

## Helix Nebula – The Science Cloud

**Title:** D4.2 Cloud Provisioning: case histories of decisions taken

**Editor:** Atos

**Work Package:** WP4 – Cloud platform and provisioning

**Deliverable:** 4.2

**Submission Date:** 24 May 2013

**Distribution:** Public

**Nature:** Deliverable



## Contents

1. Executive Summary .....	4
2. Helix Nebula Marketplace .....	6
3. Helix Nebula Services Description and Service Values .....	14
3.1 Minimal Service Descriptions and Service Values .....	14
3.2 Helix Nebula Service Descriptions and Service Values.....	20
4. Helix Nebula Flagships Services .....	21
4.1 Futuristic Helix Nebula Use Case Scenario .....	21
4.1.1. Requirements Elicited from Helix Nebula futuristic scenario.....	24
4.2 Flagship ESA SuperSites Exploitation Platform Use Case .....	28
4.2.1 Service description and Service value Mapping .....	30
4.2.2. Available Providers.....	31
4.3 Flagship CERN - ATLAS High Energy Physics Use Case .....	31
4.2.1 Service description and Service value Mapping .....	33
4.2.2. Available Providers.....	33
4.4 Flagship EMBL: Genomic Assembly in the Cloud .....	33
4.4.1 Service description and Service value Mapping .....	35
4.4.2. Available Providers.....	36
4.5 New Flagships.....	36
5. Helix Nebula Generic Service Acquisition Model.....	41
6. Helix Nebula Marketplace Catalogue.....	46
7. Conclusion and Next Steps.....	47
8. Appendix Literature list.....	48
9. Appendix Services descriptions per supplier.....	49

### Revision history table

Issue	Date	Description	Author/Partner
V0.8	19 March 2013	Initial draft	<p>Ana Juan Ferrer (Atos) Michel van Adrichem (Atos) Josep Martrat (Atos) Peter Pruijssers (Atos) Mick Symonds (Atos) Sjaak Theelen (Atos)</p> <p>Robert Jenkins (CloudSigma) Peter Gray (CloudSigma)</p> <p>Bernd Schirpke (T-Systems) Jurry de la Mar (T-Systems)</p> <p>With specific contributions from: Jonathan Graham (Interoute) Paul Parsons (The Server Labs)</p>
V0.9	8 April 2013	<p>Second draft:</p> <ul style="list-style-type: none"> <li>3.1 Minimal Service Descriptions and Service Values</li> <li>Helix Nebula Service Descriptions and Service Values <ul style="list-style-type: none"> <li>Atos</li> <li>Interoute</li> </ul> </li> </ul>	<p>Ana Juan Ferrer (Atos) Peter Pruijssers (Atos) Michel van Adrichem (Atos) Sjaak Theelen (Atos)</p> <p>With contributions from: Jonathan Graham (Interoute)</p>
V1.0	11 April 2013	<p>Final version:</p> <ul style="list-style-type: none"> <li>Revision table</li> <li>Executive summary</li> </ul>	<p>Ana Juan Ferrer (Atos) Sjaak Theelen (Atos)</p>
V1.1	19 April 2013	<p>Updated final version:</p> <ul style="list-style-type: none"> <li>Moved supplier's service descriptions from section 3 to an appendix</li> <li>Reworded some sections</li> </ul>	<p>Ana Juan Ferrer (Atos) Michel van Adrichem (Atos)</p>
V1.2	24 May 2013	Updated CERN requirements	Michel van Adrichem (Atos)

## 1. Executive Summary

Helix Nebula is envisioned as a market place where European research institutions (Demand side) access, order and use Cloud-based infrastructure services from a diversity of Cloud suppliers (Supply side). Currently, it is a relatively simple M-to-N structure where the Demand side is represented by three major research organisations, and the Supply side consists of four ‘Infrastructure as a Service’ (IaaS) providers.

This “D4.2 Cloud Provisioning: case histories of decisions taken” report focuses on the processes followed in order to map the previously mentioned M-to-N structure, Demand side requirements and Supply side offerings, thus far in the project to document it for further developments. It presents methods implemented for matching requirements to possible supply combinations in each flagship use-case and the workflow for “on-boarding” a use case onto a required set of services.

In order to do so, section 2 introduces the Helix Nebula Marketplace, defining its starting position and initial results derived from the current use case flagships Proof-of-Concept implementations, decisions taken for the design of the Helix Nebula Blue Box, as well as explored mechanisms for coordination of demand and supply side.

Section 3 focuses on the supply side, describing Services offered. It presents service classes and Service level packages in general. In addition to the general service descriptions of section 3, detailed services descriptions per supplier are presented in the Appendix, “Services descriptions per supplier”. These are presented to provide sufficient information for the services mapping presented in section 4. The detailed services descriptions are presented for five existing suppliers in the Helix Nebula ecosystem: Atos, T-Systems, CloudSigma, Interoute and The Server Labs.

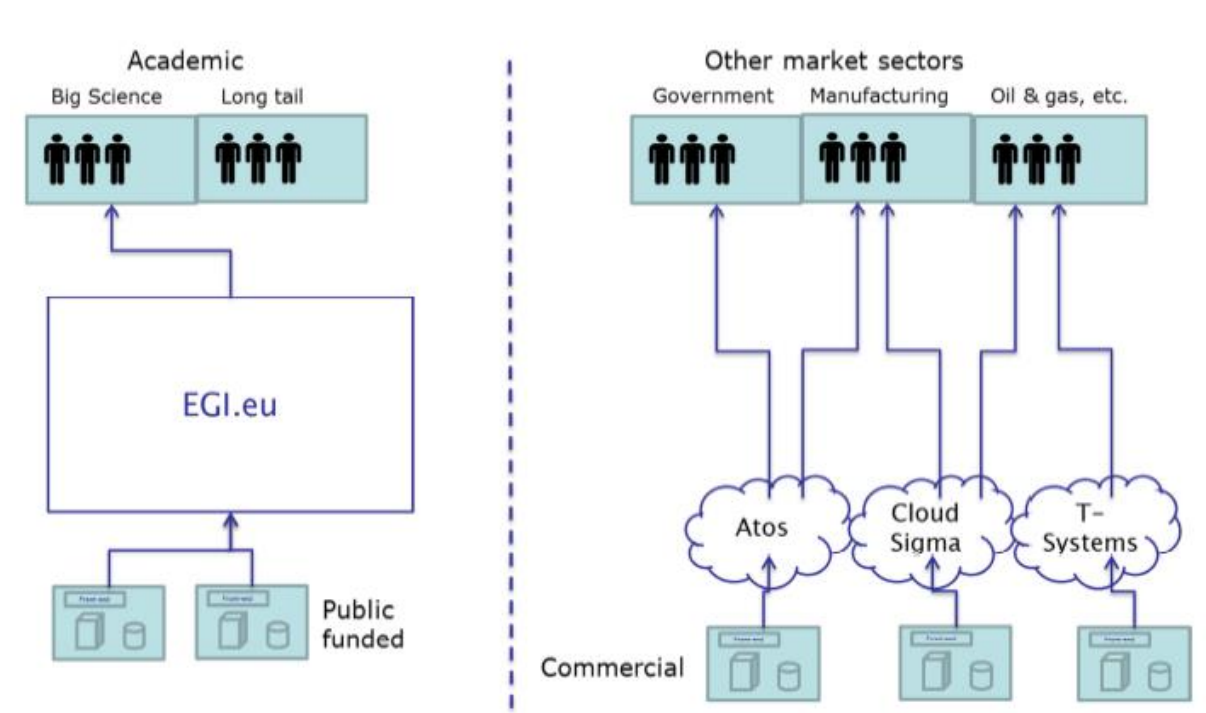
Section 4 aims to describe demand-side requirements and service requests. It describes an overall scenario with envisioned future capabilities for Helix Nebula beyond current offerings. The current situation is presented as a summary of the three executed flagships, identifying service classes and service value requirements, and identifying the suppliers’ ability to provision the required services. This exercise also includes the candidate flagships in order to present a complete and up-to-date overview of the mapping between demand side requirements and supply side ability to deliver.

Section 5 describes the generic service acquisition model for Helix Nebula. Section 6 consolidates the results of section 4 by presenting the Helix Nebula marketplace catalogue, the consolidated view of current use cases (flagships) and available services.

## 2. Helix Nebula Marketplace

### The starting position of Helix Nebula

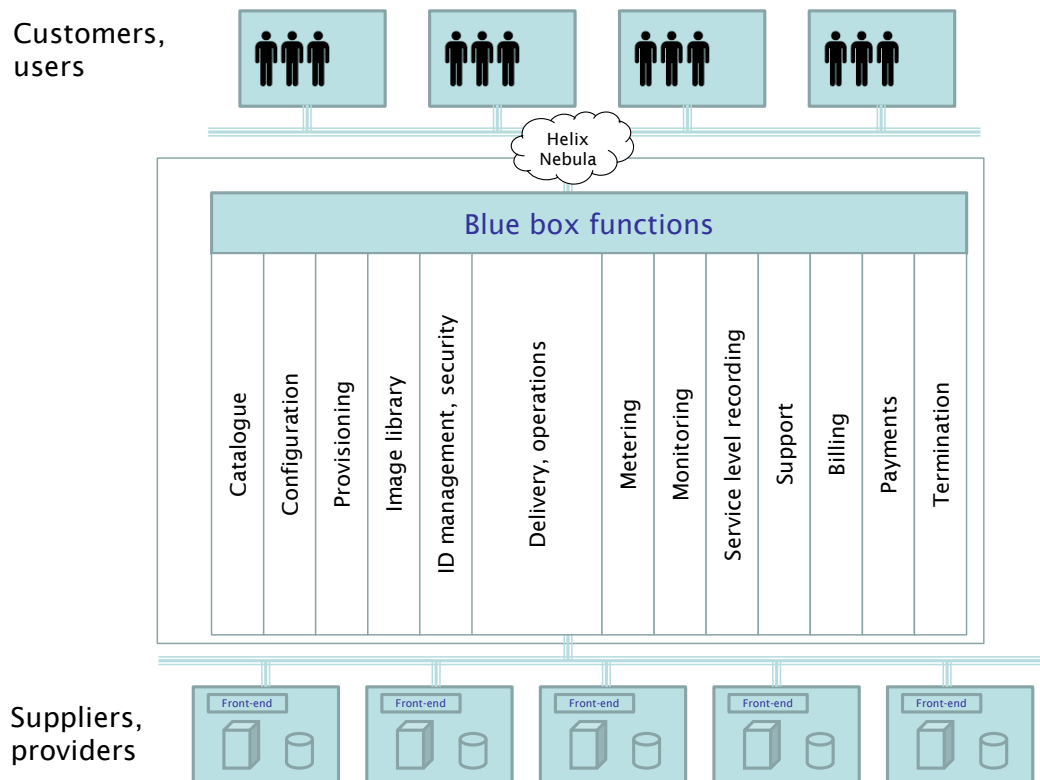
The initial customers for Helix Nebula work mostly on a do-it-yourself basis: they purchase and run their own IT facilities, with some degree of resource pooling via organisations such as the European Grid Infrastructure (EGI). The initially-intended Helix Nebula suppliers, on the other hand, were mainly occupied with delivering services to other markets. So the customers and suppliers acted in separate worlds.



