Helix Nebula – The Science Cloud

Title: Costing exercise comparing in-house vs. cloud based operation for the CERN flagship use-case and incorporation of qualitative cloud adoption criteria targeting prospect members

Editor: SAP (Julia Doll, Jochen Werner)

Work Package: 7 – Business Models

Submission Date: 20th February 2014

Distribution: public

Nature: report

Revision History			
Version	**Date**	**Description**	**Contributor(s)**
1.0	30.01.2014	Initial Draft	Julia Doll (SAP), Jochen Werner (SAP)
1.1	05.02.2014	Update	Robert Jenkins (CloudSigma), Jurry de la Mar & Bernd Schirpke (T-Systems)
1.2	12.02.2014	Update	Michael Symonds (Atos)

Table of Contents

Table of Abbreviations

AWS	Amazon Web Services
BM	Business Model
BMI	Business Model Innovation
CERN	European Organisation for Nuclear Research
e.g.	for example
EC	European Commission
EMBL	European Molecular Biology Laboratory
ESA	European Space Agency
etc.	et cetera
EU	European Union
FP7	7th Framework Programme for Research and Technological Development
FTE	Full Time Equivalent
HN	Helix Nebula
HNX	Helix Nebula Marketplace
HPC	High-Performance Computing
HTC	High-Throughput Computing
IaaS	Infrastructure as a Service
LHC	Large Hadron Collider
SLA	Service Level Agreement
SaaS	Software as a Service
TCO	Total Cost of Ownership
WAN	Wide Area Network

Table of Figures

1. Executive Summary

The aim of the Helix Nebula initiative is to pave the way for the development and exploitation of a Cloud Computing Infrastructure. It is initially based on the needs of European IT-intense scientific research organisations, but also allowing the inclusion of other stakeholders' needs (governments, businesses and citizens). This deliverable is part of Work Package 7, which aims to understand the financial implications of 'utility computing' for vendors and customers, define mechanisms for quantifying and controlling risk and assess the viability of standard cloud-service procurement templates across jurisdictions. This deliverable specifically aims to compare the TCO of current on premise deployments at CERN with potentially outsourced on demand instances offered by the participating suppliers. Originally, the use cases of the flagships were supposed to be taken as reference to perform this exercise. Because of divergent specifications compared to the HN Market Place, which was set up in January 2014, it would not have been very beneficial. Therefore, we reached out for alternatives and identified CERN, an initial member of HN, as reference case.

Gathering all relevant figures is the most important task when defining the benefit of outsourcing activities. At the same time, this is the most challenging. In order to gain insights of CERN's cost structure, SAP visited its premises. While interviewing relevant staff, the following question was raised: What is the in-house cost level of services which could be outsourced? Basically, what are the true TCO of applications and services which are currently operated on premise? Most probably, such deployments not just provide one specific service where costs can be allocated accurately, especially in IT and publicly funded organisations. That does not mean that corporates do not face these challenges, but their cost and performance accounting is supporting these kinds of assessments heavily. In contrast to them, CERN follows a budget based approach and does not track precisely which service is used by which department. Consequently, it is not clear what amount of financial and personnel resources are needed to fulfil specific services. In the past, these publicly funded organisations were not obligated to conduct hat marginal costs because there simply was no service available in the market which could replace their in-house solution. So they had to build and run most of their applications in-house. The lack of metering marginal volumes and costs of services that could potentially be outsourced is probably something, the majority of all publicly funded (research) organisations have in common at least to some extent.

CERN's IT infrastructure is huge and highly complex. Thus, we narrowed the costing exercise down to two specific services: The private cloud based on OpenStack, which is

the underlying infrastructure for the other investigated service: Zenodo, a digital repository for data storage purposes. This is a SaaS offering by CERN. The lack of metering of the cost assessment in this deliverable is causing a not explicit enough outcome. Because of the investigated services' lack of metering and cost allocation, you should not misinterpret the following statements.

The outcome of the cost comparison between the current on premise deployment and the on demand of HN suppliers is interesting, although a lot of data is still missing.
Based on volume, HN suppliers grant up to 45 per cent pricing discount. Please remember the lack of metering and cost allocation to not misunderstand the figures mentioned. Also note that the cloud instances include all cost components in contrast to CERN's in-house calculation. Hence, the gathered figures at CERN do not represent a TCO calculation. The most cost-efficient offer that we received from HN suppliers for the private cloud is approximately 3.4 times more cost-efficient than AWS. However, even the HN prices are several times higher than the current CERN in-house deployment. HN suppliers are very attractive according to the Zenodo service. Even without gathering all figures at CERN, HN suppliers are more cost-efficient when discounts are assumed (in case you consider the OpenStack cloud and Zenodo service separately from each other). In general, the longer a customer is willing to commit to use a cloud service, the more competitive offers he will obviously get. In both services, the annual discount between a one and three year commitment is roughly 30 per cent.

Especially for research organisations, HN offers the ideal combination of taking advantage of the cloud while being able to connect to other non-profit operators, e.g. EGI and GÉANT. By attracting more and more members, (scientific) knowledge is enhanced and economies of scale are achieved, which enables the suppliers to further lower their prices. In this context, it was essential to also incorporate qualitative cloud adoption criteria. Even if cloud services are mainly assessed by the economic benefit, a lot of other aspects should be mentioned, both strategically and operationally. Thus, a variety of different qualitative issues are explained.

Later on, the attractiveness of HN for supply and demand is described. By joining this initiative, suppliers are able to gain huge business opportunities. Organisations with these massive data requirements are an uncovered new target group to address. This and also the renowned flagships could and should be leveraged for other customers and public relation purposes. Through overcoming the complex challenges of interacting with the variety of members, big data know-how is enriched and competence is proven.

