Helix Nebula – The Science Cloud

Deliverable Title: D2.1_Report on the Open Market Consultation and the Results

Partner Responsible: CERN

Work Package: 2 – Task 2.1

Submission Due Date: M4 – April 2016

Actual Submission Date: 21.04.2016

Distribution: Public

Nature: Report

Abstract: Report on the market consultation and the results duly taken into account to define the tender specifications

Disclaimer

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

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

Grant Agreement Number: 687614
Start Date: 01 January 2016
Duration: 30 Months

Log Table

Issue	Date	Description	Author/Partner
V0.1	21/04/2016	Initial draft	Rachida Amsaghrou (CERN)
V0.2	17/05/2016	Input on formatting	Andrea Chierici (INFN)
V1.0	17/05/2016	Final Version	Rachida Amsaghrou (CERN)

Table of Contents

Executive Summary

This report highlights the results of the HNSciCloud Pre-Commercial Procurement (PCP) Open Market Consultation[1] (OMC), which took place on 17 March 2016 at CERN in Geneva, Switzerland.
The OMC was preceded by a preparatory workshop on 16 March during which procedures for the OMC were defined and agreed amongst the HNSciCloud collaboration board members present[2]. The preparatory morning session consisted of presentations of the procurement, legal and technical characteristics of the tender, while the afternoon session was dedicated to processing these characteristics into requirements, which were in turn grouped into use cases.
Task 2.1 "Open Market Consultation" aims at verifying the market readiness to meet the procurement needs through an open dialogue that broaches the views of the market about the intended scope. The results of this open market consultation have been duly taken into account to define the tender specifications, and filled the gap between state-of-the art industry developments and procurement needs, thus justifying the need to procure R&D services.
CERN, as project coordinator and leader of work package 2[3], applied through a management and technology consulting company[4], a unique methodology in order to identify the most appropriate use-cases and evaluation criteria to be used for this PCP project, based on an evaluation of risk (on the supply-side) and value (on the demand-side).

This consultancy resulted in the production of the following outputs:

- Main use cases prioritized
- Reference architecture for the prototype refined
- Value-Risk matrix mapped based on prioritized sources of risk
- Evaluation criteria for phases 1, 2 and 3 defined

The Open Market Consultation event brought together over 40 European cloud service providers. The main objective of the Pre-Commercial Procurement Open Market Consultation was to establish the premises of the European science cloud platform. The cloud platform will combine several sub-challenges requiring a combination of services at the IaaS[5] level, integrated into an environment supporting the full life-cycle of science work-flows.

[1] https://indico.cern.ch/e/HNSciCloud_OMC
[2] https://indico.cern.ch/e/OMC_Prep_Day
[3] Preparation of Procurement and Tendering
[4] http://www.addestino.be/
[5] Information As A Service

Helix Nebula - The Science Cloud PCP Open Market Consultation

All information contained herein are for discussion purposes only and shall not be considered a commitment on the part of CERN or the Buyers group.

1. Introduction

Context

This report summarizes the results of an open market consultation for the **Pre-Commercial Procurement** (PCP) of the **Helix Nebula Science Cloud project** (HNSciCloud).

The goal of the HNSciCloud project is to create a **hybrid cloud platform** that can be used to deploy a **massive** amount of applications to perform **high performance computing** on **large datasets**.

Prior to the open market consultation, a workshop was organized on **March 16th** from 13h till 16h30 at the CERN headquarters in Geneva. During this moderated workshop, procurers analyzed and discussed the different requirements of the project. From these requirements, different **use cases were identified and prioritized** based on their **value for the end-user**. These use cases are discussed in more detail in section 2. An overview of the procurers that attended the workshop can be found in Appendix 1.

The open market consultation itself took place on **March 17th** 2016 from 13h30 till 17h, also at the CERN headquarters in Geneva. During this day a moderated workshop was organized to identify and **estimate the risks** for each of the previously identified **use cases**. These risks are discussed in more detail in section 3. An overview of the suppliers that attended the workshop can be found in Appendix 2.

Methodology

To assess the innovation potential of a use case for the PCP, two dimensions need to be considered. On one hand there is the added **value for the end-user**, on the other hand the estimated **risk for the supplier**. During the moderated workshops on March 16th and March 17th, procurers and suppliers were respectively asked to estimate the value and risk of the identified use cases using "planning poker".

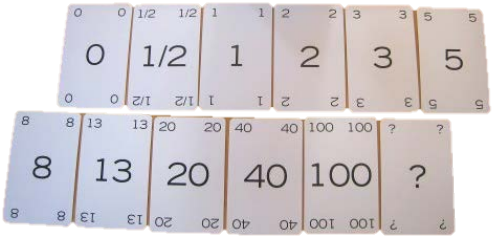

Planning poker is a "best practice" methodology to estimate e.g. added value, level of complexity, required implementation effort or risk. The technique is based on domain expert evaluation and achieving consensus. It uses a Fibonacci sequence to reflect the inherent uncertainty in estimating larger items.

Procurers were asked to estimate the **value** of a specific use case using the following scale:

0 - 2	**Could not care less**, not needed.
3 - 5	**OK, why not**, could perhaps bring some value.
13	**Must have**. Solution should have this requirement.
20+	**WOW, home run, this would really shift the boundaries!**
?	**Don't know**, no experience with this subject.