### The initial vision

The initial goal of Helix Nebula is to develop a market place where European research institutes can order and use Cloud based infrastructure services from several suppliers. In its initial development period, Helix Nebula is starting off as a relatively simple M-to-N structure: on the Demand side there are three major research organisations, and the Suppliers consist of a few (currently four) sizable Infrastructure as a Service (IaaS) providers, supported by various other SME's with specialist skills and knowledge. When a single customer uses infrastructure services from more than one supplier, then some sort of coordination between the customer and the suppliers is needed. The TechArch group within Helix Nebula envisioned a "Blue box", positioned in between the customers and suppliers, to handle that coordination. In addition to the technical functionality of the blue box, some service oriented functions were also identified as part of the Blue box. Combining the M-to-N market structure with the coordinating

functions of the Blue box gave the following design of the Helix Nebula market place, with some of the relevant facilities and processes shown, in approximate sequence.

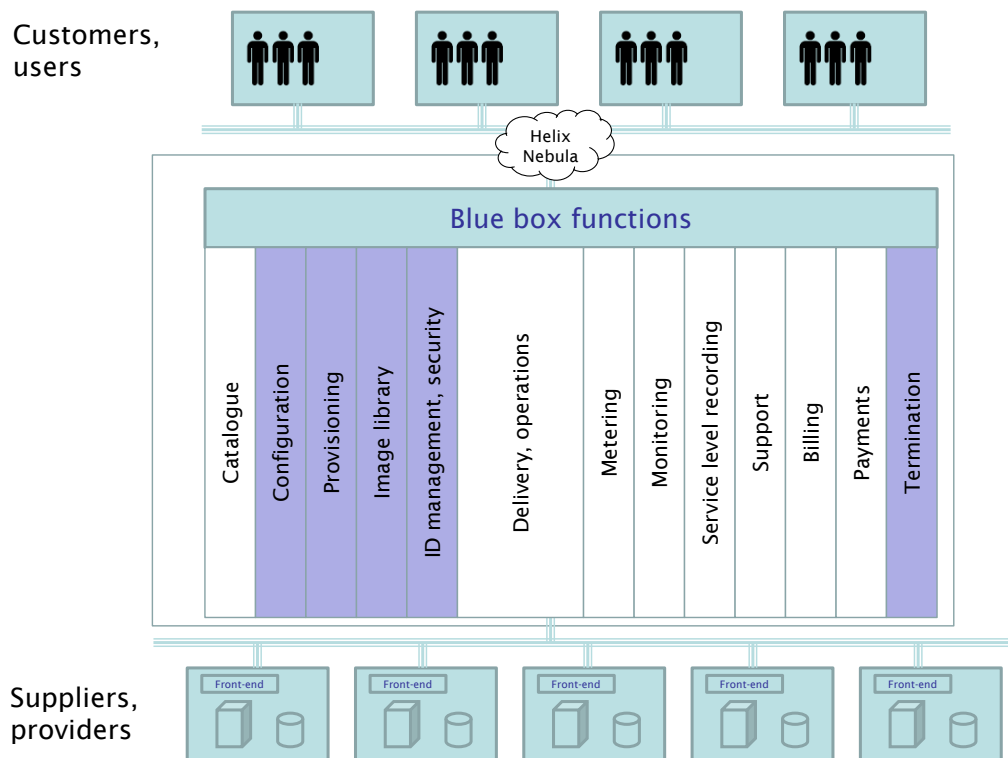


## The initial results

The initial focus of the developments was to determine the technologies needed to establish a working Proof of Concept (PoC). This focus has led to an insight into the technical requirements of the current flagships. Less attention was given to the non-technical requirements: the service requirements and the financial / business requirements. The requirements gathering have also been focussed on the chosen flagships. There has been little consideration for the need to allow for new, and as yet unknown, requirements from future flagships, or for the need to take into account industry-standard requirements for potential future users. The table below shows the current situation regarding the requirements. A green cell signals sufficient insight in the needed requirements, a red cell signals insufficient insight and an orange cell signals the opportunity leaves room for new or industry-standard requirements.

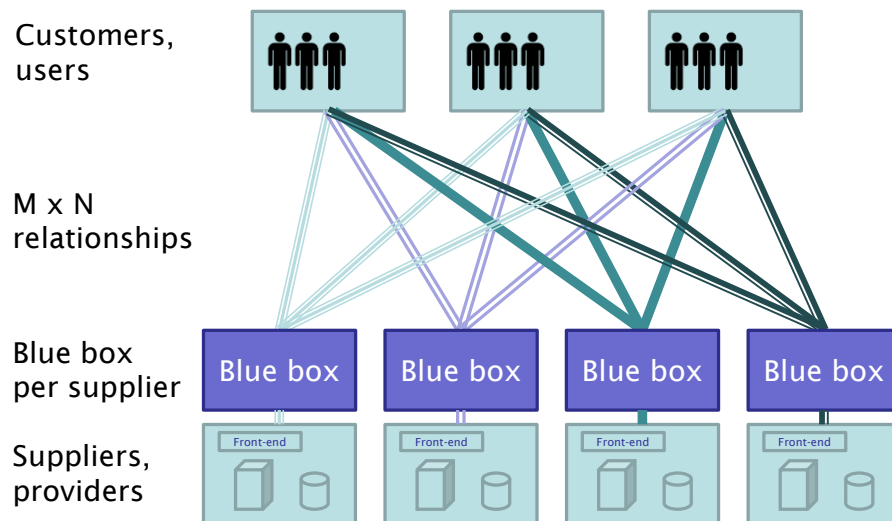
Requirements from	Requirements perspective		
	Technical requirements	Service requirements	Financial / Business requirements
Current flagships	Requirements available	Requirements are needed	Requirements are needed
Future flagships	Need to allow for new requirements	Need to allow for new requirements	Need to allow for new requirements
Potential future users	Need to predict industry-standard requirements (e.g. from ODCA <sup>1</sup> )	Need to predict industry-standard requirements (e.g. from ODCA)	Need to predict industry-standard requirements (e.g. from ODCA)

The focus on the technical requirements and the goal to prove the concept has determined the priorities of the Blue box implementation for the Proof of Concept. The priorities were set as the functions to configure, provision and terminate access for the customer to the needed infrastructure, including its image. So from all functions mentioned in the Blue box design, the functions marked blue in the image below were implemented and tested during the Proof of Concept.



<sup>1</sup> Open Data Center Alliance, see: <http://www.opendatacenteralliance.org/>

The first PoC's showed that the technology worked to allow multiple users to access multiple suppliers' clouds: but that is no real surprise because industry-standard systems were being deployed and used. What was more of a concern was that each of the users had to deploy their workload on each of the suppliers in a different way: there were  $M \times N$  1-to-1 relationships required as depicted in the image below.



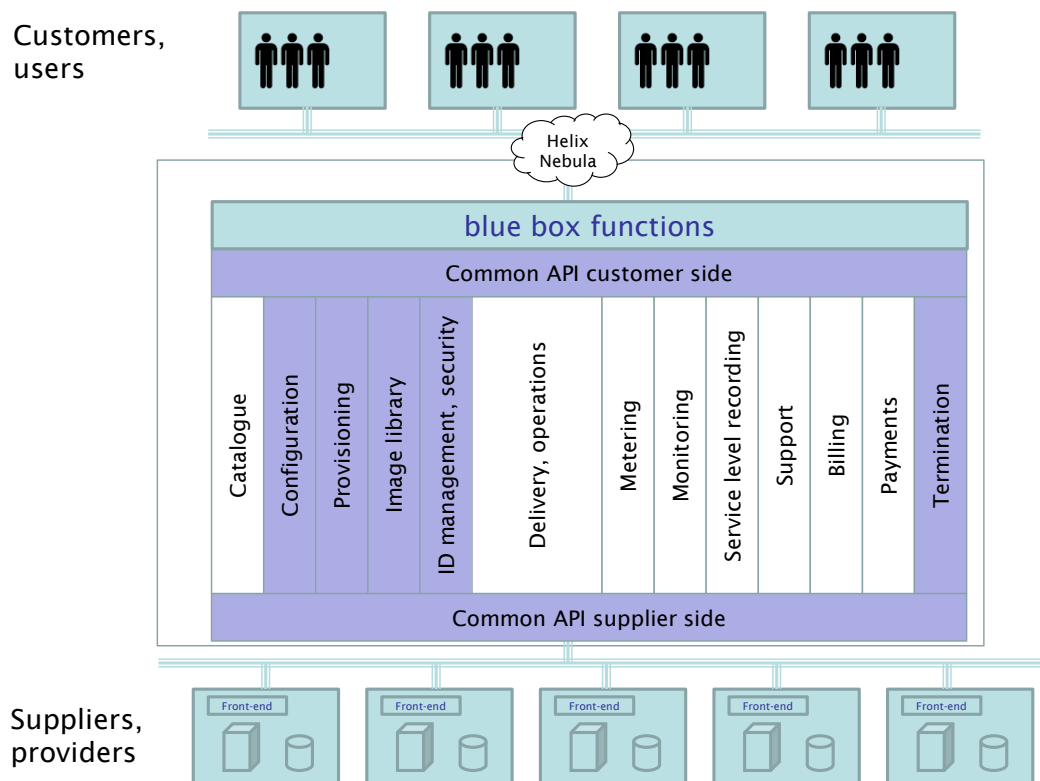
The original intention of the PoC's was to “run each of the Flagship workloads on each of the supplier's clouds”, which was fulfilled. However, it emerged that common ways of interfacing with any supplier's platform, including API's to invoke the services, and standard ways of keeping and transferring virtual machines images, were required if the services were to be successfully federated and used as a whole resource.

### The direction of the future developments

The limitations of the chosen priorities and the impractical  $M \times N$  relationships have led to two decisions:

1. The functional priorities of the Blue box remain as the earlier selected functions. However, the other functionalities will get more attention in the pilots of the current flagships. This will put more attention on the service aspects of the Helix Nebula infrastructure services.
2. Two common API layers will be introduced in the Blue box: the first on the customer side of the Blue box; the second on the supplier side of the Blue box. The common API layers on both sides of the Blue box will support the desired common way of interfacing for customers and suppliers.

The consequences of both decisions led to the following design of the Blue box.



### The service dimension of coordinating demand and supply

The federated nature of the Helix Nebula cloud service means that several suppliers will deliver a part of the cloud service for a flagship to the relevant institute. This leads to a demand to coordinate certain activities between the demand and supply. Where the Blue box fulfils the technical aspects of the coordination, there is also a need to fulfil the service aspects of coordination. For the service aspects various roles, with possible added value to the Helix Nebula partners, can be envisioned:

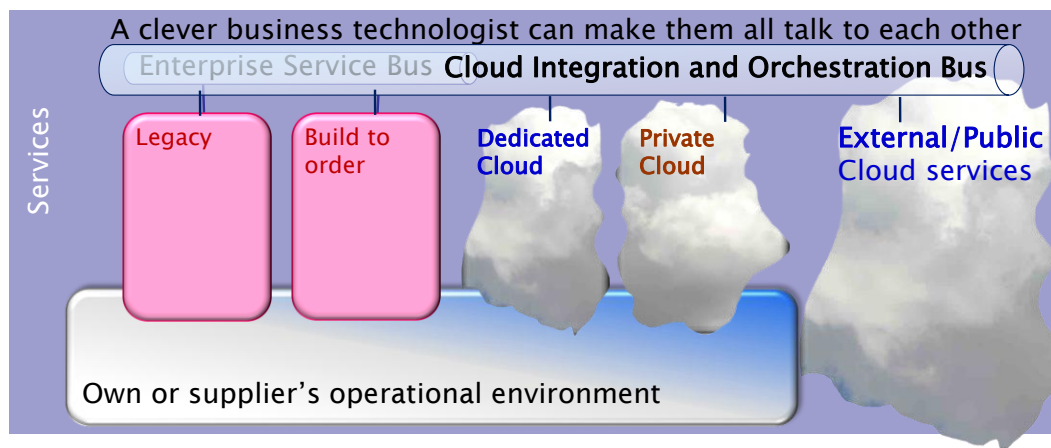
- Broker: finds the best combination of service items and suppliers to deliver the required environment for any particular usage, e.g. the best quality at the lowest price, or the distribution of a large request over multiple suppliers. The broker role may include a payment/finance model.
- (Systems) Integrator: makes multiple new cloud and pre-existing non-cloud systems able to communicate with each other, from a technical and systems perspective, often by means such as the extension of an Enterprise Service Bus (ESB) into the cloud environment.
- Orchestrator: having integrated new facilities and arranged that distributed environments (e.g. servers at multiple suppliers) can find and communicate with each other, this role then coordinates changes to the environment(s) during operations.
- Aggregator: makes multiple running services which are interacting do so seamlessly, from an end-user perspective, so as to deliver an on-going “end-to-

end” service, e.g. with integrated service reporting and billing. This role is also known as Service Integrator.