The flagships take advantage of dynamic computing. They can rely on professional (European) providers to not only overcome performance peaks, but also use permanent services. Within HN, they have the opportunity to actively shape a European counterpart to leading American technology companies. Furthermore, collaboration with other organisations is simplified and more efficient while gaining high level data security of European providers.

Finally, should be stated that a lot of data is still missing in order to gain a proper cost comparison. The confidentiality that CERN is required to comply with makes this cost assessment even more difficult.

2.   Introduction

This deliverable aims to perform a costing exercise for the BM Generic Cloud Computing for European Science. After previously researching a variety of potential BMs with the BMI approach of SAP, this one was characterized as worth pursuing and identified as the stepping stone to establish the other chosen BM Information as a Service. We are here concentrating on the first mentioned, which is further illustrated by the following figures.
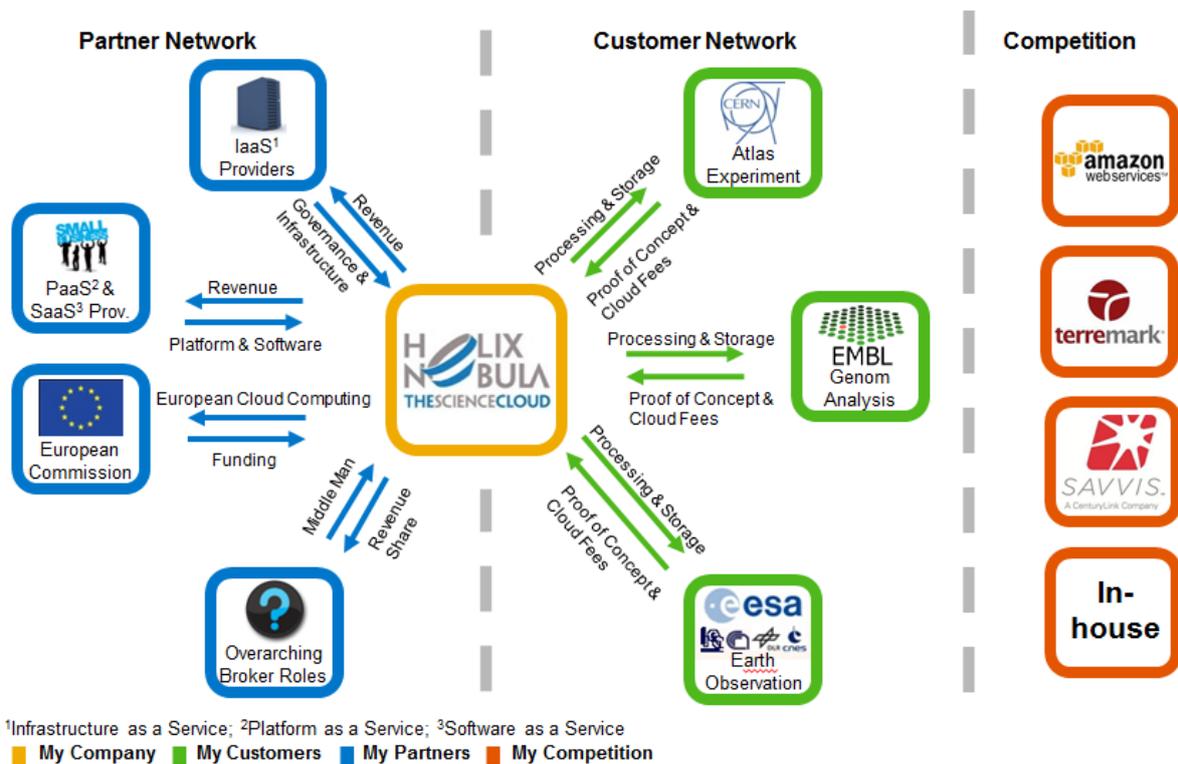
Figure 01:1 Generic Cloud Computing for European Science (Network View)
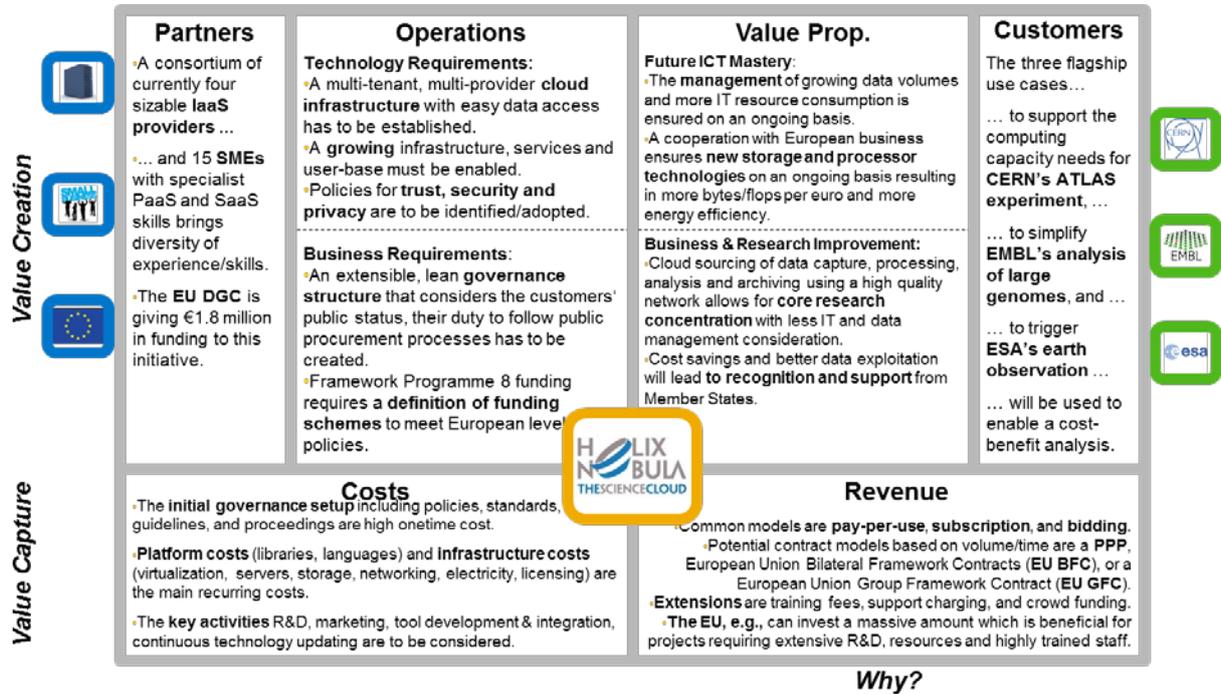
Figure 02:2 Generic Cloud Computing for European Science (Enterprise View)
EU BFC = European Union Bilateral Framework Contracts
EU GFC = European Union Group Framework Contracts