Suppliers were asked to estimate the **risk** of a specific use case using the following scale:

0 - 2	No worries, **off-the-shelf** standard solutions exist.
3 - 5	**A frequent problem**, certainly solvable.
13	**Absolutely not a standard problem.** A solution requires important choices, thorough elaboration, and specific expert effort. Success can be achieved with significant time and effort.
100	**Impossible**, requires breakthroughs of physical laws.
?	**Don't know**, no experience with this subject.

By combining both scores, it's possible assess the innovation potential of a use case. This is done by plotting each use case on two-dimensional graph. The vertical axis represents the added value for the end-user whereas the horizontal axis reflects risk. Within the PCP, the highest priorities for the envisaged solution are the items that realize the highest added value for the end-user, while holding a moderate to high risk.

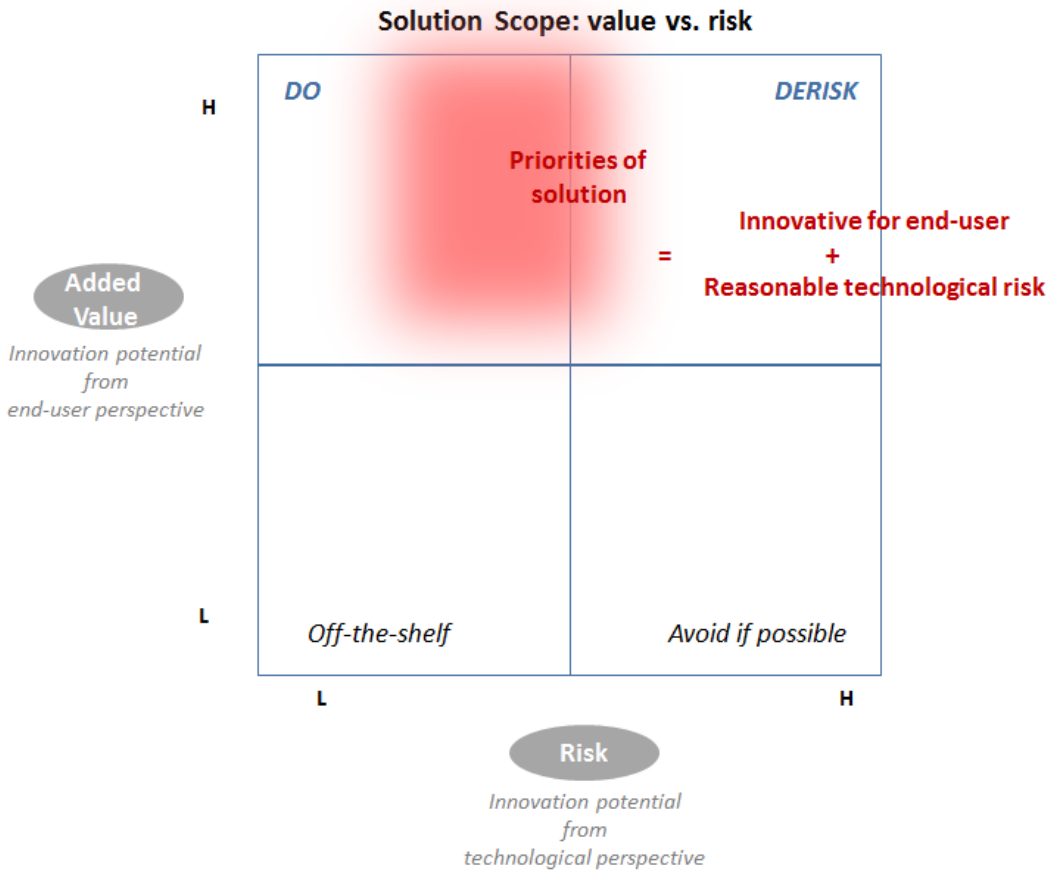

2. Overview of use cases and their value for the end-user

Based on the requirements for the HNSciCloud project, the procurers identified multiple **use cases**. Each of these use cases essentially describes specific functionality that is of value to the end-user (e.g. a research scientist, an IT manager, a procurer…).

Procurers estimated the **value** of each of these use cases using the planning poker methodology mentioned in the previous section. This value indicates the importance of a use case for the end-user. Below is the full list of use cases, **prioritized by value**. Use cases that offer **significant value** to the end-user are listed **first**, while less important use cases are listed last.

Use case 16: Efficient data access

Value assessment: 40

As a user
I can rely on the cloud provider to implement optimal strategies for handling very large datasets (e.g. via persistent caching, pre-staging)
so that I can reach high performance at an optimal cost (storage, network, CPU…).

Use case 25: Transparent data access

Value assessment: 40

As a user
I can have an application running in the cloud which sees the full datasets/environment in exactly the same way as on-site
so that I don't need to modify my application.

Use case 26: Configure once, deploy everywhere

Value assessment: 40

As a user
I can write the deployment configuration (including possible performance strategies) of my application once
so that I can easily deploy my application to all cloud providers using the same configuration.

Use case 8: Cloud economic models and procurement procedures of public organizations

Value assessment: 20

As Procurer
I can choose from different commercial delivery models e.g. reservation-based, spot-market, night-time bulk processing…
So that I can experiment with these and select the best model for my project.

Use case 13: Federated AAI and Credential Translation schemes (for resource management)

Value assessment: 20

As IT Manager or user
I can use my own credentials
so that the IaaS can easily be accessed and managed.

Use case 17: Seamless service

Value assessment: 20

As end-user (scientist)
I can use the service provided by multiple providers seamlessly
so that I don't to be aware that the service is being provided by multiple providers.