It will be seen that there is some overlap and the potential for complementarity between these roles, but that they need not be unified. Further, it is possible for multiple such brokers, etc. to coexist. For the ease of reading this document will use the term ‘Broker’ for any and all of these roles.

### Combining Helix Nebula services with existing IT services

The Helix Nebula cloud services are an addition to existing IT services used by the institutes. Combining the (new) Helix Nebula services with the (traditional) IT services requires provisions to integrate and orchestrate the services. An example to integrate several cloud and non-cloud services is given below.



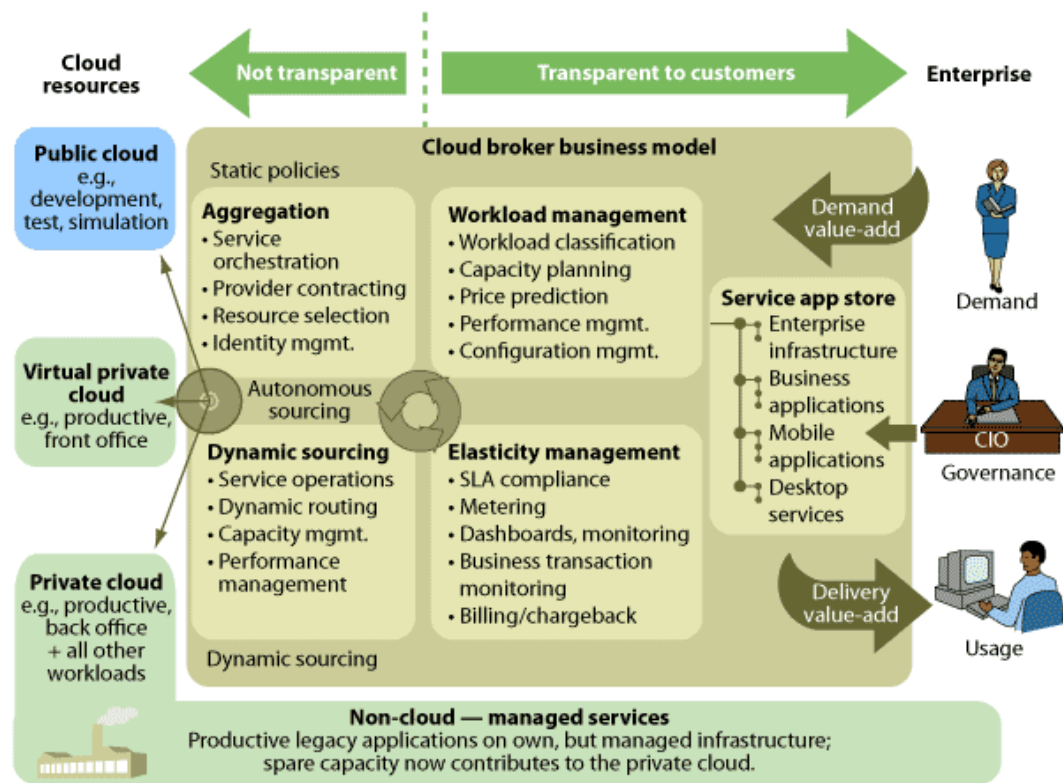
The earlier-mentioned focus on technical requirements makes it premature to launch Value Added Services on Helix Nebula, like a shared Cloud Integration and Orchestration Bus. Such an integration bus could handle information exchange between services running on the platforms of the traditional IT and the multiple Helix Nebula suppliers. This has led to the following decision:

- The focus of Helix Nebula will remain on first providing infrastructure: processing, storage and network connections. Development of Value Added Services like a Cloud Integration and Orchestration Bus will be postponed until the infrastructure services are well established.

### The broker role

The federated nature of the Helix Nebula cloud services introduces coordination activities. In turn these activities need to be executed by one or more partners in the Helix Nebula consortium. This raises the need to assign coordinating roles and

describes their activities, so that the institutes can order and use the Science cloud services as a single and integrated service. The Description of Work mentions the cloud broker role, but also other roles need to be considered. The figure below shows an example of possible areas where a broker or other roles can add value to the Helix Nebula cloud services.



57809

Source: Forrester Research, Inc.

## The initial approach

The do-it-yourself basis that was mostly used by the initial customers reflects on their requirements of the Helix Nebula services. The requirements focus on the ability to order and configure processing capacity, storage space and network connections. The individuals from the demand side that order the Helix Nebula infrastructure services are experienced IT personnel, who are used to determine the required IT solution in detail. This means that currently the user wants to select the required hardware and image from a list of available brands, types and configurations (e.g. processor type and speed). This approach leaves little room for a broker to add value in matching available solutions to service oriented requirements. However, a required activity of the broker, as put forward by the demand side, is to aggregate the billing of an ordered

infrastructure service that is provided by multiple suppliers into one bill for the service. This contrasts with receiving a bill per supplier for part of the service.

With the above mentioned requirements in mind the following decisions concerning the broker role during the pilot phase have been taken:

1. The user selects his or her preferred hardware and image directly from lists of available options in the Blue box. The broker has no active role in matching the available options from the different suppliers to the requirements of the user.
2. The user selects his or her preferred supplier for the delivery of the infrastructure service directly from lists of available suppliers in the Blue box. The broker has no active role in matching the available suppliers to the requirements of the user<sup>2</sup>.
3. The user can define configurations of sets of IT components in the Blue box. For instance one or more web servers that are connected to a database server. In the configuration the user can define the hardware (brand, type and configuration) and image of each IT component, the way they are connected with each other, and possibly the supplier to deliver (part of) that configuration. The broker has no active role in defining configurations or matching suppliers that are able to deliver (part of) that configuration.
4. The broker will aggregate the bills for an infrastructure service that is delivered by more than one supplier into a single bill for the user. The broker will pay each of the suppliers with the payments the broker receives from the user. The terms for this service, including the exact process of billing and settlement, the mitigation of risks for the broker and how the costs of the broker for this service will be covered, need to be determined during the pilot phase.

---

<sup>2</sup> Note that in one version of the Blue box implemented, the user first selects the hardware and then the supplier, as described; in the other, they select the supplier first, then the hardware.

### 3. Helix Nebula Services Description and Service Values

From the sources below service description can be found:

- Helix Nebula: Service Architecture Operational Model [2]
- Helix Nebula Service Properties Catalogue[3]

#### 3.1 Minimal Service Descriptions and Service Values

The following are very minimalistic service descriptions, fitting for the Proof of Concept phase; at the end of this chapter we will suggest more details to be added, to develop towards a more fitting service description for the Operations phase after the PoC phase.

#### Service Classes

The following are the resource-classes currently available through Helix Nebula:

- VM provisioning – includes all services which allow the management of VM lifecycle.
- VM image transformation – transforming a VM image from one format to another<sup>3</sup>.
- Compute capacity (processing) – includes all services which perform operations on data.
- Storage – includes all services which store and retrieve data.
- Network – includes all services which transfer data from one storage or memory resource to another.
- Random-Access Memory

#### Service level packages

In order to simplify the negotiation and management of service delivery, service level packages are defined as an additional abstraction layer. Service level packages group together ranges of service level values for multiple attributes into a single package. The following are some examples:

- Gold reliability
  - o Availability above 99.99%
  - o MTBF > 2 months
  - o Data durability 10<sup>9</sup>/Year (in 1 GB, 1 bytes lost in 1 year)
  - o Support response < 2 hours
  - o Incident response < 2 hours
- Bronze reliability

---

<sup>3</sup> At present, the service is to use “recipes” to build user-specific images from local basic images.

- Availability > 98%
  - No other guarantees
- Gold security
  - ISO 27000 family, PCI-DSS certified
  - Incident response < 1 hour
  - All crypto and key mgmt. modules FIPS 3+
  - RPO 1 hour
  - RTO 10 minutes
- Network gold
  - Bandwidth 10 GB
  - Latency 10 ms

## VM Provisioning

VM provisioning consists of full lifecycle services: ordering, activating, stopping, re-starting and releasing (multiple) VM's; with choices in availability configurations and with supporting services like:

- ability to install software/services, giving access, seeing usage via easy-to-use command-line tools for system administrators and users
- Marketplace with Base OS images available for operating systems such as CentOS 6, Fedora, OpenSuSE, Windows
- Marketplace allowing user communities to create, search for and share customized digitally-signed virtual machine images
- A virtual machine manager, based on Open Nebula

Service options like:

- Gold reliability [6] is the ability of a system or component to perform its required functions under stated conditions for a specified period of time.
  - Availability above 99.99% [6] is the degree to which a system or component is operational and accessible when required for use.
  - MTBF > 2 months
  - Data durability  $10^{-9}$ /Year (in 1 GB, 1 bytes lost in 1 year) [6] is defined as the probability of data loss
  - Incident response < 2 hours, time to start of support after being notified of a service disturbance
  - Support response < 2 hours, time to start of support after being notified
- Bronze reliability
  - Availability > 98%
  - No other guarantees

- Gold security
  - ISO 27000, Wikipedia: ISO/IEC 27000 is an international standard entitled: Information technology — Security techniques — Information security management systems — Overview and vocabulary.
  - PCI-DSS certified, Wikipedia: PCI DSS is the Payment Card Industry Data Security Standard (PCI DSS), [which] is a proprietary information security standard for organizations that handle cardholder information for the major debit, credit, prepaid, e-purse, ATM, and POS cards.
  - Security Incident response < 1 hour, response to critical security incidents is set to within 1 hour.
  - All crypto and key mgmt. modules FIPS 3+, where FIPS is standard for encryption.
  - RPO 1 hour, with RPO defined as a “recovery point objective”, defined by business continuity planning. It is the maximum tolerable period for which data might be lost from an IT service due to a major incident.
  - RTO 10 minutes, The recovery time objective (RTO) is the duration of time and a service level within which a business process must be restored after a disaster (or disruption), in order to avoid unacceptable consequences associated with a break in business continuity.