The initial goal of the assessment is to gather and analyse the economic figures for providing a quantitative assessment of HN for the participating members. The suppliers invested a reasonable amount of resources in order first to develop and then to demonstrate the attractiveness of HN but it is not proven yet. At this stage, it is still not verified, if the HN investment of the participating parties will finally pay off for them. Thus, it is not clear, if HN will turn into a worthwhile business and the suppliers' investment will pay off. It is indeed a good learning exercise which could lead to a competitive advantage.

The flagships committed to using the HN Market Place as soon as its technical maturity is given. Therefore this deliverable could rather be used as showcase for prospect members.

While all members are busy with setting up the HN Market Place, the argument was raised, as to whether the costing exercise would make sense at this stage. Probably this task should be carried out once more when the HN Market Place is used for a period of time

3. Establishing the Helix Nebula Marketplace

After completing the pilot deployments the use cases of each flagship successfully, the HN suppliers set up a productive infrastructure environment in January 2014. This was an ambitious goal, taking the short time since the third General Assembly in September 2013 and the technical complexity into account.

The HN Market Place press release states: "The HN suppliers Atos, CGI Logica, CloudSigma, Interoute, SAP, Sixsq, TSL and T-Systems have expressed the intention jointly to start a European Cloud Market Place to offer a European Cloud services alternative that fully complies with EU regulations and legislation. The Cloud Market Place will be branded HN Market Place to communicate its Vision and Mission, following the FP7 project of that name that was made possible with a grant from the EC. The HN Market Place will deliver easy access to a range of commercial Cloud Services through the innovative broker technology that has been developed within the HN project over the last two years. The participating suppliers will compete within the HN Market Place and offer their Cloud services at competitive prices. The HN Market Place will be open for new cloud providers that are able to participate on an equal basis."

The Cloud services are initially targeted at the global scientific community, for both publicly funded and commercial Research and Technology Organisations, within e.g. oil and gas, high tech, and manufacturing industry, offering large-scale, HTC and HPC-type deployments from the start.

With HN Market Place current members are able to choose between various suppliers or combinations of suppliers, to buy, use and manage Cloud Services seamlessly. Following an initial period of operations, the partners are also determined to extend the HN Market Place with Information as a Service offering, making open research data and algorithms from various Research and Technology Organisations available to the public in a pre-configured processing environment. Eventually, the HN Market Place will lead to a Digital Energy Market Place, where users will have easy access to a wide range of services including digital infrastructure, tools, information and applications."

Thus, this market place will be the starting point for dynamic offering and purchasing of cloud computing capacities within HN. The interface will allow prospect customers to choose the best suitable offer for them in real-time. Since the technical specifications of the use cases were unique concerning volumes and granted discounts, taking them as reference for the productive environment of the HN Market Place, would not be useful. Therefore we will use a different approach which will be outlined in the next chapter.

4. Cost Assessment for Generic Cloud Computing for European Science

4.1 Methodology

At first, we assess CERN's in-house costing level, secondly we calculate the respective prices of HN suppliers and AWS. Furthermore, it is described what a newly installed, self-constructed on premise IaaS deployment would cost. All this together gives a comprehensive overview about the different options every organisation has.

4.2 CERN IT Background Information

CERN is the world's largest particle physics lab and home of the Large Hadron Collider (LHC), the world's most powerful accelerator providing research facilities for High Energy Physics (HEP) researchers across the globe. LHC experiments will run up to 1 million computing tasks per day and generate around 15 petabytes of data per year over a 10-15 year period. CERN leads the World-wide LHC Computing Grid project (WLCG), to provide computing resource to store, distribute, analyse and access (in near real-time) LHC data for a community of more than 10,000 physicists worldwide. It is astonishing that only around 8 per cent of the computing takes place at CERN itself. The rest is done by associated data centres around the world.

This introduction impressively illustrates the importance of sophisticated IT systems in general, as well as data centres in particular, needed to fulfil these tasks. As a consequence, CERN is used to invest heavily in state of the art technology. Because of its extraordinary reputation, some hardware and software manufacturer provide their solutions with discounts. Marketing and technical issues are two reasons for that. If CERN is using a solution, the manufacturer can use it as a reference to showcase the maturity and capability of its product. Additionally, CERN has long-time experience of testing new technology which is not market ready yet. With their feedback, they can further improve the product.

