophysics, Life Science and Photon Science have requirements for their resources to scale massively, being able to share the data and to some extent being available on demand. All of this being deployable in potential world wide scientific collaborations including the 'long tail of science'. All represented scientific communities (the buyers group) expect benefits from hybrid cloud solutions especially for:

- economy of scale - the costs for public cloud resources have demonstrated significant reduction within the recent years. It is expected that this trend continues.
- more elasticity - adaptivity to changing demands for faster and more efficient scaling.
- more concentration on science-specific services/demands not covered by public cloud providers
- standardization - make computational and storage infrastructure changeable and replaceable - benefit from real market for standardized resource providers

### *2.1. Limitations*

The overall assumption for the execution of the project was that the basic IaaS services, widely available today, are close to the requirements of the represented scientific communities. In order to fill the small, but existing, gap and to prove scaling of the solution, the PCP has an R&D part and a service part. According to the PCP rules the R&D part has to be at least 50% of the overall budget. To balance the 'innovation' effort against the service (real resources) is not easy and led to compromises in a few areas.

## 3. Technical Challenges

The project starts with some identified technical challenges, mostly derived from the demand to build a hybrid cloud and second, to easily integrate existing scientific applications without any changes. The initial call, for the design phase, expresses the simple statement that scientific applications should run without changes on the hybrid cloud infrastructure. This implies, especially for the data access, more than just simple IaaS services and a closer connection between the storage systems at the buyers side, to those deployed at the public cloud providers. The 'seamless' integration also requires an integration in other areas, i.e. authentication and procurement.

### *3.1. Compute - HPCaaS*

Apart from the standard IaaS compute services (i.e. standard Intel based two socket machines) which cover the magnitude of the expected scientific workload, a growing demand for HPCaaS like resources exists. For specialized (i.e. new accelerator technology) simulations, classical HPC resources are required - several 10K cores, high bandwidth/low latency interconnect (i.e. InfiniBand, OmniPath) and multiple TB of memory. Beside that, communities like photon and neutron science are asking for HPC component based architectures to allow efficient data analysis. The analysis core runs parallel (MPI enabled), but requires only a few 100 cores. The connection to on-site storage should allow for several GB/sec for a single stream and application. In both cases, the usage of cpu accelerator technologies, like GPU and FPGA, are possible and welcome.

IOP Conf. Series: Journal of Physics: Conf. Series **898** (2017) 052040 doi:10.1088/1742-6596/898/5/052040

### *3.2. Containers*

For this project, the buyers group is in the active process of learning and understanding the benefits from container technology. Although our workloads are not based on 'micro services' - as the main origin motivations for container development - the expected benefits are covering the following areas:

- containment - in scientific communities the user of software is mostly also the developer. Containers should allow for perfect separation of work and responsibilities between developer, local service provider (IT group) and public resource providers. The achieved 'runtime containment' allows a maximum of independence for the development and, at the same time, the required standardization and security for the resource providers.
- low overhead - especially for HPC like resources or high bandwidth storage connection, the overhead is lower compared to virtualization.
- better scaling and interfacing to local scheduler. Shipping containers instead of jobs, allow clients to interface with services at a higher level (higher than IaaS) in a standardized way.

### *3.3. Data Access*

While the project proposal was being set up, it became clear that all represented scientific communities depend on file based data access. The minimum set of functions derived from the POSIX standard are grouped around the basic access primitives, explicitly excluding sophisticated locking or parallel access controls. The requirements express that no modification of the scientific applications should be required and the connection to the storage systems at the buyers group should be minimally invasive. No further details have been requested in the call, aiming for a wide range of technical solutions. Two modes of operation regarding the caching of data are envisaged - pre-fill (conditioning) of data copied from the source, or data fetching on demand. The first will allow further optimizations once the datasets are known prior to execution, the second allows ad-hoc job scheduling by fetching data on demand.

### *3.4. eduGAIN/ELIXIR integration, SAML2*

To allow integration of individual users accessing public resources, these resources must be 'embraced' with the same federated identity service as the buyers resources have in place. Although the scientific community has not (yet) finalized on a single identity service, the call asked for eduGAIN integration, more generally any SAML2 based federated identity system (e.g. EMBL ELIXIR).

### *3.5. Network*

In order to demonstrate a hybrid cloud covering 5% of the overall scale, the project asked for a 10Gbit connection from the public cloud provider to GEANT (European science network) for the functional testing throughout the prototype phase. This will be scaled up to a 40Gbit connection for the pilot phase

### *3.6. semi technical*

Mapping the world of scientific computing, with all its cultural and different wording compared to the industrial world, poses quite a few cultural issues which need to be addressed throughout the project. Beside the different worlds of industry and academia, the different scientific communities represented in the buyers group, imposes similar challenges.

The rules and obligations of the project type - a Pre-Commercial-Procurement project (PCP), differ from normal procedures following the business models of the related cloud industry.

IOP Conf. Series: Journal of Physics: Conf. Series **898** (2017) 052040 doi:10.1088/1742-6596/898/5/052040

## 4. Summary and Outlook

The project has finalized phase 1 (design) and started phase 2 (prototype) which will include real integration tests on all buyers use-cases with all contractor supplied prototypes. The prototype phase will be the first time the public resources together with the buyers resources will form a hybrid configuration with location independent data access and authentication. First impressions from the supplied design documents show the transparent data access being the largest part of the R&D activities for all contractors.

## References

[1] The H2020 project Helix Nebula Science Cloud (HNSciCloud), Co-funded via H2020 as a Pre-Commercial Procurement (PCP) project: Grant Agreement 687614, http://www.hnscicloud.eu/
[2] The Helix Nebula Initiative (HNI) http://www.helix-nebula.eu/