Use case 18: Metrics for aggregate performance, benchmarking

Value assessment: 20

As IT Manager
I can retrieve performance and benchmark metrics
so that effective end-to-end delivered service capacity & performance can be assessed.

Use case 20: Multi-cloud management frameworks

Value assessment: 20

As IT Manager
I can use a single multi-cloud management framework via open APIs
so that the integration costs of managing multiple clouds can be reduced.

Use case 22: Support

Value assessment: 20

As IT-Manager or system administrator
I can access adequately skilled support groups at the cloud provider
so that I can solve or act as coordinator/supervisor for the solution of complex problems running complex scientific workflows/applications.

Use case 23: Container support

Value assessment: 20

As a user
I can develop, maintain, manage and run my containerized applications (with all local dependencies) in the hybrid cloud via standardized tools
so that I minimize the deployment effort for me and for the scientific workgroups.

Use case 24: Scaling of containers

Value assessment: 20

As a IT Manager & Middleware Developer
I can deploy massive amounts (> 1000) of containerized applications on clouds from various providers meeting requirements on stability & performance
s**o that** I can easily and safely achieve massive scaling of scientific workloads with identical user and management experience.

Use case 19: HPC

Value assessment: some 13, some 20

As Procurer or IT Manager
I can in addition to typical configurations, also provision large memory, high core counts and/or larger/more performant local storage high performance backpanel
so that some end users are able to run their massively parallel applications.

Use case 1: Orchestration

Value assessment: 13

As IT Manager
I can provision resources (CPU, Storage, Network) through open APIs
so that transparently expanding my existing computing infra into Cloud is possible.

Use case 2: Monitoring and Dashboards

Value assessment: 13

As IT Manager
I can extract in real-time (< 10 min) resource usage metrics, through a dashboard or via APIs
s**o that** I can see the allocated resources and adjust these if necessary.

Use case 3: Reporting and accounting

Value assessment: 13

As IT Manager or Procurement Service
I can have a breakdown of the used resources over a period of time
so that an assessment of the accumulated cost of the resources consumed is possible.

Use case 4: Service level agreements

Value assessment: 13

As IT Manager
I can have a set of standard service definitions (Scope, Responsibilities, KPIs, etc.)

So that service level is defined, understood, agreed between parties for each service.

Use case 5: Network IP addresses

Value assessment: 13

As IT Manager
I can use Public IPv4 addresses (or equivalent addressing techniques) from within the private cloud
so that majority of applications such as databases and remote read-only file systems e.g. CERN CVMFS) that do no yet support IPv6, can be used.

Use case 6: Ordering, billing and invoicing mechanisms

Value assessment: 13

As IT Manager or Procurement Service
I can pay for and can get a receipt of the paid services
so that total expenditure of the used services is documented.

Use case 7: Performant WAN connectivity via GEANT

Value assessment: 13

As IT manager
I can efficiently move and access data from local resources to commercial clouds which are connected to GEANT at a speed of 40 to 100 Gbits/s
so that I have a seamless, performant and reliable computing experience.

Use case 9: Data privacy

Value assessment: 13

As IT Manager
I can restrict access to confidential data to specific users (e.g. LHC experiments)
so that data access can be restricted to an identified user community.

Use case 10: Pay-as-you-go

Value assessment: 13

As IT manager and Procurement Officer
I can have credits per cloud provider
so that I can see how many credits have been consumed for each cloud provider.

Use case 11: Pay-as-you-go refill

Value assessment: 13

As IT manager and Procurement Officer
I can refill the credits in my account and continue using cloud provider resources

so that I don't need to undertake an new procurement procedure and run the risk that contract conditions have changed.

Use case 12: Combination of firms/Sub-contracting

Value assessment: 13

As IT manager and Procurement Officer
I can expect suppliers that depend on subcontractors to have agreements with each subcontractor to provide at least 30% of the total required capacity
so that uninterrupted service can be ensured.

Use case 14: Helpdesk and Computer Security Response Teams

Value assessment: 13

As IT Manager
I can use a common Service Desk and Security Response Team
so that reporting and troubleshooting operational/security problems across the full infrastructure can be guaranteed.

Use case 15: Alerts

Value assessment: 13

As IT Manager
I can receive meaningful notifications and alerts about the resources I'm using
so that I have prompt information about infrastructural problems l cost (storage, network, CPU…).

Use case 21: Consultancy & training

Value assessment: 8

As a end-user or middleware developer
I can get help in migrating my application to the hybrid cloud and I can get training on how to use this hybrid cloud
so that I can effectively use the platform for my research.

3. Risk assessment of the use cases

For each of the user cases listed in the previous section, **suppliers estimated the risk** using the planning poker methodology mentioned in section 1. These risks could be technical, economical, organizational... To make it easy to look things up, the use cases below are **ordered** in exactly the same order as in the previous section (**by value**), despite having different risk values.

Use case 16: Efficient data access

Risk assessment: 8

The performance of accessing data (read/write) is influenced by many factors such as: concurrency, data size, network latency, architecture of the application, structured vs unstructured data etc.

The current wording of the use case suggests that it is up to the cloud provider to come up with an optimal solution. The cloud provider typically has no knowledge about the internals of an application however. This makes it particularly hard (and costly) to come up with such an "automagic" solution. Because each cloud provider uses different in-house optimization techniques (that will continue to advance in the future), a standardized solution is also unlikely to happen.

The general consensus of the industry is that collaboration with the developers of the institutions is necessary in order to achieve high performance at an optimal cost.

A responsible from T-Systems estimated this as a zero-risk use case. This person claimed to already have a solution available for this problem but was not permitted to discuss it in public.