Given these facts, why did we choose to assess the costs of such an already high efficient organisation? Publicly funded organisations, especially those who have been running their own data centres for decades, tend to not include all relevant aspects into their IT costing exercise. Therefore it is highly recommended that an expert outside of the

organisation takes an insight look at the TCO. We recognize the challenge of identifying hidden costs in an organisation like CERN. It could be an eye-opener for others to join the initiative. Especially for those who are convinced about their unbeatable IT cost level. Additionally, we will outline what the next steps are to be taken in order to be able to make a decision for or against cloud services.

The CERN on site visit included meetings with a variety of staff, involved in the running of IT services. It proved what was to be expected: Not surprisingly, publicly funded organisations mostly do not have cost and performance accounting systems like corporates. Hence, they neither know specific cost figures, nor have defined service levels. That does not make this cost assessment easier.

Assessing the costs of specific corporate services is not even that easy, but the design of their accounting systems helps a lot. It is a common procedure at corporate outsourcing projects, to let external consultants take a look at every cost type which affects the calculation. Such exercise can easily take months and keep a reasonable amount of staff busy on both sides, also depending on the company's size.
In contrast to them, publicly funded organisations do not aim for profit maximizing. Because there was no alternative available in the market in the past, which could have replaced the self-operated service, determining costs was apparently not first priority.

Today is a different story. Data centres of technology companies like Google and Amazon driven by consumer demand and the leading global ICT providers driven by outsourcing contracts nowadays exceed data capacities of CERN for the first time since its foundation. For decades, CERN had to build the necessary applications because there simply was no alternative. This has changed with the rise of these technology companies. The shift from developing local solutions completely from scratch to questioning the need and seeking the ideal solution is massive. This not only affects the IT department but the whole organisation. It can easily be defined as change management because dealing with service providers requires different skill sets. Programming effort declines, the level of required governance raises, just to name an example.

In that context, it would be useful for potential new members to clarify, if IT services are considered a core competency. Maybe IT services will not always belong to that category for CERN in the future, but scientific services might. If this holds true, it does not consequently mean to not use on demand offerings at all, for example for covering peaks it could still proof beneficial.

The biggest challenge is not creating the business case itself but gathering the figures necessary for it. Thus, the challenge at CERN was to determine marginal volumes and costs of services that could potentially be outsourced. Basically, which portion of which service is consumed by whom? Again, take the complexity of CERN`s IT infrastructure into account.

With that in mind, we narrowed the whole costing exercise down to two specific services: the private cloud based on OpenStack (IaaS) and Zenodo, a digital repository that offers permanent storage for research data for any science domain (SaaS). The private cloud is relying on other departments responsible for tendering, purchasing, customizing, deploying, operating, monitoring and maintaining (new) hardware. The SaaS solution Zenodo is in turn naturally depending on the Infrastructure (OpenStack cloud). Again, how can you determine the cost of these in-house services, if interdepartmental interactions and volumes are not metered?

Because of contractual obligations with third parties, CERN is not allowed to disclose detailed figures. This is the reason why areas of investigation cannot be disclosed outside CERN and SAP.
Therefore we will perform the cost comparison with the suppliers not with absolute but relative numbers. By doing so, we can benchmark the findings and protect confidentiality at the same time.

4.3   AWS Background Information

To benchmark the findings with current market offers outside of HN, we calculated the OpenStack cloud and Zenodo service for AWS. They offer three different kinds of services which are comparable to the HN offerings:

a) On demand instances
b) Reserved instances
c) Spot instances

a) The on demand instances provide compute capacity that is cleared on an hourly basis without long-term commitment. Of course, you do not have to worry about planning, purchasing and maintaining hardware so you have variable costs instead of high upfront costs. On the other hand, prices are relatively high.

b) Reserved instances have a contract period of e.g. minimum one and maximum three years. You have to pay one-time to reserve each instance you want but therefore gain a high discount on the hourly price paid for on demand instances. Available instance types are Light, Medium and Heavy Utilization that enable you to pay in advance. In case you do not use the instances you purchased, you have to pay them anyway until the end of the contract period, which makes this offer inflexible.

c) Spot instances are based on supply and demand and change periodically. Customers can bid on unused capacity and run those instances as long as their bid exceeds the current spot prices. For non-time critical applications this can be a good alternative because this could be the most competitive service offer but also the most volatile. Therefore it is not predictable.

The reason for providing this background information is that there is not only one price for cloud computing as can be seen by the variety of these services. The more and longer a customer is committed to pay for a certain level of resources, e.g. via bundles consisting of a fixed amount of computing, storage and database capacity, the less he has to pay. Hence, there will be significant differences between the different purchase options of AWS. Similar services are offered among the HN suppliers.

The calculations in the cost assessment were all made with reserved instances. Using on demand instances for that amount of computing would be a lot more expensive. In other cases, spot instances could be a reasonable alternative, but not in our case here.

Please note:
In case of unexpected downtime or even data loss, AWS grants a "Service Credit" applicable against future AWS payments only. The customer has to prove the not-flawless incident. It could be doubted, if that is an ideal customer experience.
Probably data loss is very unlikely looking at their level of redundancy, but still a considerable issue.

4.4 Cost Assessment

With the hurdles described above, we were able to conduct a cost calculation which is based on the few numbers we could gather. This gives us an idea of the rough cost level for the two mentioned services.

Please note: Lack of detailed volumes and figures
Of course, this outcome may not be a proper apple-to-apple comparison. Given the fact that it was not possible to get detailed figures of CERN's IT TCO, you should not misunderstand the outcome. On the one hand there is an incomplete, thus imprecise cost assessment of the current in-house deployment and on the other hand, a fully covered service, offered by the provider.