[3] Volatile demand gold, setting thresholds for the [6] elasticity in provisioning

- Max provisioning rate VMs, 2000/hour, so the maximum volume of virtual machines per hour will be 2000
- Max VM's provisionable 50,000. The maximum of virtual machines that can be provisioned is 50,000
- The VM provisioning availability up to limit, 99.99%, meaning infrastructure to request and deliver additional VM's is available 99.99%

Service attributes that can be chosen per VM:

- Compute capacity (processing) – Adding, changing, releasing # CPU's, within the limits of the hardware configuration
- Random-Access Memory– Adding, changing, releasing # Ram memory, within the limits of the hardware configuration

### Storage, backup and replication

It includes all services which store, backup and retrieve data for the full service lifecycle: ordering, activating, stopping, re-activating and releasing

Service options like:

- Gold reliability
  - Availability above 99.99%
  - MTBF > 2 months
  - Data durability  $10^{-9}$ /Year (in 1 GB, 1 bytes lost in 1 year)
  - Support response < 2 hours
  - Incident response < 2 hours
  - Bronze reliability
  - Availability > 98%
  - No other guarantees
- [3] Volatile demand gold, setting thresholds for the [6] elasticity in provisioning
  - Max provisioning rate storage, 10TB/hour
- Gold security
  - ISO 27000, PCI-DSS certified
  - Security Incident response < 1 hour
  - All crypto and key mgmt. modules FIPS 3+
  - RPO 1 hour
  - RTO 10 minutes

### VM image transformation

This is transforming a VM image from one VM hypervisor format to another, making it possible to move images to other infrastructure. The OVF Open Virtualization Format makes it possible for customers to migrate more easily, independent from hypervisor or processor, within a vendor or even to another (so avoiding lock-in).

- Transformation within x hour.

### Network

It includes all services which transfer data from one storage or memory resource to another, including Network access to and between cloud services for the full service lifecycle: ordering, activating, stopping, re-activating and releasing.

Service options like:

- Gold option[3]
  - Bandwidth 10 GB, reserved maximum available network usage capacity
  - Latency 10 ms, network delay from point of entrance of supplier network towards service component
- Gold reliability

- Availability above 99.99%
- MTBF > 2 months
- Data durability  $10^{-9}$ /Year (in 1 GB, 1 bytes lost in 1 year)
- Support response < 2 hours
- Incident response < 2 hours
- Bronze reliability
  - Availability > 98%
  - No other guarantees
- Gold security
  - ISO 27000, PCI-DSS certified
  - Security Incident response < 1 hour
  - All crypto and key mgmt. modules FIPS 3+
  - RPO 1 hour
  - RTO 10 minutes

## Service Description

This document works towards more formal Service Descriptions (SDs), so these should be on a more detail level, stating exactly what can be expected. The SDs include a generic part and specific details per service.

Within Cloud both human and machine readable SDs will be needed. For machine-readable SDs, several research project have contributed to developments in this area. One example is SLA@SOI where a proposal for Service Description was offered.

Generic SD content:

- Opening hours; availability of support for the service
- Incident resolution response and associated priority levels
- Customer responsibilities
- Vendor responsibilities
- Problem management
- Innovation/Release/Change management, for customer and vendor
- Service level reporting
- Compliance management (Security monitoring)
- Configuration/usage monitoring
- Elasticity
- Volumes, pricing
- Billing, payment
- User management

## Continuity management

### Service specifics

#### Processing Service:

Capacity variants: in # CPU's, # allocated memory in Gb, speed

Availability

Automated requests; for adding, changing, removing processing or processing service options

#### Storage Service:

Type of storage; e.g. SSD/Hard disk based storage systems

Availability

Backup and restore

Automated requests; for adding, changing, removing storage or storage service options

#### App store:

OS types and versions;

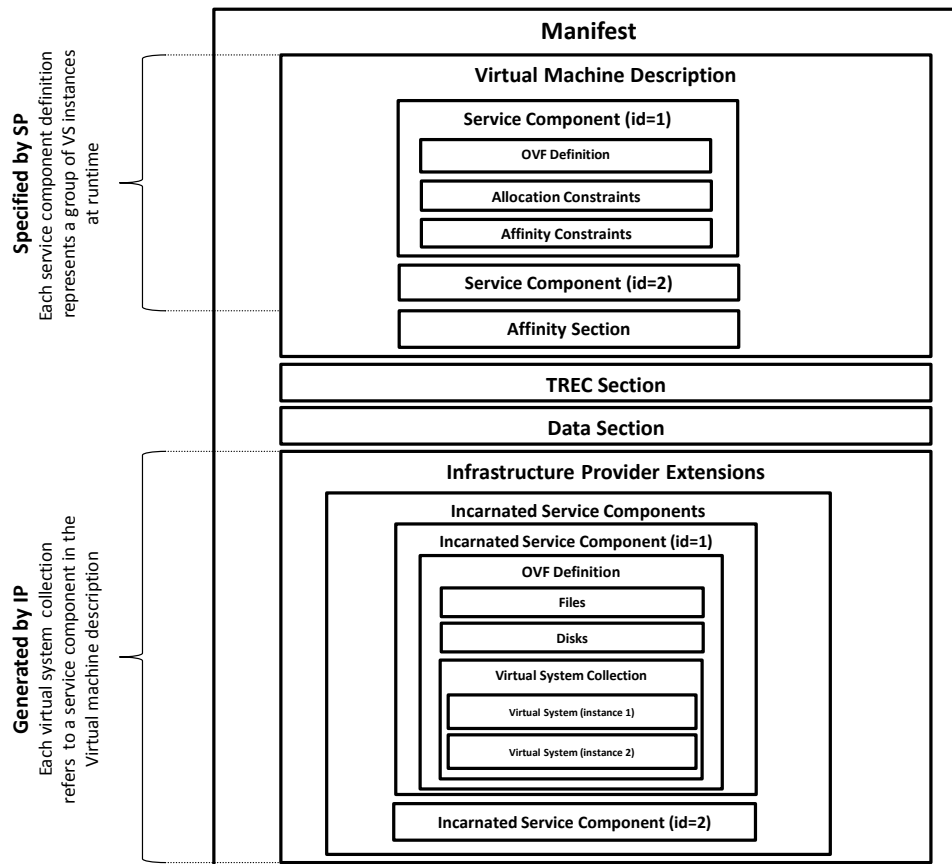
Other software images

Default software, e.g. monitoring & anti-virus software

Automated requests; for adding, changing, removing services or service options

Other projects such as OPTIMIS rely on the concept of a Service Manifest to describe the requested service in detail:

- Provides a Component Model (composed service)
- Multiple sections
- Information only used in the service provider domain,
  - e.g. SW license tokens for license-protected applications
- Information only used in the infrastructure provider domain, e.g. OVF definition of files, disks
- Shared information (passed from SP to IP), e.g.
- Trust, Risk, Eco-efficiency and Cost values
- VM description
- Elasticity rules
- Rules for data protection and data security



### 3.2 Helix Nebula Service Descriptions and Service Values

The appendix “Service description per supplier” presents in detail some of the Cloud Services offered by suppliers in Helix Nebula. The ability of the suppliers to fulfill the requirements of the flagships, which is presented in section 4, has been determined on the basis of the descriptions presented in the appendix; because of the length and level of detail of the descriptions they have been placed in the appendix. The detailed services descriptions are presented for the five existing suppliers in the Helix Nebula ecosystem: Atos, T-Systems, CloudSigma, Interoute and The Server Labs.

## 4. Helix Nebula Flagships Services

This section presents Helix Nebula Flagships Services, using a twofold methodology: First, in section 4.1, Futuristic Helix Nebula Use Case scenario, views on a potential approach for future developments of Helix Nebula are provided, in order to elaborate on long term requirements and capabilities.

After this, in sections 4.2 to 4.5, the concrete requirements and services requested by the three Flagship Proof-of-Concepts developed so far, together with initial requirements extracted from the next wave of candidate flagships, are presented.

### 4.1 Futuristic Helix Nebula Use Case Scenario

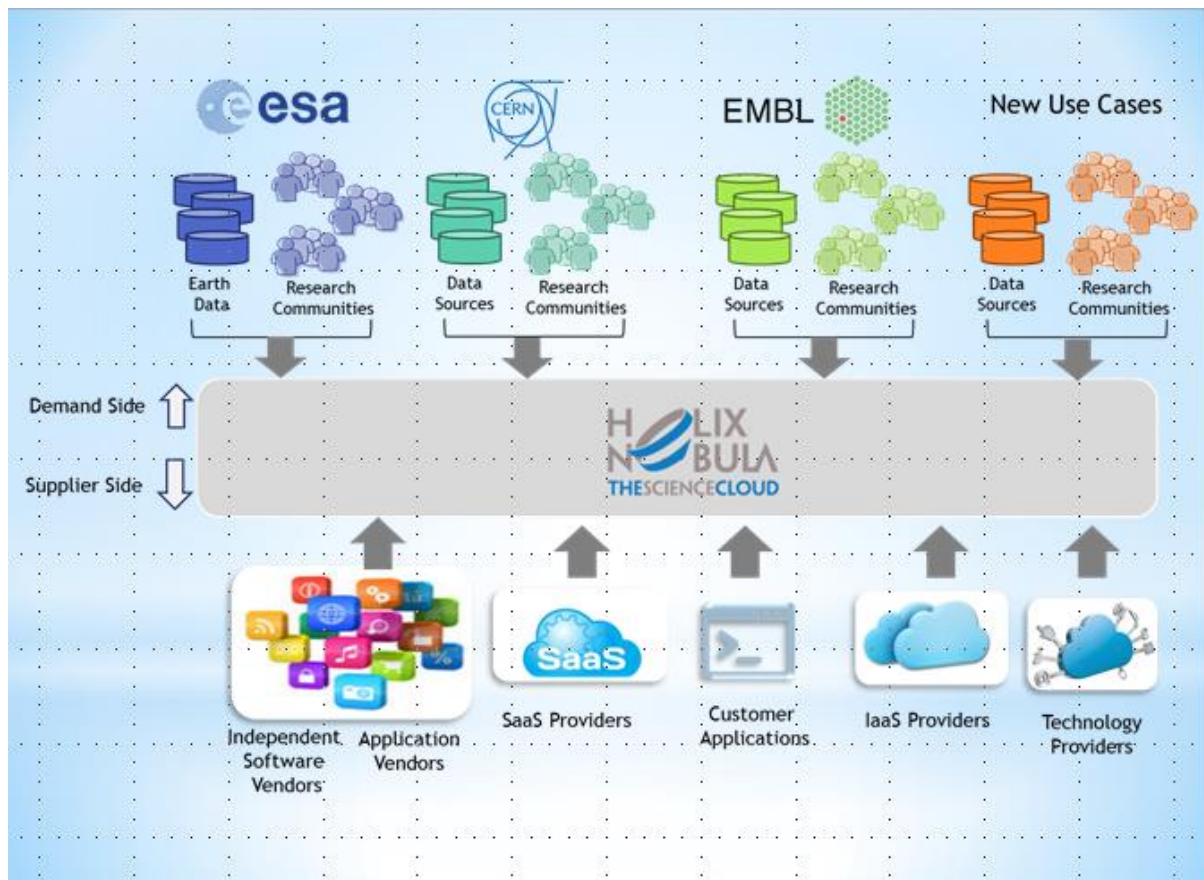
This section presents an overall Helix Nebula scenario that goes further beyond currently-existing Proofs of Concept. The intention is to present a potential approach for usage from which to extract additional requirements and capabilities necessary for future developments of Helix Nebula.

In current Proofs of Concept, Helix Nebula can be compared to a virtual data center. Helix Nebula is, at its current stage, an advanced virtual data center with image / application store, virtual network set-up capabilities and IT component profiling capabilities. The results so far (PoC) are that the advanced capabilities only work within a single supplier. Capabilities such as image exchange between suppliers, virtual networks set-up spanning multiple suppliers, and deployment of profile of an IT solution containing multiple components (e.g. webserver and databases) distributed over multiple suppliers, were not supported.