Use case 25: Transparent data access

Risk assessment: 13 with a large variance

Several suppliers identified this use case as low risk (< 3). One solution for this use case was the use of a virtual file system (VFS). This ensures the data is accessible in the same way regardless of where the application is running. This solution supports different protocols and also supports additional authorization rules (e.g. ACL). From the developer's perspective, this would result in minor to no code changes. Object storage is an alternative solution to VFS that can be used to store data either locally or remotely, such that the end-user can have transparent data access.

Other suppliers identified this as a high risk use case. While it is definitely possible to make data access transparent in a single-cloud environment, there is currently no standardized way for migrating data between multiple cloud providers due to different data formats, conversions, protocols etc. There might also be potential geographical or political limitations on where data can be stored. Depending on how the data is organized in the private cloud, moving data into the cloud might require identity management to become a responsibility of the cloud provider.

One risk that every supplier agreed on was the potential difference in performance. Being able to access the data in a transparent way is easy, but maintaining the same performance (whether locally or remote) adds additional complexity.

As a summary, it can be agreed that (as a separate use case) it's definitely doable using existing technologies. When adding the performance requirement, the use case becomes riskier. In combination with other use cases (identity management etc.) the risk of this use case is further increased.

A responsible from Safe Swiss Cloud claimed to already have a solution available for this use case, but was not permitted to explain this solution in public.

Use case 26: Configure once, deploy everywhere

Risk assessment: 8

The primary risk is the absence of a standardized deployment configuration. Cloud providers have different interfaces and configurations that prevent applications from being configured once and deployed everywhere.

A suggested solution would be to define a template describing the application requirements, including functional but also non-functional requirements (e.g. SLAs). Each application then needs to be modelled in this template. Before an application is deployed, every cloud provider will be checked for compliancy with this template. Only cloud providers who comply with the application's requirements will be allowed to deploy it.

This will require a one-off API integration per cloud provider. Hence one may think to create a separate entity responsible for defining the template (could e.g. be CERN). This puts the burden of compliancy on the cloud provider, but not the definition of the template itself.

Use case 8: Cloud economic models and procurement procedures of public organizations

Risk assessment: 13

Each cloud provider currently has their own set of business models using different systems, pricing, APIs… These business models are also optimized for the industry and not for research institutes (difference in CPU usage for example). To prevent cloud providers from (blindly) implementing all possible business models, research institutions should openly discuss with the cloud providers which business models make sense and how to use them.

Use case 13: Federated AAI and Credential Translation schemes (for resource management)

Risk assessment: 3 to 8

The risk value for this use case depends on the complexity of the solution. Managing an AAI as a single cloud provider is doable (3). Doing the same in a multi-vendor cloud environment (m institutions x n cloud providers) is considerably more difficult (8). Next to this, there's also the question of who will take up ownership of this feature. Further collaboration between research institutes and cloud providers is required.

Use case 17: Seamless service

Risk assessment: 3

Assuming there are no significant performance differences between cloud providers, no risks were identified that should prevent the user from receiving a seamless service.

Use case 18: Metrics for aggregate performance, benchmarking

Risk assessment: 5 to 8

Measuring the performance of cloud providers requires reliable and repeatable benchmarks. In practice, benchmarks might not always reflect the actual application behaviour and might be influenced by many factors such as: concurrency, data size, network latency etc. Next to this, the benchmark results will also differ per cloud provider due to different in-house optimization techniques.

While it is possible to agree on a common definition of performance metrics and benchmarks, care should be taken when interpreting the results.

Use case 20: Multi-cloud management frameworks

Risk assessment: 2

Currently each cloud provider has their own API(s). Defining a cloud management framework that supports multiple cloud providers poses few technical risks. The OpenStack framework for example already provides connectors to interact with e.g. Amazon EC2.

Use case 22: Support

Risk assessment: 2

No primary sources of risk. Cloud providers already have support groups available to provide assistance with technical problems.

Use case 23: Container support

Risk assessment: 5

Although the industry is moving towards adopting container technologies, it's still a relatively new technology that is not yet supported by every cloud provider. In practice it's definitely doable, but procurers should be aware that using container technology will require software changes on their side.

Use case 24: Scaling of containers

Risk assessment: 8

As a single-cloud environment that supports containers, it should be possible to deploy a massive amount of these containers. In a multi-vendor cloud environment however, scaling an application across multiple cloud providers might introduce additional risks concerning security, performance etc. Despite these risks, this use case remains doable.

A responsible from T-Systems claimed to already being able to scale their applications.

Use case 19: HPC

Risk assessment: 8

Not all cloud providers support high core counts, performant storage solutions etc. at the moment. There's also an absence of a standardized deployment configuration that specifies requirements (RAM, CPU etc.) for each application. The template used to solve use case 26 could overcome this challenge and would allow applications to only be deployed on cloud providers that match the requirements.

Use case 1: Orchestration

Risk assessment: 5

The primary source of risk for provisioning resources (CPU, storage…) through open APIs is the lack of a standardized API across cloud providers. This could be solved by implementing a generic API that transforms requests to the APIs of the individual cloud providers.

Use case 2: Monitoring and Dashboards

Risk assessment: 3

Every cloud provider is already capable of generating resource usage metrics for their own cloud. In a multi-vendor cloud environment however, these metrics may differ in content, format, granularity etc. due to a lack of standardization. This could be solved by a separate layer that transforms metrics from different cloud providers into a standardized format.