Investigating the following cost types is necessary due to lack of metering:
- Cost and performance accounting within the different groups: greatest challenge to allocate FTE and infrastructure investment volume consumed by OpenStack cloud and Zenodo
- Potential sunk cost in data centre: e.g. server, switches, routers, racks and marginal volumes like cooling, electricity, uninterruptible power supply
- Potential sunk cost per office m^2: e.g. heating, security, workplaces, cleaning
- Staff training
- Other marginal volumes: Networking WAN, office LAN, value added networking, data centre LAN, 3rd party and remote access connection, storage and backup, service infrastructure, helpdesk, database management, technical application management, configuration management assets, governance, innovation aspects in general

= The total IT budget is the sum of all man hours (FTE), supporting contracts and investments following the principle: In order to deliver a service, you have to spend working hours, use support from external service suppliers and use investments required to offer the service. Most of these costs are overheads and hard to be broken down to costs "caused by" single services.

Overheads are of significant relevance for determining TCO. Platform costs (libraries, languages) and infrastructure costs (virtualization, servers, storage, networking, electricity, licensing) belong to the main recurring cost components. Research about operating hyper-scale data centres outlines significant percentages of roughly 90 per cent for data centre operating costs (Miller 2011). Therefore, overheads have to be

allocated accurately. Typically, suppliers will apportion them across services by some suitable proxy metric, e.g. revenues, incident numbers, level of changes.

Another study by Hamilton (2011) about reviewing costs assumptions for an AWS data centre (eight megawatt), which could include 46,000 servers, outlines monthly operating costs for a facility of that size. The main cost drivers are the servers (57 per cent), followed by power and cooling (18 per cent) and electric power (13 per cent). Ultimately, he assumes total operating cost at around 88 million USD per month (about 11 million USD per megawatt).

In order to produce a precise quote, information providers need the required resilience level of the storage, the requested or supplied CPU generation, actual volume (critical for volume discounting) or hardware lifetime. Please note that cloud cost is a TCO offering which includes much redundancy, maintenance and hosting .

Additionally, there are some peculiar specifics for CERN:
- Subsidised internet connection
- Highly competitive electricity rates

Technical specifications:
A) **Private cloud (OpenStack)**: A virtual node with 2 cores, 8 GB memory, 80 GB local disk space (one 8th of a node)
B) **Zenodo service**: 10 VMs, 1 PB of disk space (Dropbox-like), 1.5 PB of Archive Storage

These were the figures we benchmarked.

Based on the volume discounts (up to 45 per cent are granted by the HN suppliers), we calculated with the given volumes. For AWS (S3 Normal Storage), there is no volume discount offered by AWS. In order to protect confidentiality, we always took the most cost-efficient offer from the HN suppliers.

Findings for:

A)   Private cloud (OpenStack)

In contrast to CERN in-house	AWS	Helix Nebula
Open Stack IaaS Service 1 year commitment	+ 19.4 times	+ 5.6 times
Open Stack IaaS Service 3 year commitment	+ 14.6 times	+ 3.9 times

+   = more expensive   - = more cost-efficient

The most cost-efficient offer for a one year commitment that we received from the HN suppliers is approximately 3.4 times as cost-efficient as AWS (3.8 with a three year commitment). Again, we have to take into account that we did not have access to an exact cost assignment of overheads at CERN. But in any case, HN suppliers are more than just a reasonable alternative, especially given the qualitative adoption criteria as outlined in chapter five.

A three year commitment at HN represents a price reduction of around 31 per cent compared to a one year contract price, AWS respectively grants 25 per cent.

Some HN suppliers use other solutions than OpenStack for their IaaS calculation, e.g. VMware vCloud (DSI vCloud).

Benchmark with self-construction/ on premise deployment

What if an organisation wants to set up such a completely new infrastructure on site? It would be interesting to know how much a totally newly installed on premise equivalent costs. An example would be:
One petabyte in self-construction: 8 Backblaze Storage Pods with 4 GB Seagate Drives under ZFS RaidZ3 configuration:

8 Backblaze Storage-Pods		
(www.45drive.com)		= 8 x 6488 = 51904 USD
8 x 45 Seagate ST4000DM000 4 TB Drives		
(www.amazon.com)		= 8 x 45 x 155 = 55800 USD
		= 107704 USD
		= **approximately 79706 EUR**

This brief calculation features the costs pertaining to acquisition but not the operational costs such as electricity. A comparable on premise deployment needs sixty KW per petabyte (https://archive.org/web/petabox.php). Depreciation, interest, housing and FTE costs are other components to be taken into account, just to name a few. A complete cost assessment questionnaire has to include total budgets relating staff, supporting contracts and investment, following the mantra: People delivering services through working hours, using support from outside suppliers and using resources required to offer the service. Hence, just comparing investment costs with TCO is not possible. The outlined calculation should just provide a rough idea of the investment level needed to set up such an on premise deployment.

This deployment would not have the same features of HN suppliers or AWS, e.g. disaster recovery.

This investment exceeds hugely the annual costs for CERN's in-house spending, respectively for cloud services. This is not surprising. However, it may be doubted that small or medium-sized organisations are able and willing to invest such amount of resources. HN is more than just a reasonable alternative.

B) <u>Zenodo service</u>

In contrast to CERN in-house	AWS	Helix Nebula
Zenodo SaaS Service 1 year commitment	+ 30%	- 13%
Zenodo SaaS Service 3 year commitment	+ 30%	- 36%

+ = more expensive   - = more cost-efficient

In this case, the most cost-efficient offer with a one year commitment received from the HN suppliers is 33 per cent more cost-efficient than AWS (51% with a three year commitment). AWS offers the same conditions regardless of the contract period.