The decision taken is to extend the current Blue box capabilities for the pilot phase so the capabilities include:

- Make exchange of images possible, so an image created at one supplier can be used for an IT component that is deployed at another supplier
- Make it possible that virtual networks can span multiple suppliers
- Make it possible that IT solutions containing multiple IT components can be deployed in a distributed way over multiple suppliers.

The futuristic vision for Helix Nebula is to go beyond pure infrastructure provisioning, but to develop strategies for large and highly distributed heterogeneous scientific data, and a marketplace of applications and processing infrastructure able to exploit this data deluge.



**Figure 1 Envisaged Helix Nebula Ecosystem Actors**

Helix Nebula aims to become a research and business ecosystem enabled through a combination of capabilities, data and services from several actors. By enabling research communities to access, analyse and curate diverse and enormous data sources, science is expected to deliver significant progress and innovation. While, at the same time, innovative business models can appear out of this data explosion or through provision of application or infrastructure services. In order to enable this futuristic scenario, new actors need to be introduced in the Helix Nebula ecosystem, with specific emphasis on supplier side. Identified actors at this stage are the following:

- Independent Software Vendors / Application vendors: Companies whose primary function is produce and distribute software products. Under this category, also open source application communities are considered.
- Software as a Service (SaaS) Providers: Similar as above, but in this case the application is delivered in a pay-per-use, on-demand model.
- Customer Applications: Under this tag, applications owned by Use Case organisations are represented. These applications can be shared at different

levels: to the entire Helix Nebula ecosystem (then becoming another provider) or internally to a single use case.

- IaaS Providers: Providers that offer infrastructure resources in an on-demand pay-per-use model, including storage, processing and networking components.
- Technology/Solution Providers: Technology / Solutions providers offer products to be used to define the technology stack necessary to build the envisioned Helix Nebula Solution.

The success of Apple's iTunes Application Store and Android's Play in the mobile world has shown the enormous possibilities application marketplaces offer. The combination of well-designed, integrated and easy-to-use marketplaces, together with an effective payment system based on micro-transactions, and a broad selection of interesting content, delivers a very effective distribution channel for services provided by IT service providers, as well as SMEs. Application Marketplaces is a new model for the delivery of application and services, and a clear example of how new ecosystems and service models are emerging.

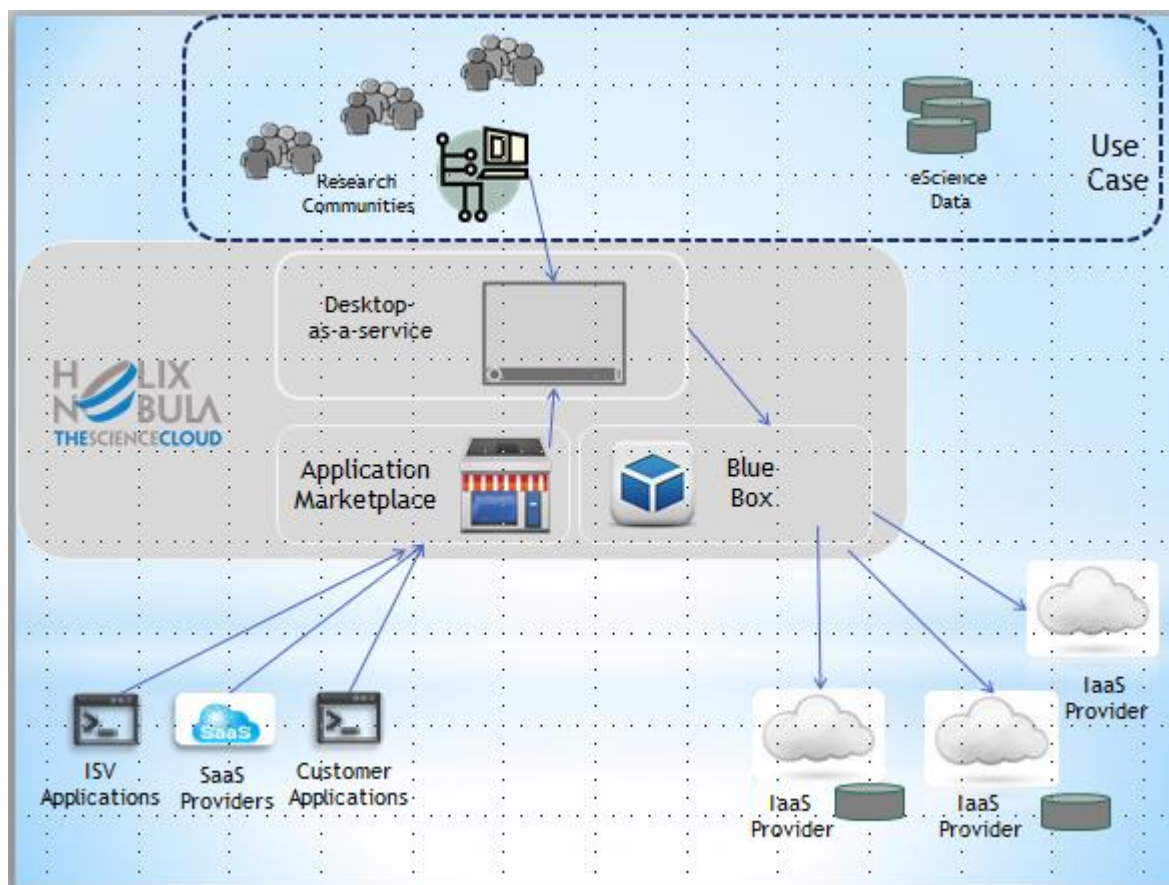
In the context of Helix Nebula, it could be made an analogy to Virtual Desktops. The objective, therefore, is to provide the necessary platform that enables Helix Nebula to create and operate Advanced Virtual Desktops (AVDs) for research communities and scientists, which includes the possibility to access applications offered through the marketplace, access to available data sources and processing power.

There are some basic concepts in Desktop Virtualization technologies to be taken into account:

- VDI (Virtual Desktop Infrastructure) is an infrastructure (hardware, software and networking) supporting Desktops or virtual PCs. With this technology, it is possible to host the desktops in private or public clouds.
- Hosted Virtual Desktop (Cloud Desktop or Desktop as a Service):
  - Virtual PCs are hosted in the cloud. The customer is using a desktop image, all the computation process are realised in the server.
  - Reduction of hardware, more computing capability.
  - Two kinds of desktops: shared and dedicated: In the shared desktop, several users share the same server although they are isolated. In dedicated desktop, the customer has available all the resources.

Figure 2 Overall Scenario represents the envisioned usage scenario. The entry point for research communities' members would be a simple Desktop-as-a-Service facility. From there, researchers would have access to authorized data-sources from a Data-Catalog, making accessible both existing Data Sources in Use case facilities or, through the Helix

Nebula Blue Box, stored in suppliers' facilities. Available Data Sources can present different processing levels, ranging from pure row data to pre-processed data.



**Figure 2 Overall Scenario**

Researchers can access applications in the Application Market place in a simple manner, by means of the Virtual Desktops, similar to the iPhone or Android way of working: the researcher is able to select the applications in their desktop (ISV applications, SaaS providers or even own-developed customer applications). These applications will make use of the Helix Nebula Blue Box, to enable customers to manage their infrastructure, resizing against peaks or valleys, and paying for its use or using a flat fee depending on their needs. Provisioning methods should include automatic provisioning methods of applications and desktops, as well as an integration framework between all the components: applications marketplace, VDI and Blue Box.

#### 4.1.1. Requirements Elicited from Helix Nebula futuristic scenario

### Applications Virtualization:

- Application virtualization concept is important in cloud desktop technologies. When the app is virtualized, it is kept in a repository and it will be delivered when the customer needs it.
- It allows not having to install applications on every desktop, with an important resources saving and simplifying the pay-per-use model.
- The virtualization process of an application could be very complex; it depends on the application.

### VDIs and current approaches

On one hand, there are currently two models for the remote delivery of software and apps:

- Cloud based: software runs in the servers of the software vendor. Users access the apps through an Internet browser via the HTML protocol. Software upgrades are transparent, since they don't require new installations, and benefit all customers at once. Some SaaS providers have developed a catalogue of apps around key products, since it is a provisioning model which fits perfectly with the proposed concepts.
- Virtualization based: apps, licenses for which have been bought from the software vendor, run in the servers of the host. Users access apps via a proprietary protocol and software that takes control of the client desktop. This model is commonly called VDI, standing for Virtual Desktop Interface; main players are Citrix and VMware.

The virtualization based delivery model's success owes to a few key benefits: cost savings in service and support through the central control of software, the opportunity to deploy cheap and energy-friendly thin clients vs. high powered PCs, and the possibility to virtualize any application, even a highly-customized piece of software unique to one company. In the vast majority of cases, the host and the users happen to be the same company.

Solutions such as EyeOS, iCloud or Global Hosted Operating System (G.ho.st) enable the possibility to access web applications, instant messaging tools and office suites remotely in a concrete OS, while remote storage can be used for saving and sharing files with others. Moreover, they offer the possibility to access through mobile devices (such as smartphones) by using light clients. Some of them even provide a development environment for building new applications adapted to the VDI solution.

On the other hand, traditional VDI vendors like Citrix and VMware are moving away from the fully-fledged thin client model. Workforces are more mobile and cannot avoid working offline sometimes. The thin client value proposition don't fare well when users

go offline. For that reason, the latest trends observed in the market point to the virtualization of applications vs. the whole PC, that can run from any PC, such as Citrix XenDesktop and VMware View, that are generic solutions for any OS. So VDIs can be created in any OS (Windows, Linux, etc...) and applications can be selected later, for customizing the final VDI to be delivered.

It is safe to assume though that software vendors will become better at modularizing their offering, helping customers to build customized solutions by combining the appropriate modules. The Cloud based SaaS model will make possible and profitable the delivery of these modules. Companies will find this alternative more and more attractive to substitute for their high maintenance legacy apps. This trend will gain even more momentum in SMEs, where legacy apps are rarely highly customized and complex. The challenge then shifts to the ability of companies and their providers to integrate all these apps under the same roof, so that complexity is removed from the user experience, licenses management is transparent to the company and customer support is centralized. This integration will require high customization capabilities, making it possible to change dynamically one of the selected applications, even if it is better to do it due to the context. For instance, if we use a complex tool for designing a product, probably we will not need so much complexity when accessing from a mobile device, but we will need an application with more visualization options (because of screen limitations).

### Licensing management and other platform services

In the case of using licensed software applications, software should be protected to avoid the usage without license. This is the case for the applications to be deployed in the Advance Virtual Desktop. Current license management can be divided in several categories:

- No license. The software is not protected at all and it is freely installed. This situation is commonly used to spread the product in the market (for example, very old versions of Microsoft products), or for attracting users by providing less-capable versions (for example, some antivirus which have free versions and commercial distributions).
- Short temporary license. The software is provided with a temporal license, usually used for the testing and verification.
- Software key. A key is provided to the final user to be inserted during the installation. Usually, in the last few years, when the key is inserted, the software connects outside using the Internet to confirm the validation of the key.
- Hardware desktop key. Using a dongle or special hardware in the device. In the first case, the software is provided with a dongle which must be present in the desktop to execute the software. In the second, a key is generated specifically for the device using specific hardware or some identification provided by the installed devices such as the network MAC address.