Use case 3: Reporting and accounting

Risk assessment: 5 (remark: don't be stupid)

The primary source of risk for extracting a breakdown of the used resources is the lack of a standardized approach. In practice this can definitely be solved. However, it will be important to collaborate with the procurers to validate if the existing solutions are sufficient for them.

Use case 4: Service level agreements

Risk assessment: 3, except for two 20s and one 13

Every cloud provider already has their own set of service definitions and SLA's. There is however no common definition across cloud providers. Coming to a cross-provider agreement will be hard. Procurers are recommended to not impose an (additional) definition on the cloud providers.

Use case 5: Network IP addresses

Risk assessment: 2, except for one 8 and two 13s

Although cloud providers currently support IPv4, they highly recommend to migrate towards IPv6 because public IPv4 addresses are becoming increasingly more expensive.

Use case 6: Ordering, billing and invoicing mechanisms

Risk assessment: 1

Every cloud provider has ordering, billing and invoicing mechanisms available. No primary sources of risk were identified.

Use case 7: Performant WAN connectivity via GEANT

7a) In case the use of the GEANT network is mandatory (recommended approach):

Risk assessment: one group of 1 and one group of 8 to 13

The primary source of risk is the 100 Gbit/s connection to the GEANT network. The expensive network costs for the cloud providers could potentially kill the use case.

7b) In case the GEANT network cannot be used, using the available PCP budget:

Risk assessment: 13 to 20

If the GEANT network cannot be used, cloud providers are responsible for setting up separate connections to each research institute, which could potentially become very costly. An alternative solution that was proposed is to use CERN as a gateway to connect to the individual research institutes.

It can be concluded that further discussions between procurers and suppliers are necessary to define scope, requirements, budget, solutions etc. to come up with a cost effective solution that is beneficial to all parties.

Use case 9: Data privacy

Risk assessment: 3

No primary sources of risk identified. Cloud providers already provide mechanisms to restrict data access to confidential data.

Use case 10: Pay-as-you-go

Risk assessment: 5, except for three 8s

No primary sources of risk identified, assuming that each cloud provider is only responsible for managing his own credits. It is not required to compute the credit consumption for other cloud providers.

Use case 11: Pay-as-you-go refill

Risk assessment: 3

No primary sources of risk identified.

Use case 12: Combination of firms/Sub-contracting

Risk assessment: 3

No primary sources of risk identified, but this rule might hinder innovation.

Use case 14: Helpdesk and Computer Security Response Teams

Risk assessment: 8 to 13

Easy in a single cloud environment, but harder in a multi-vendor cloud environment where cloud providers have different definitions of service desk. Security response teams are also not standard for every cloud provider.

A responsible from T-Systems estimated this as a zero-risk use case. They implemented a solution using managed cloud services.

Use case 15: Alerts

Risk assessment: 5

Every cloud provider is already capable of generating meaningful notifications and alerts within their own cloud. In a multi-vendor cloud environment however, notifications and alerts may differ in content, format, granularity etc. due to a lack of standardization. This could be solved by using a message broker that transforms messages from different cloud providers into a standardized message.

Use case 21: Consultancy & training

Risk assessment: 2

Several use cases already identified the need for a tighter collaboration between research institutes and cloud providers. This should help research institutes to get a better understanding of which steps are required to migrate applications to the hybrid cloud.

From the perspective of the cloud provider, there are also no technical risks involved to provide documentation, training or other services to the research institutes.

4. Value vs Risk Prioritization Matrix

After assessing the value (section 2) and risk (section 3) of every use case, it is now possible to plot the results on a **value-risk prioritization matrix**. This is a two-dimensional matrix of which the **horizontal axis** indicates **risk** and the **vertical axis** represents **value**. By plotting each use case in this matrix, clusters with innovation potential can be identified.

The value-risk prioritization matrix itself can be divided in **four quadrants**. The **bottom left** quadrant contains use cases that offer little value for the end-users and few risks for the suppliers. Such use cases typically have very low innovation potential. It's very likely that ready-made solutions are already available on the market.

The **top left** quadrant contains use cases that offer a high amount of value for the end-users but with few risks for the suppliers. Ready-made solutions to solve these use cases might already be available on the market.

The **bottom right** quadrant contains use cases that offer low value for the end-users but with many risks for the suppliers. Very little value is created in relation to the costs of development.

The **top right** quadrant contains use cases that offer lots of value to the end-users but also many risks for the suppliers. These are the use cases that typically have a high innovation potential.