Here, a three year commitment at HN represents a price reduction of around 27 per cent compared to a one year contract price, AWS respectively grants no discount.

General Issues

The service levels are an important issue. If they are missing or not defined precisely, it is hard to gain a meaningful comparison with providers. This relates again to the culture of an organisation. If employees are expected to answer inquiries on their free time, and free of charge how can an external service compete with that? To offer an equivalent cloud service, high service levels with reasonable personnel resources must be guaranteed, in case they are required. This is by far not always the case which is important to mention here. Of course, not every organisation needs a platinum service for everything all the time.

The whole calculation is a lot more beneficial for organisations with less on premise infrastructure already in usage. This includes associated researchers with limited or no infrastructure at all. They sometimes only need IT resources for a short period of time. In this case, there is no reasonable alternative to cloud services.

To summarise this chapter, HN supplier are definitely able to compete with AWS. This could have been surprising for those who thought that AWS is the market leader concerning competeveness. Therefore, the initiative is more than just another alternative within the cloud market, especially for research organisations. For example, it is not possible to connect to EGI and Dante via AWS, which makes HN even more attractive. This is a major advantage for HN and therefore should be leveraged to attract more (research) organisations.

Also keep in mind that a single provider would struggle to offer the massive IT capacities the flagships already have. Data generation will increase significantly within the near future. Even if an organisation identifies a provider outside of HN, which is able to offer extensive IT resources, the question could be raised if it ensures equivalent non-monetary benefits. This is further explained in the next chapter.

The combination of having the choice to select the best offer among different leading suppliers via a single interface only, while having the ability to customize it, is unique. Winning additional members means enriching knowledge as well as achieving

economies of scale which is favourable for both sides. Consequently, suppliers can offer even better rates to the demand side.

Even without gathering all cost components, HN prices for the Zenodo service are very competitive. This could be one more indicator for the not (correctly) metered overheads of the OpenStack cloud, because Zenodo (also) offers infrastructure ("Dropbox-like" storage). But as already outlined several times, it was not possible to gather these figures.

5.   Qualitative Criteria for Cloud Computing Adoption

Beyond the quantitative cost assessment, taking also qualitative aspects into account is essential. Aiming for a win-win situation for supply and demand is the ultimate goal.

General/ Economic perspective

The decision for the usage of cloud services hosted by a 3rd party should not be made by creating business cases only. However, this is an important criterion because the benefit is easy to assess (if you have all necessary figures). Determining setup, monthly, migration, internet, switching, support costs while eliminating main IT maintenance are just a few important figures necessary to be able to make a decision for or against the cloud. Replacing capital expenditure with operating with the goal to lower overall IT costs is often the mantra to follow. Especially service providers like to reduce capital expenditure. However, there are some organisations which prefer to have capital on their books, e.g. banks and telecommunication companies.

Effective marketplaces are a key enabler and prerequisite also for the public sector to be able to set up new procurement processes and be able to reap the benefits of competition in a more dynamic way.

There are a lot of not easily quantifiable benefits which come along with using the cloud, e.g. elimination of capacity guessing. Some of them influence each other mutually.

Flexibility

The increased competitiveness is not the only strategic advantage. In today's highly changing world, running experiments often and at low costs is critical to survive. If you fail, fail fast and quick. There is hardly an alternative than sticking to this principle, because no one can anticipate the change in customer needs by not asking customers and constant testing. Hence, cloud computing supports more innovation by letting companies experiment efficiently and at the same time concentrating on their core competencies. On the other hand, internal IT know-how decreases, but that is a general phenomenon of outsourcing activity. Consequently the organisation can fully focus on the core business and take advantage of service provider's competence.

Migrating to the cloud is not an ultimate step of dismissing on premise facility. Both can be combined by running a hybrid model. This is probably the most dominating deployment model of the future. Furthermore, if you do not want to rely on a single provider only, set up or join a multi-cloud deployment like the HN Market Place.

Through defined portability and data availability in the SLA, the vendor lock in risk is rapidly lowered.

Expertise

By taking part in ecosystems like HN, organisations have the unique opportunity to benefit from such emerging cloud value networks which strengthen the level of innovations by bringing a variety of areas of expertise together. The more members, the greater the chance of establishing a sustainable first mover ecosystem.

Variety of Services

Taking advantage of cloud computing, opens up new opportunities such as integration and interoperation of SaaS applications. Again this is a point for concentrating on core competencies since you can choose among a variety of highly specialized SaaS providers with scalable feature modules. The desire to receive a service which is scalable to meet demand is no longer a boundary, given our information and communications technology infrastructure scalability.

Moving from on premise data centres into the cloud is a major shift within an organisation, e.g. not having to handle technical operational inquiries but having more governance involved.

Higher flexibility is also achieved by the following: modifiable contractual parameters at later stages (e.g., cancellation period), various payment options (e.g. payment/billing options) and potentially deciding between application customization through configurability and limited customization. Of course, functional aspects which were already mentioned but strongly relate to that field like scalability, interoperability or modularity of the application should not be forgotten.

Collaboration Potential

This aspect is connected to the one mentioned above. HN promotes inter- and intra-organisational collaboration, through working together with other organisations and within an organisation through universal file access. This supports collaboration of internal and external projects and avoids unnecessary double efforts.

Ease of Use

Frameworks such as ITIL or ITSM which help to manage services to migrate to private, publicly, or hybrid clouds should be evaluated. There is also a desire to receive technical services which reduce unnecessary complexity in the service they provide. Thus, non-pertinent complexity should be hidden, particularly the complexity of managing the underlying IT infrastructure and software.