- Distributed license. In this case, the license of the software is on a remote server, usually from the company. When an authorized user demands the usage of the software, the application asks for the license from this remote server. The licenses in the server are usually protected using keys attached to the hardware. This is the case of Flexera software.

In the last few years, due to the limitations of the last model, used mainly in technical software, there were some projects to allow the usage of the software outside the company network dynamically. For example, software solutions allowing the usage of software in Cloud environments. It has been designed to take into account several scenarios, such as the usage inside the company, where policies can be established for the access rights to the licenses, as well as brokering of licenses, elastic management of number of tokens, License as a Service (i.e., providing the license as a service by an ASP), etc. The software allows the companies to control the usage of licenses or provide open services in the Internet market.

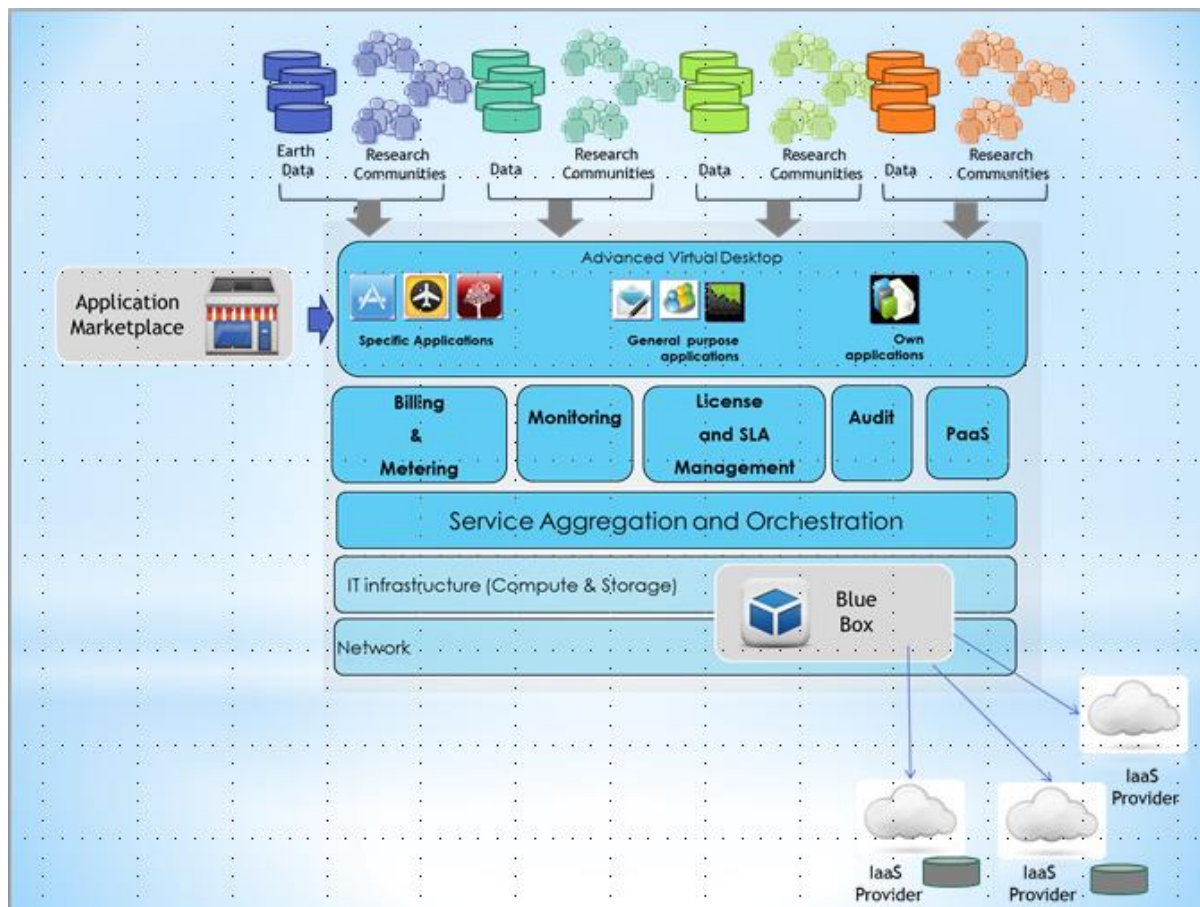
Apart from applications, there are also Digital Rights Management systems, which are more general than the software protection but based on the same principles. They are used mainly for contents such as videos, games and ebooks.

## Cloud computing and virtual networks

The emergence of the Cloud Computing paradigm has highlighted the importance of the research done on virtual network technologies and paradigms. There is an increased need to multiplex several logical usages of a shared physical infrastructure, while still offering solid user isolation, flexibility and performance.

Project like PHOSPHERUS (EU), G-Lambda (Japan) or Enlightened Computing (US) pioneered this area. Currently, projects like MANTYCHORE and GEYSERS are further exploring the interrelation of virtual networks and cloud computing. The GEYSERS project is focusing on the design and implementation of a suitable architecture for resource provisioning of both optical and IT resources (computing or storage). On the other hand, the MANTYCHORE project is implementing a software stack that satisfies advanced Cloud Computing scenarios such as network demarcation, virtual private networks and energy-based scheduling.

Figure 3 Helix Nebula Futuristic Scenario Building blocks represents the capabilities described in this section in building blocks.



**Figure 3 Helix Nebula Futuristic Scenario Building blocks**

## 4.2 Flagship ESA SuperSites Exploitation Platform Use Case

The European Space Agency (ESA), in partnership with the Centre National d'Etudes Spatiales (CNES) in France and the German Aerospace Center (DLR) is collaborating with the National Research Council (CNR) in Italy, to create an Earth observation platform focusing on earthquake and volcano research. Geohazard Supersites is an international partnership of projects, organizations and scientists, involved in monitoring the dynamic and complex solid-Earth system and the assessment of geohazards. Its mission is to support the Group of Earth Observation (GEO) in the effort to reduce loss of life and property caused by the natural disasters of earthquakes and volcano eruptions. The Geohazard Supersites project will advance scientific understanding of the physical processes which control earthquakes and volcanic eruptions as well as those driving tectonics and Earth surface dynamics.

### Use Case Objectives:

- To provide an open source, unified e-infrastructure for solid Earth data,
- Improved data products for solid Earth monitoring,
- Secure solid Earth data sharing on the cloud,
- Enabling international collaboration,
- Enlargement of the science user base.

### Use Case Requirements:

#### *Overview*

- To create an Earth Observation platform, focusing on earthquake and volcano research.
- [Helix Nebula EMBL - ISC12 - Lueck v1.1] Enabling earth data from satellites to be re-used by putting it in the Cloud for scientific research and application and also to enable more commercial applications.
- Expectation to use Cloud for accommodation of peak loads. The scaling frequency will be quite low, varying the number of VMs in matter of days (in relation to medium term processing requirements).

Case 1: Data dissemination: 2VM and 20TB of storage always on

Case 2: Processing campaigns: from 10 to 50VM on for a period from 15days to 3months

Case 3: Sandbox service: 2VMs on during working hours (ad hoc machines for people to use)

[Helix Nebula EMBL - ISC12 - Lueck v1.1] CNR SBAS processing 1,5TB satellite images, 150 hour processing

#### *Sizing*

Processing: Min 5 seconds; Likely 10 seconds; Max 1 minute per VM (and/or VM modification operation)

Storage: Large data transfer will occur only at start-up, to fill the system with the ESA database of EO data (around 20TB). After that time, data transfer will limit themselves to smaller amounts. Incremental updates <1TB month. Enhance? In 2,5Tb blocks. Min 30TB, Likely 200TB, Max 300TB

Network: bandwidth 10 Gb, Jumbo frame, one isolated network with possibility to define more zones

#### *Capacity Estimates*

Infrastructure per request:

Processing: Server type A: Min 10, Likely 50, Maximum 100

Storage: In 2,5Tb blocks. Min 30TB, Likely 200TB, Max 300TB

Network: Min: 1GB, Likely 3GB, Max 10GB

*Other service aspects:*

- API: OCCl, Amazon EC2 preferred
- Web console contextualisation mechanism for VM's, manage firewalls
- Data in member states, No disk access to provider.
- Thorough data wipe after de-provisioning
- Helix-Nebula Single-sign-on or ESA EO SSO single sign on support is preferred (Shibboleth)
- 99,99% availability
- SLA: incident severity1/Blocking: 1 hour (Work around), 3 day solving
- High: 1 day Workaround, 7 days solving
- Medium: 3 day Workaround, 14 days solving
- Low: On agreement
- Helpdesk, Email, FAQ, Client Exposed Ticketing system. Coverage: normal working hours

#### 4.2.1 Service description and Service value Mapping

Service Classes	Service Value
VM Provisioning	Min 5 seconds; Likely 10 seconds; Max 1 minute per VM (and/or VM modification operation)
Compute Capacity	Case 2: Processing campaigns: from 10 to 50VM on for a period from 15days to 3months Case 3: Sandbox service: 2VMs on during working hours (ad hoc machines for people to use) [Helix Nebula EMBL - ISC12 - Lueck v1.1] CNR SBAS processing 1,5TB satellite images, 150 hour processing
Storage	In 2,5Tb blocks. Min 30TB, Likely 200TB, Max 300TB Large data transfer will occur only at start-up, to fill the system with the ESA database of EO data (around 20TB). After that time, data transfer Incremental updates <1TB month. Enhance In 2,5Tb blocks. Min 30TB, Likely 200TB, Max 300TB
Network	bandwidth 10 Gb, Jumbo frame, one isolated network with possibility to define more zones
<b>Identified GAPS</b>	<b>API access</b> <b>Single-sign-on</b> <b>Network Jumbo Frame</b> <b>User Roles</b>

#### 4.2.2. Available Providers

Service Classes	Atos	CloudSigma	T-Systems	Interoute	The Server Labs
VM Provisioning	☑	☑	☑	☒	☒
Compute Capacity	☑	☑	☑	☒	☒
Storage	☑	☑	☑	☒	☒
Network	☐	☐	☐	☒	☒

☑ Covered / ☒ Not covered / ☐ Partially Covered

#### 4.3 Flagship CERN - ATLAS High Energy Physics Use Case

ATLAS is a particle physics experiment in the Large Hadron Collider at CERN, Geneva. The ATLAS detector is searching for new discoveries in the collisions of protons of extraordinarily high energy. ATLAS will learn about the basic forces that have shaped our Universe since the beginning of time and that will determine its fate.

The ATLAS experiment is currently running a large scale distributed computing system to process the data acquired by the experiment. The ATLAS Distributed Computing environment (ADC) consists of several pieces: a Distributed Data Management component (DDM), a distributed production and analysis system (PANDA), and associated tools as well as the processing and analysis binaries. This flagship is part of a wider project within ATLAS, to research the applicability of cloud computing to ATLAS computing.

“CERN’s computing capacity needs to keep-up with the enormous amount of data coming from the Large Hadron Collider and we see Helix Nebula- the Science Cloud as a great way of working with industry to meet this challenge,” says Frédéric Hemmer, head of CERN’s IT department.

##### Use Case Objectives:

- Evaluating the use of cloud technologies for ATLAS data processing,
- Transparent integration of cloud computing resources with ADC software and services,
- Implementation of the ATLAS cloud computing model within DDM, Panda, and related tools and services.

## Use case requirements

### *Overview*

The use of IT for CERN revolves around tasks. Tasks are CPU bound and this is reflected in the requirements.