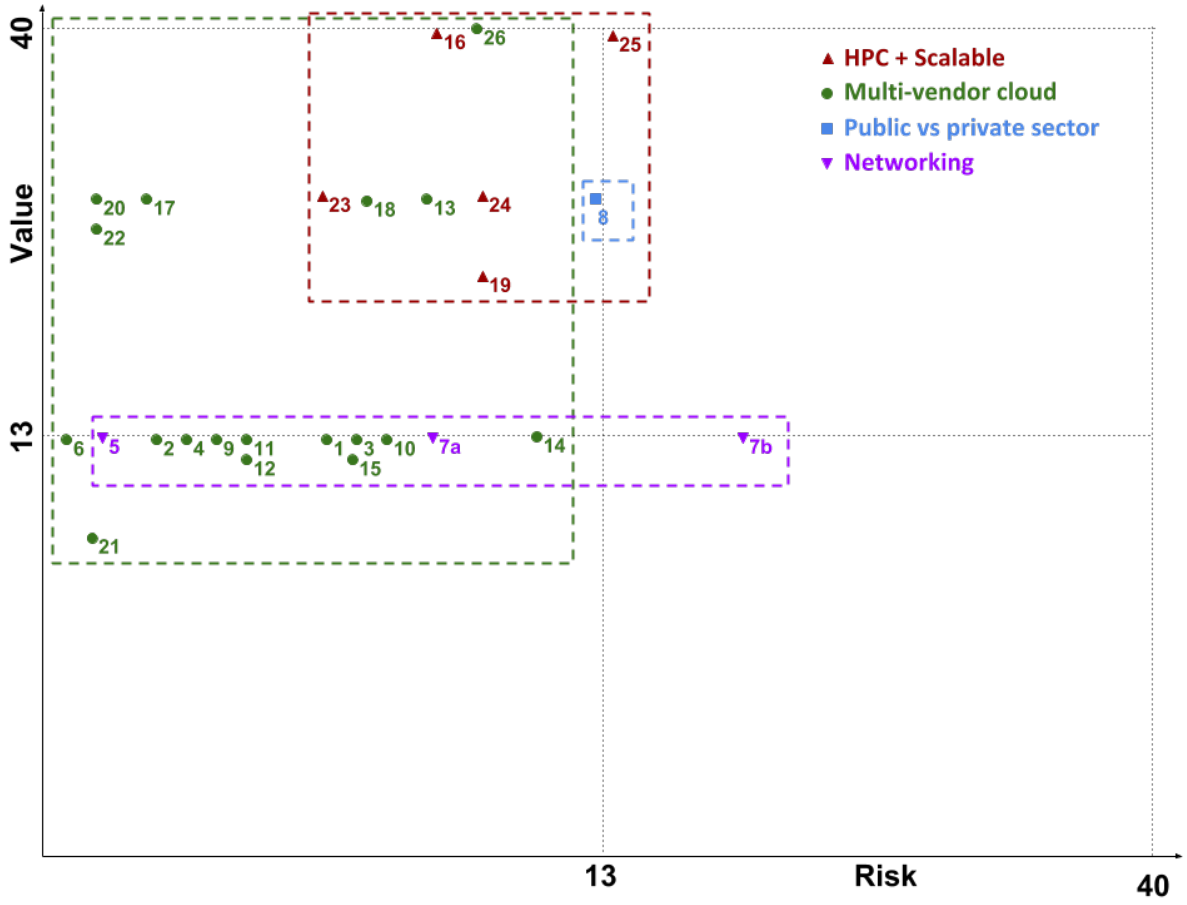

From the matrix, **four modules** can be identified:

- High Performance Computing (HPC), including scalability
- Network addressing and connectivity
- Multi-vendor cloud platform
- Public vs. private sector

High Performance Computing (HPC), including scalability

This module contains all the use cases related to the development of a High Performance Cloud Computing (HPCC) platform that is also scalable. From the prioritization matrix it is clear that the majority of use cases in this module offer high value with moderate risks. This module has the highest innovation potential.

The following use cases are part of this module:

- Use case 16: efficient data access
- Use case 19: High Performance Computing (HPC)
- Use case 23: support for deploying containers
- Use case 24: massive scaling of containers
- Use case 25: transparent data access

Network addressing and connectivity

This module contains all the use cases related to network addressing and connectivity. From the prioritization matrix it is clear that the majority of use cases in this module offer moderate value with moderate risks. Some of the requirements are trivial to solve, others will require more innovative solutions.

The following use cases are part of this module:

- Use case 5: network IP addresses
- Use case 7a: performant WAN connectivity with GEANT
- Use case 7b: performant WAN connectivity without GEANT

Multi-vendor cloud platform

This module contains all the use cases related to building a multi-vendor cloud platform. From the prioritization matrix it is clear that the majority of use cases in this module offer moderate to high value with fewer risks. As a single cloud provider, each use case is perfectly solvable using solutions that are already available on the market. In a multi-vendor cloud environment however, the innovative value will be created by defining an optimal collaboration model between the different cloud vendors.

Deciding on an optimal collaboration model that encapsulates the underlying technical challenges will be critical

The following use cases are part of this module:

- Use case 1: orchestration
- Use case 2: monitoring and dashboards
- Use case 3: accounting and reporting
- Use case 4: SLAs
- Use case 6: ordering, billing and invoicing mechanisms
- Use case 9: data privacy (ACL)
- Use case 10: pay-as-you-go
- Use case 11: pay-as-you-go refills
- Use case 12: subcontracting constraints
- Use case 13: federated AAI and credential translation
- Use case 14: helpdesk
- Use case 15: alerts and notifications
- Use case 17: seamless service
- Use case 18: performance metrics and benchmarks
- Use case 20: multi-cloud management framework
- Use case 21: consultancy/training
- Use case 22: support
- Use case 26: configure once, deploy everywhere

Public vs. private sector

This module identifies the differences between the requirements of the public sector and the capacities of the private sector. It only consists of one use case: use case 8, cloud business models. The innovative value comes from connecting both worlds.

5. Conclusion

For the PCP **three main challenges** can be identified for the construction of the hybrid cloud. First and foremost is the development of a **High Performance Cloud Computing (HPCC) platform**. This platform must support a range of **Virtual Machine and container** "configurations". It is expected that minimal to no code changes will be required to deploy existing applications to this platform. Since the majority of the applications are **data intensive**, efficient data access will be key in meeting the acceptance criteria. **Petabytes** of data are expected. Suppliers are free to **design a "full solution"** and, if necessary, collaborate with other suppliers in order to address the many use-cases that are part of this challenge.