A visually appealing and sympathetic user interface as well as a user-friendly navigation structure and search functionality which can be adopted easily to user feedback favours maintainability.

The provider has a lot of possibilities to support its customer by offering plenty of features: data reporting and importing, SaaS application's configuration and help functionalities, dashboard with metrics measuring customers' service usage as well as performance monitoring.

Of course, no service provider should frustrate customers by limiting customization possibilities or data import, incompatibility with business processes or new technology, using proprietary technologies, offering insufficient maintenance and performing updates at unfavourable times. This all refers to the degree to which the key functionalities and design features of a SaaS application meet the business requirements of a customer.

High Data Security

Recent incidents in governmental surveillance programs have shown that this could also be seen as a technological shortcoming. But on premise data centres are not excluded from these programs. It is important to stress here that this not a technological but a political and service issue. There is relatively little new technology in the cloud. This issue, which mainly US providers are affected by, could turn into an advantage for Europeans. The geographical legal differences have to be examined at this point. Privacy regulations and data protection laws are much stricter within the European Union. Customer data must be encrypted, protected, back -up, anti-virus protected as well as addressable, traceable and recoverable. That surely includes providing a secure physical environment (e.g. secure data centre) which has to be audited on a regular basis. Standards such as ISO27000 and provisions such as ISAE3402 have to be considered. That is how most businesses assure themselves as to their security and integrity.

Another aspect is data integrity. Can the transferred data be manipulated or accidently modified in internal systems or elsewhere?

Each provider has to secure data confidentiality and has to prevent eavesdropping communications, disclosing sensitive or internal system data. Access control measurements and accountability have to be put in place to track and identify data thefts, access without authorization, insufficient user separation and logging of actions as well as missing logging of actions in internal systems.

So the whole data security issue includes all aspects to ensure that regular and preventive measures (e.g. regular security audits, usage of encryption, anti-virus technology) are taken to avoid unintentional data breaches or corruptions through loss, theft or intrusions.

Most probably the level of data security a professional cloud provider is able to offer exceeds the capabilities of on premise deployments.

High Responsiveness

Cloud computing providers with highly specialised, constantly improving processes, achieve high connectivity and access, assumed that the latency is minimal. The system availability and liability should not be lower than in conventional data centres. Service continuity, sufficient availability of internal systems, higher hardware and software redundancy, prevention of data loss access or data in general are characterized competitive advantages of the cloud compared to on premise systems.

Outsourced computing consists of all aspects of a SaaS provider's ability to ensure that the availability and performance of the SaaS-delivered application (e.g. through professional disaster recovery planning or load balancing) as well as the responsiveness of support staff (e.g., 24‑7 hotline support availability) is guaranteed. An adequate number of service personnel can be dedicated to customer care. But in most cases service personnel is aligned horizontally. Additionally, support should be up-to-date using cutting-edge hardware and software, provided via multichannel. They are in the position to tell their users exactly when services will be performed, so business continuity are ensured.

Specified Reliability and Support

Once again the support issue has to be mentioned here. Provider management contains support and contact information of the provider. This criterion includes all facts regarding support and customer service, e.g. which support is offered and under which conditions. Furthermore, it contains information about the internationality of the

provider and its offered services, e.g. multilingual support, several offices or local contact options. Providers can offer consumer support in different ways, if their consumers need help. Standard electronic support includes Frequently Asked Questions and online documentation. Within individual electronic support, such as e-mail, live-chat or a forum, consumers get answers tailored to their specific needs. Phone and personal visits belong to the group of individual personal support, where consumers get personal and direct contact to a provider's employees. Service management includes all activities necessary to control and manage the obtained cloud services which are subsumed in this criterion, e.g. monitoring of services and volume control via APIs.

In order to get the cloud manageable, consulting and migration support for cloud implementation projects is essential. This is described by the term transformation management.

Reliability is another critical aspect beyond data security. The SLAs have to be transparent to ensure the provision and performing of services at the promised time in a correct manner, e.g. error-free services, appropriate quality certifications (if needed), accurate budgetary controls. Again, the customer has to define what suits his needs best in order to balance the service level against the cost. Fulfilling all contract obligations generates customer trust and strengthens the reputation of the provider.

To sum this up, cloud computing drives efficiency, contingency and replace policy and scalability. Especially the last point is supposed to be one of the most attractive aspects for moving into the cloud. Scaling is easy, fast and seamless which also increases flexibility and cost-efficiency. You will very likely face performance problems at some point when scaling massively with internal systems. The cost and effort of capacity guessing is entirely replaced for the customer. If he needs more cloud resources, he just scales to the amount needed.

By offering a great product, cloud providers will attract more and more customers. As a consequence, they achieve economies of scale.

6. Attractiveness of Helix Nebula

6.1 Supply Side

The cloud computing industry will probably further benefit from on-going outsourcing, not only due to cutting costs but also by focussing on core competencies. There are several advantages that come with taking part in HN for the suppliers:

1) Gaining huge business opportunity
2) Enriching big data know-how
3) Leveraging HN for public relations

1) According to McKinsey ("Winning in the SMB Cloud") the global total public cloud computing market is expected to reach a size of between 40 and 50 billion USD by 2015 which would be around five times the size of 2010. This number outlines the huge growth opportunity of that industry. By joining HN, the involved suppliers are collaborating with organisation having massive data requirements. The go live of the HN Market Place in January 2014 is just the starting point.