### *Capacity estimates*

Infrastructure per request:

- Typical VM: 4 cores \* 2-3 GHz CPU / core, 8 GB RAM (= 2 GB RAM per core), 80 GB HDD (=20 GB HDD per core). Scientific Linux OS with the CernVM image.
- Number of VM's per request: Minimum 50 VM's, most likely 250 VM's, maximum 2500 VM's.
- Storage: Local spinning disks with 80 GB capacity per VM suffices. Size per request is minimum 4 TB, most likely 20 TB, and maximum 200 TB. There is no need for shared network disks, SDD, volume storage or object storage.
- Network: Access via GEANT is preferred. The required bandwidth is 1 Gbps per 100 VM's with an estimated peak load of 1 Gbps per 25 VM's. Per request the required bandwidth is minimum 1Gbps, most likely 1 Gbps, maximum 10 Gbps.
- Provision lead time: depends on the number of requested VM's: 3 days for 50 VM's, 10 days for 250 VM's, and 30 days for 2500 VM's. The more dynamic the resource lead times are the better, order of a few days is tolerable.

### *Other service aspects*

- Access method: API via HTTPS with 99% uptime of the API / web console, EC2 tools / OpenStack tools suite (nova, glance).
- Content delivery: API to query machine status and provision / deprovision machines.
- Software deployment via the software deployment infrastructure built into the CernVM image.
- System /network / performance monitoring with Ganglia, ATLAS PanDA monitor
- Single integrated portal presenting the service levels of all suppliers (ideally).
- Support via helpdesk, email or client exposed ticketing system.
- Single Sign On, Roll Based Access Control, Secure remote access via SSH with key length of 2048/4096, key management via API, security related logging, signed timestamping, WORM functionality.
- Data preferably in Europe.
- The contract price shall be net, firm and inclusive of all costs relating to the performance of the contractor's obligations under the contract and take into account CERN's exoneration from VAT and import duties.

#### 4.2.1 Service description and Service value Mapping

Service Classes	Service Value
VM Provisioning	Provisioning lead time depends on number of requested VM's: 3 days for 50 VM's, 10 days for 250 VM's, 30 days for 2500 VM's. The more dynamic the resource lead times are the better, order of a few days is tolerable.
Compute Capacity	Minimum 50 VM's, most likely 250 VM's, maximum 2500 VM's with 4 cores each. A typical machine has 4 cores * 2-3 GHz CPU / core, 8 GB RAM per VM (= 2 GB RAM per core), 80 GB HDD per VM (=20 GB HDD per core). The VM's OS is Scientific Linux with the CernVM image.
Storage	Shared network disk is not required, only local system disk with 20 GB per core is required. Given the compute capacities in a request this leads to: minimum 4 TB, most likely 20 TB, maximum 200 TB of spinning disks (SSD not required, tasks are CPU bound).
Network	Minimum 1Gbps, most likely 1 Gbps, Maximum 10 Gbps (=1 Gbps per 100 VM's). Estimated peak load is 1Gbps per 100 cores (= 1 Gbps per 25 VM's). Connection via GEANT is preferred.
<b>Identified GAPS</b>	<b>API access</b> <b>Single-sign-on</b> <b>Single integrated portal for service levels</b> <b>User Roles</b>

#### 4.2.2. Available Providers

Service Classes	Atos	CloudSigma	T-Systems	Interoute	The Server Labs
VM Provisioning	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Compute Capacity	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Storage	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Network	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

☒ Covered / ☒ Not covered / ☒ Partially Covered

#### 4.4 Flagship EMBL: Genomic Assembly in the Cloud

Next Generation DNA Sequencing technologies have had a huge impact on how biological and medical research is performed today. It has helped large-scale sequencing to become affordable and revolutionized the sequencing of complete genomes. Scientists today easily generate huge amounts of quality sequence data within a few days. The analysis and management of these vast data sets, however, require high

performance computing and fast data storage infrastructures as well as bioinformatics expertise, which are often challenging for many labs.

The European Molecular Biology Laboratory (EMBL) is developing a portal for cloud-supported analysis of large and complex genomes. This will facilitate genomic assembly and annotation, allowing a deeper insight into evolution and biodiversity across a range of organisms.

“The quantities of genomic sequence data are vast and the need for high performance computing infrastructures and bioinformatics expertise to analyse these data poses a challenge for many laboratories. EMBL’s novel cloud-based whole-genome-assembly and annotation pipeline involves expertise from the Genomics Core facility in Germany, EMBL’s European Bioinformatics Institute, and EMBL Heidelberg’s IT Services. It will allow scientists, at EMBL and around the world, to overcome these hurdles and provide the right infrastructure on demand” said Rupert Lueck, head of IT services at EMBL.

#### Use Case Objectives:

- Open up new possibilities for scientists to perform large-scale genomic analysis without making large capital investments in computing infrastructure, thereby making de novo assembly and genome annotation affordable to many more laboratories,
- Provide a leading bioinformatics pipeline to perform fast and on-demand genomic data analysis,
- Provide a basis for future extension of genomic research using cloud computing infrastructures.

#### Use Case Requirements

Part of EMBL’s mission is to provide services to the scientific community in Europe. We are clearly motivated to doing so and in the pilot phase our major focus is on

The long term objective is to make de novo assembly and genome annotation widely available to the scientific community maximize the scientific and social benefits we believe this service can bring. Also it proposes to expand into comparative assemblies of human genomes given that the confidentiality and legal issues of dealing with potential clinical data can be satisfied.

Average 42 analyses runs per month

Each Analyses in 1 DC, so typically one separate supplier. All infrastructures per analyses will be within one provider.

*Infrastructure per analyses (duration):*

1. Compute: 50 VM’s for compute, plus 4 VM’s for storage, speed to deliver: Minimum seconds; most likely 30 sec; maximum 2 min

2. Storage: persistent for upload 0,5 – 1 TB, download 10 -100 GB, during calculation 1,2 TB mirrored (so 2,5 TB) + 500Gb permanent
3. Network: Large data transfer will occur when customers (such as EMBL or other labs) will upload Next Generation Sequencing data (e.g. 0.5-1.0 TB) for analysis. During calculation analyses 7+ Gbps. After an analysis customers will download the results (e.g. 10-100 GB). Access to analyses output: ...central EMBL on-site storage?

*Overall capacity during pilot, possibly spread over suppliers:*

1. Compute: 50 – 100 VM's, 54 needed to host 1 run, Minimum seconds; most likely 30 sec; maximum 2 min
2. Storage 3Tb x av. 42/month = approx. 80 TB, e.g. duration min: 5 days; mean: 10 days; max: 6 weeks
3. Network bandwidth: 1 – 10 Gbit

*Other service aspects:*

- RTO < 1 buss day
- RPO < 1day
- MTBF > 6 – 12 months
- Capacity planning based on capacity usage report per analyses
- Monitoring
- SSO
- Service interface, helpdesk & support organization
- Compute & Data behind firewall
- Single point of contact, 1 portal, 1 catalog
- Data in EU
- Jumbo frame

#### 4.4.1 Service description and Service value Mapping

Service Classes	Service Value
VM Provisioning	50 VM's for compute, plus 4 VM's for storage, speed to deliver: Minimum seconds; most likely 30 sec; maximum 2 min
Compute Capacity	50 – 100 VM's, 54 needed to host 1 run, Minimum seconds; most likely 30 sec; maximum 2 min
Storage	persistent for upload 0,5 – 1 TB, download 10 -100 GB, during calculation 1,2 TB mirrored (so 2,5 TB) + 500Gb permanent
Network	Large data transfer will occur when customers will upload Next Generation Sequencing data (e.g. 0.5-1.0 TB) for analysis. During calculation analyses 7+ Gbps. After an analysis customers will download the results (e.g. 10-100 GB).

Service Classes	Service Value
	bandwidth: 1 – 10 Gbit
<b>Identified GAPS</b>	<ul style="list-style-type: none"> <li>– <b>Capacity planning based on capacity usage report per analyses</b></li> <li>– <b>Single point of contact, 1 portal, 1 catalog</b></li> <li>– <b>Jumbo Frames</b></li> </ul>

#### 4.4.2. Available Providers

Service Classes	Atos	CloudSigma	T-Systems	Interoute	The Server Labs
VM Provisioning	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Compute Capacity	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Storage	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Network	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

☒ Covered / ☒ Not covered / ☐ Partially Covered

#### 4.5 New Flagships

The following table presents the initial high level initial requirements for three candidate flagships under evaluation. Based on the expressed high level requirements a rough estimation of potential service Classes to be used is done per each case.

#### PIC Neuroimaging Center on Cloud

Area	Requirements (per work unit)
CPU requirements (total CPU hours)	2000 hours on 3 GHz CPU 1000 hours on GPU-assisted nodes
Peak requirements in terms of CPU, RAM, disk storage (GHz, GB, GB)	CPU – 3 GHz RAM – 2.5 GB per core Disk – 0.5 GB per core
Data requirements (quantity of data, files/databases accessed/produced)	Maximum input data: 7 GBytes Maximum output data: 60 GBytes Number of files: 400
Single/multiple binaries (indicate dependencies between multiple binaries)	Execution of single binaries with internal calls to other binaries.
Peak server size required per resource (server size required in terms of the largest CPU configuration, largest RAM configuration etc.)	Same as peak requirements above.
Programming model (serial, parallel, shared memory, threads, data parallel etc.)	Serial process, parallel process and data parallel
Style of interaction with the user (interactive/batch)	Interactive using PICNIC web portal Batch back-end with emphasis on fast turnaround time. Batch results are accessed interactively by the user for visualization and statistical analysis.

<i>External connections (e.g. data input/output rates and access to external databases)</i>	2 options to be evaluated: Access to external database and data storage. Access to external database only (storage in the cloud).
<i>Expected external connectivity volume (GB/TB per which time frame)</i>	67 Gbytes per dataset reprocessing. This should be small compared to the total turnaround time. Assuming 200 jobs can be executed in parallel, the turnaround time at 100% efficiency would be 10 hours. Taking a maximum of a 10% impact on turnaround, the 67 Gbytes need to be moved in 1 hour,
<i>Peak expected bandwidth to/from a cloud server (Mbps/Gbps)</i>	0.15 Gbps effective, real, end-to-end throughput. If interactive visualization of data directly stored in the cloud, the peak bandwidth necessary would be 100 times larger.
<i>Security requirements (e.g. confidentiality of data, algorithms)</i>	Confidentiality of data. We require that all the data is encrypted.
<i>Licensing aspects (third party packages requiring licenses)</i>	Several applications need Matlab.
<i>Operating environment (operating system and version, libraries etc.)</i>	For porting the current environment: -Scientific Linux (64 bit), Centos 5.6 (64bit), HDF5, MNI, VXLm TCL and VTK libraries. - FreeSurfer, SPM or FSL among other imaging software. - NVIDIA drivers, CUDA Toolkit 4, CUDA SDK. For adding additional packages: - Windows
<i>Other technical requirements</i>	Some algorithms require the use of GPUs (GTX Family or Tesla) Internal Memory > 1.5 G Stream Processors > 480 Processor clock > 1.4 MHz

Service Classes	Service Value
VM Provisioning	CPU – 3 GHz RAM – 2.5 GB per core Disk – 0.5 GB per core
Compute Capacity	2000 hours on 3 GHz CPU 1000 hours on GPU-assisted nodes
Storage	Encryption ~ 1 petabyte Maximum input data: 7 GBytes Maximum output data: 60 GBytes Number of files: 400
Network	0.15 Gbps effective, real, end-to-end throughput. If interactive visualization of data directly stored in the cloud bandwidth requirements to be 100 times larger.