Secondly are the **network addressing and connectivity** challenges. This requires the construction of a **high bandwidth network** connection across country borders using the services of different public and private suppliers. Suppliers are **recommended** to use **the GEANT network** to overcome this challenge. This module also requires IPv4/VPN addressing solutions and support for federated identity management using **eduGAIN**.

Thirdly are the **differences** between **the public and private sector**: how to match the expectations of the public sector with the fast evolving capabilities of the private sector? Various purchasing options (pay-as-you-go, spot-market, reservation-based, over-night/low-priority scheduling, etc.) should be investigated to determine the most appropriate ones for the public research sector.

In addition to the already known list of use cases and requirements, it's expected that **no infrastructural changes** (either hardware or software) are required on the procurers' site. In case such changes are necessary for one or multiple use cases, an in-depth analysis will be required to assess the current and future implications of this approach.

For the PCP, the following **deliverables** are expected:

Initially a **design** should be proposed that explains how the platform will be able to support all the use cases (in particular the **HPCC** and **networking challenges)** given a range of purchasing options. This deliverable typically consists of datacenter and cloud architecture diagrams, network diagrams and many others.

In the second phase, a **prototype** should be developed that provides modest compute, storage and network capacity. The main objective of this prototype is the ability to run a suite of **tests** (one test per use-case) that **validates** whether the prototype meets the expected requirements. These acceptance tests will be provided by the procurers and made available to the suppliers. Procurers should be able to run the tests themselves to verify the prototype.

Finally, a **pilot** should be developed on which procurers can **deploy and test** their **applications**. Compared to the prototype, this pilot is expected to have significantly more compute, storage and network capacity available. Procurers will keep the required capacity within reasonable limits

for the suppliers, however. This pilot will also be used for performing **scalability and reliability tests**. The various purchasing options will be evaluated as well. This should lead to accurate cost estimates for running applications in production.

Appendix A: Procurers Attending the Workshop on March 16th

Below is an overview of all procurers that attended the workshop on March 16th. The procurers identified and prioritized use cases based on the value for the end-user. The majority of these procurers also attended the open market consultation on March 17th.

Consortium Partner	Number of attendants
CERN	15
INFN	3
DESY	3
CNRS	2
KIT	1
SURFsara	1
STFC	-
EMBL	3
IFAE	2
ESRF	2
Trust-IT	1
Stitchting EGI.eu	1

Appendix B: Suppliers Attending the Workshop on March 17th

Below is an overview of all suppliers that attended the open market consultation workshop on March 17th. These suppliers were responsible for evaluating the risks of every use case.

Company	Number of attendants
100 Percent IT Ltd	1
Atos - BULL	1
AWS	1
BIOS IT	4
Cloud Security Alliance	1
Clouditalia Telecomunicazioni SpA	1
Cyfronet	1
DataCentred Ltd	1
Forum Virium Helsinki Oy	1
GEANT	1
GMV	1
Google Switzerland	1
Huawei	2
IBM	5
Indra Sistemas	2
iNNOVO Cloud GmbH	2
Insight Technology Solutions	1
Intel Corporation	1
Microsoft	3
Oracle	1
ORANGE Cloud for Business	3
Prologue	1
Rakspace International GmbH	2
Sansom IT	1
SixSq	1
T-Systems International GmbH	1
Telefonica	2
The Server Labs Ltd.	1
TIS	1
Vancis	1

Appendix C: Type of Participation

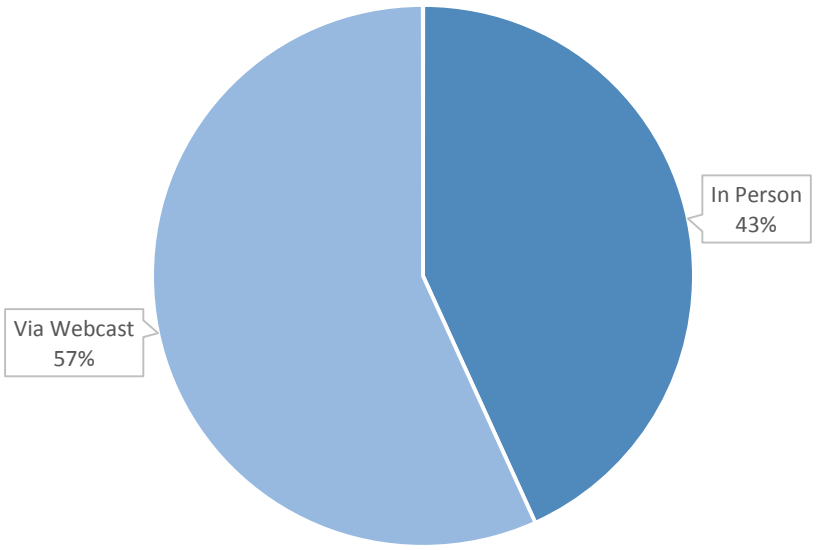

Appendix D: Geographical Distribution of Participants (webcast and in person)

*11/20 participants from CERN

Appendix E: Views of Event Recordings (from 17 March to 08 April 2016)