A single provider could actually handle the huge data that is produced by the flagships. Most customers seem to want a choice of multiple providers. Thus, the initiative among big cloud providers is the ideal way to balance the risk and benefits. Every supplier can provide services. Hence, according to its expertise, a supplier can strengthen its services or try to diversify by also offering capacities in areas of potential growth. There are already further flagships strongly interested in joining the initiative. It is very likely that more organisations will join HN in the future, which favours again the accomplishment of economies of scale.

Adding the Information as a Service BM in a later stage will increase the revenue potential as outlined in the previous deliverable will be achieved. Therefore the more players on the demand side, the better, because their scientific and economic value is created by merging several different data sets with each other.

The overall ambitious goal of establishing a worldwide all-in-one enterprise cloud for organisations, enterprises and governments should be a strong motivator for more suppliers to join this initiative, given the huge business potential.

The variety of commercial, non-commercial suppliers and flagships with specific needs increases the complexity involved in overcoming the technical challenges. HN is a unique opportunity for suppliers to enrich cloud computing knowledge,

especially concerning big data. There were a lot of obstacles to overcome in order to run the pilots. The gained knowledge is a competitive advantage, which of course can be leveraged for one's corporate goals.

2) They also position themselves as first mover and innovation friendly players in the market. HN represents an attractive alternative to current players in the market such as AWS. The suppliers will also strongly benefit from the remarkable reputation of the participating flagships, which should be leveraged with public relation activities.

6.2   Demand Side

1) Taking advantage of dynamic computing
2) Actively shaping a European counterpart to dominating American providers
3) Simplified and efficient collaboration with other organisations

1) The first advantage for the demand-side members is cloud computing in general. Let us keep the key distinction of that concept in mind which is outlined as follows: "Cloud computing is a model for enabling ubiquitous, convenient, on demand network access to a shared pool of configurable computing resources (e.g. networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction' (Mell and Grance, 2011)." As a consequence they can take full advantage of that increasing trend. They can extend their computing capacities without having high upfront costs and implementation effort as well as need for long-term commitment. When more suppliers achieve economies of scale combined with more (cost-) efficient hardware, prices continue to decrease, which benefits for both sides.

2) Recent incidents concerning data security triggered a global outcry. Even though global data surveillance programs like PRISM and Tempora damage the trustworthiness of cloud computing as a whole, the European providers have a great chance to benefit from it. HN is the perfect stepping stone to position itself as the global centre for reliable and secure cloud computing services.

The flagships are less affected by the US Patriot Act. The disclosure of sensitive data is therefore a notable competitive advantage to their cloud providers such as AWS which should be capitalized.

Taking part in this initiative entices benefiting from the opportunities of cloud computing and at the same time not depending on a single supplier.

Another aspect is aimed at organisations which have been recently founded. Because they do not have to make the shift from on premise computing to cloud computing, they not have to face the cultural change organisations with a long tradition of using own data centres. "Leap-frogging outsourcing" is the term for describing this phenomenon.

It is planned to establish a broker within HN which is supposed to act as intermediary between supply and demand. This includes metering the purchased capacities, sharing generated revenue among suppliers, acting as overall billing institutions and provisioning a trading exchange platform between suppliers to allow resource trading. By having this broker as an intermediate, transparency and consequently credibility of HN is increased.

A broker could also prevent this initiative from being perceived as a cartel, especially through transparency measures and not being involved in offering or purchasing computing services.

3) By joining HN, development of inter-organisational collaboration is fostered. This in turn favours the generation of (scientific) knowledge which is their overall goal. The HN cloud infrastructure is a platform for these organisations to interact and collaborate with each other. The resulting opportunities cannot be foreseen now.

7. Next Steps

To achieve a proper cost comparison of on premise and on demand deployments, all relevant volumes and figures have to be gathered and taken into account. This is the most important obstacle to overcome. The lack of metering and cost detail assessment in this deliverable is causing a not explicit enough outcome. This is obviously easier said than done. Allocating the true cost of a service is very challenging. Organisations often only receive one consolidated invoice for a service a lot of different groups are using. It has to be discussed, if the benefit of accurate costing, which is achieved by metering everything in detail, outweighs the effort of these complex tasks. Another approach is to do rough costing exercises for all services of an organisation, like previously described here with CERN's private cloud and the Zenodo service. If there is a huge gap between the total sum all services cost combined, and the total IT budget, it is an indicator for ignoring significant, relevant cost components. This is reversed, a more practical approach.

With a clear cost/benefit assessment, the sustainability and attractiveness of HN can be showcased, especially when the HN Market Place is running successfully for a while. It also allows you to examine the margin available for a broker between supply and demand.

Furthermore, it would definitely make sense to carry out the same task for small and medium-sized cases. Even the case of a single or small group of researchers with the desire to connect to HN is worth thinking about. By addressing the long-tail of research, even more scientific findings are shared among all members. Especially this target group could be very attractive for the suppliers because they do not operate major data centres like the current flagships. Thus, the requirements and the budgets are not comparable with the existing members at all.

Keep in mind that cost assessments are not about figures only. The strategic and operational perspective has to be taken into account, too.

A powerful and user-friendly broker-service is key to establish a large user uptake and competitive service. The retail market has shown how important it is to have strong global leaders to grow the demand and supplier base. Europe could still take a leading role in this new domain, fostered by its experience in managing fragmented markets. And the political agenda could stimulate the establishment of European cloud broker champions, avoiding an initial momentum created with e.g. the Helix Nebula Initiative would be lost due to defocus. In any case, the solutions have to prevent cartel like behaviour, motivate other suppliers to join the initiative based on transparent criteria and at the same time avoid unnecessary complexity for customers. In this case, sustainability for all members of HN could be proven.