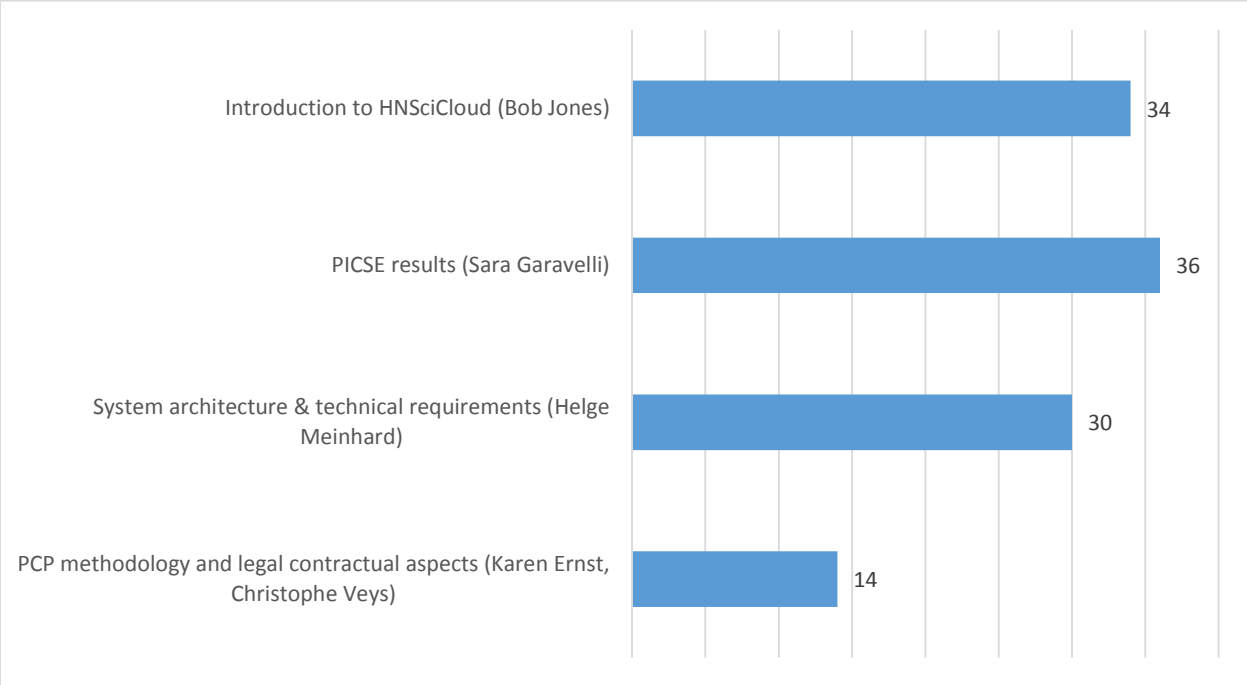

Appendix F: Distribution of participants (78 total)

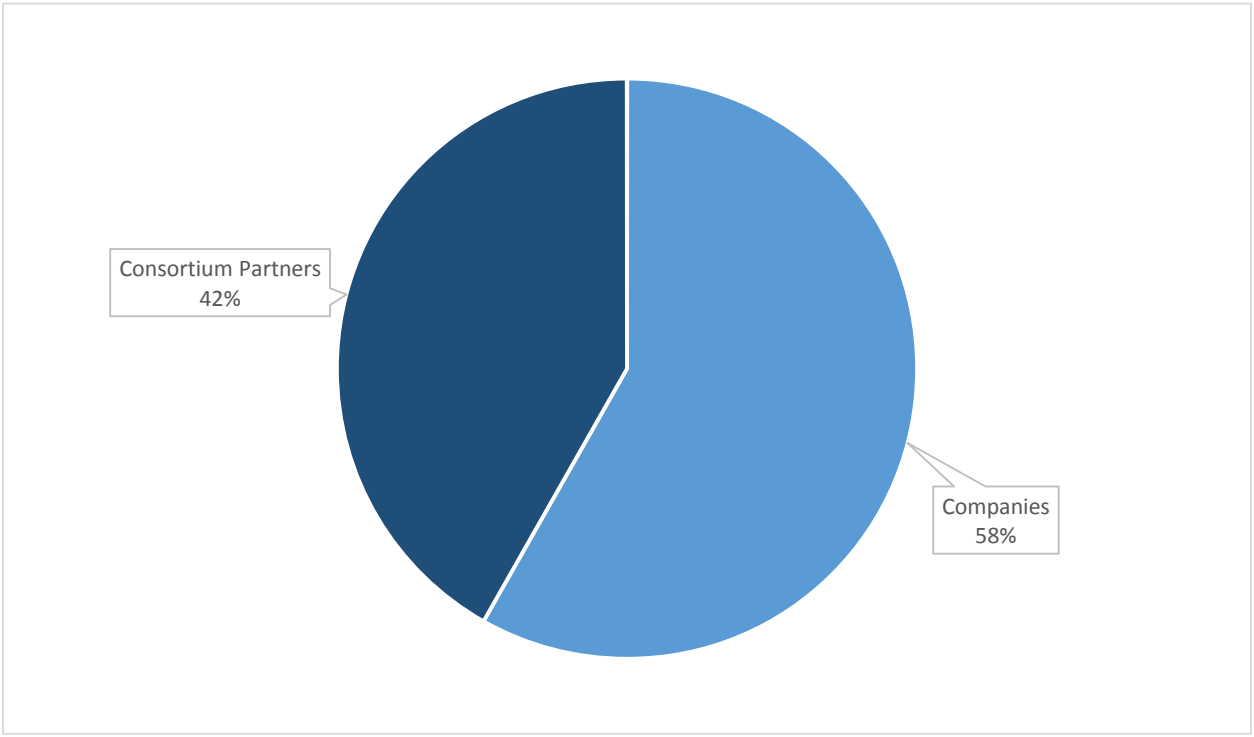

Appendix G: Size of Entreprises