#### Identified GAPS

**Software License management (Matlab)**  
**Use of use of GPUs**

Service Classes	Atos	CloudSigma	T-Systems	Interoute	The Server Labs
VM Provisioning	☑	☑	☑	☒	☒
Compute Capacity	☐	☐	☐	☒	☒
Storage	☐	☑	☐	☒	☒
Network	☑	☑	☑	☒	☒

☑ Covered / ☒ Not covered / ☐ Partially Covered

## Ocean & Coastal Information Supersite

Area	Requirements (per work unit)
CPU requirements (total CPU hours)	20 000 single-core-hours / processing day
Peak requirements in terms of CPU, RAM, disk storage (GHz, GB, GB)	CPU – 3 GHz RAM – 2.5 GB per node Disk – 0.5 GB per node
Data requirements (quantity of data, files/databases accessed/produced)	~ 1 petabyte
Single/multiple binaries (indicate dependencies between multiple binaries)	Multiple binaries Apache Hadoop
Peak server size required per resource (server size required in terms of the largest CPU configuration, largest RAM configuration etc.)	All nodes equally configured according to peak requirements
Programming model (serial, parallel, shared memory, threads, data parallel etc.)	Massive parallelisation, MapReduce Data local processing
Style of interaction with the user (interactive/batch)	Interactive through Web Application + CLI for batch modes
External connections (e.g. data input/output rates and access to external databases)	Output: 1 TB / day Input: only once (full mission datasets)
Expected external connectivity volume (GB/TB per which time frame)	1 TB / day
Peak expected bandwidth to/from a cloud server (Mbps/Gbps)	100 Mbps
Security requirements (e.g. confidentiality of data, algorithms)	none
Licensing aspects (third party packages requiring licenses)	none
Operating environment (operating system and version, libraries etc.)	Scientific Linux, Centos 5.6, Fedora, RHEL, HDFS, Apache Hadoop 1.0, python, gcc, java
Other technical requirements	Bandwidth between nodes is critical and shall be as large as possible; a good bandwidth is 10Gbps.

Service Classes	Service Value
Compute Capacity	20 000 single-core-hours / processing day
Storage	~ 1 petabyte Output: 1 TB / day Input: only once (full mission datasets)
Network	100 Mbps

**Identified GAPS** – **Open Source software repository**

Service Classes	Atos	CloudSigma	T-Systems	Interoute	The Server Labs
VM Provisioning	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Compute Capacity	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Storage	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Network	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

☒ Covered / ☒ Not covered / ☒ Partially Covered

## Weather Data Information Supersite

Area	Requirements (per work unit)
CPU requirements (total CPU hours)	10 000 single-core-hours / processing day
Peak requirements in terms of CPU, RAM, disk storage (GHz, GB, GB)	CPU – 3 GHz RAM – 2.5 GB per node Disk – 500 GB per node
Data requirements (quantity of data, files/databases accessed/produced)	~ 1 petabyte
Single/multiple binaries (indicate dependencies between multiple binaries)	Multiple binaries
Peak server size required per resource (server size required in terms of the largest CPU configuration, largest RAM configuration etc.)	All nodes equally configured according to peak requirements
Programming model (serial, parallel, shared memory, threads, data parallel etc.)	Massive parallelisation Data local processing
Style of interaction with the user (interactive/batch)	Interactive through Web Application + CLI for batch modes
External connections (e.g. data input/output rates and access to external databases)	Input: only once (full mission datasets) Output: 1 TB / day
Expected external connectivity volume (GB/TB per which time frame)	1 TB / day
Peak expected bandwidth to/from a cloud server (Mbps/Gbps)	100 Mbps

Security requirements (e.g. confidentiality of data, algorithms)	None
Licensing aspects (third party packages requiring licenses)	None
Operating environment (operating system and version, libraries etc.)	Scientific Linux, Centos 5.6, Fedora, RHEL, HDFS, python, gcc, java
Other technical requirements	Bandwidth between nodes is critical and shall be as large as possible; a good bandwidth is 10Gbps.

Service Classes	Service Value
Compute Capacity	10 000 single-core-hours / processing day
Storage	~ 1 petabyte
Network	100 Mbps

#### Identified GAPS

–

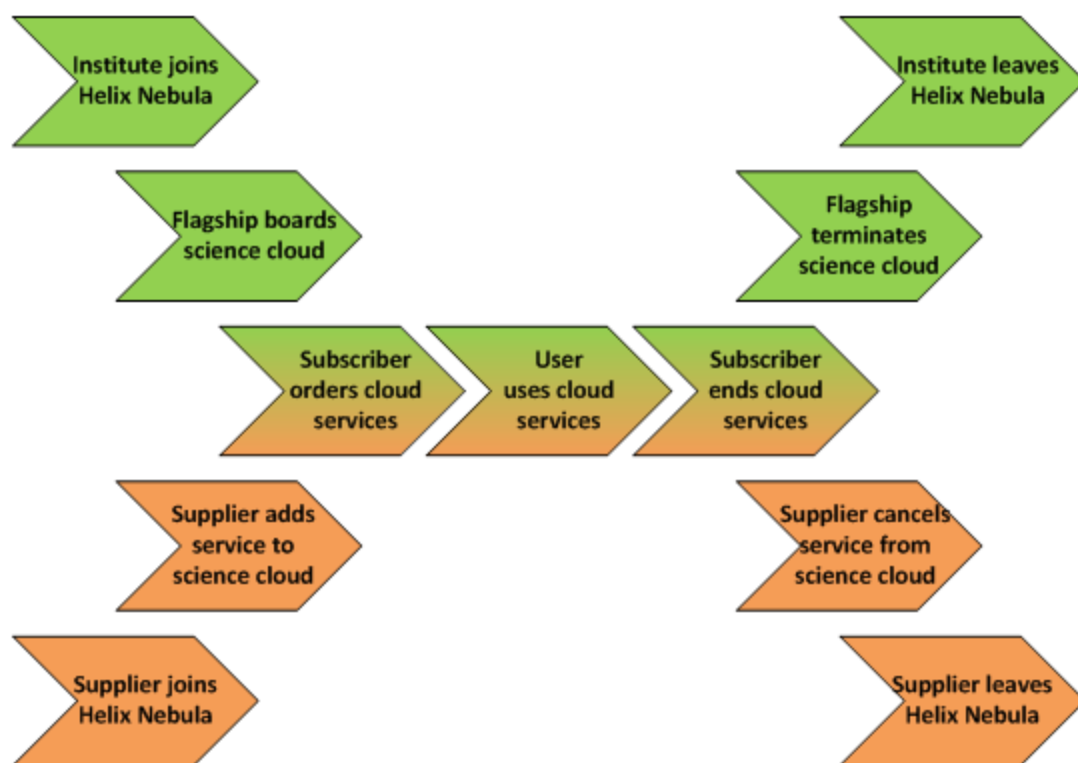
Service Classes	Atos	CloudSigma	T-Systems	Interoute	The Server Labs
VM Provisioning	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Compute Capacity	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Storage	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Network	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

☒ Covered / ☒ Not covered / ☒ Partially Covered

## 5. Helix Nebula Generic Service Acquisition Model

### The high level processes of Helix Nebula

Before a flagship can use cloud services from a supplier a number of prerequisites need to be fulfilled through preceding activities. Also when a flagship stops using the cloud services a number of activities will need to take place. All activities taken together will show the lifecycle of institutes, flagships, services and suppliers within Helix Nebula. The high level processes connect the lifecycles of these elements.



**Figure 4 High level processes of Helix Nebula**

From the demand side (top left) an institute joins Helix Nebula before a flagship can board the science cloud. From the supply side (bottom left) a supplier joins Helix Nebula before it can add services to the science cloud. Then (in the middle) a flagship can order, use and end the use of cloud services provided by the supplier. At the end of the lifecycle a flagship can terminate the science cloud and when there are no more flagships from the institute an institute can leave Helix Nebula (at the top right). Similar a supplier can cancel a service and when he provides no more services he can leave Helix Nebula (at the bottom right).

## Joining Helix Nebula

The Helix Nebula initiative is aimed at European research institutes and other governmental organizations. It is therefore not open to everyone. Up till now all parties that are involved in developing and (ultimately) using Helix Nebula need to join the Helix Nebula consortium. A partner that joins Helix Nebula needs to sign a membership agreement and a non-disclosure agreement.

In the future it is expected that there will be users that intend to use the service 'as is' and will not take part in further developing the service. Those users don't need to know and are not interested in the inner workings of the Helix Nebula services. They can use the services with the knowledge of the publicly available service descriptions.

With the new type of users in mind the questions arose what should be required from a customer (a research institute or governmental organization) when they want to use the Helix Nebula services. The decisions were taken that:

- A customer of the Helix Nebula services needs to join Helix Nebula and therefore sign a membership agreement.
- It should be made possible that customers of Helix Nebula can join Helix Nebula without signing a non-disclosure agreement. This means that those users will not have access to information about the inner workings of Helix Nebula.
- For developers of Helix Nebula services that do need to know the inner workings of the Helix Nebula services the non-disclosure agreement is still required.
- Terms of use for the Helix Nebula services must be developed. These terms of use apply to all users of the Helix Nebula services and should describe the terms of using Helix Nebula services 'as is'.

## Service On-boarding Process (Flagship boards science cloud)

- User management process and assignment of roles
- List of Applications Available for Researcher communities
- Service subscription
- Contracting (SLAs (infrastructure, Application, License Accounting, referencing to service catalogue)
- Technical implementation (network)
- Data Source Access
- Definition of Data Sources replication and back/up recover policies
- Allowed providers per Data sources
- Definition of selected (if there is restriction) IaaS providers to use.
- Definition of SLA templates, Violation policies and penalties.

At the current state in Helix Nebula PoC these processes are not implemented and each supplier is offering its own procedures.

### Service Operations

- Automation of new user definition and application access.
- Automatization of Desktop provisioning and application availability
- Ability to monitor and audit applications usage, license usage, resource usage per provider.

- Ability to access consolidated billing information,

Research communities will be offered to Advance Virtual Desktop: A virtual Desktop where the applications can be easily configured (as in an iPhone) and where they are charged by the actual use. It could contain a combination of own developed applications (through the PaaS layer) and third party applications that could be grouped and offered through the Applications Marketplace.

- Service Aggregation and Orchestration: Integration third party offerings for SW products and IaaS providers
- Helix Nebula operators do not have the need to provide nor infrastructure neither SW packages by themselves, but they become a market place or ecosystem that allow aggregation of services from different sources.
- PaaS: Specific PaaS environment to enable Research Communities or independent commercial IT providers to build their specific applications or in general, to develop applications to be sold through the Helix Nebula market place.

As it is possible to identify that some applications which will be quite common and others which will be more specific. Research communities would have to:

- Select a set of fixed applications for management purposes and other applications which always apply a given research community
- Select a set of variable applications depending on the needs
- Assign a set of resources where all the applications will be deployed

The way to provide the adequate context for those applications and services which may be integrated will be a public category-based marketplace, where all software elements may be published.

Finally, the project will look for a way to manage licenses in such a way that we move from traditional licensing models to more flexible ones like SaaS-based licenses and others. This will allow Helix Nebula to act as a license broker, as a way to offer customers all the needed functionalities without having to cover the whole cost of the original software..

Additionally, Helix Nebula has to provide accounting service which permits the aggregation of application records from several providers.. The accounting solution can be integrated additionally with a billing system based on rules.

Currently, the main dynamic on-demand application system works mainly for allocation time of the license independently of the final productivity or performance of the application. However, for example in the technical application environment, the performance and time to solution of one dataset depends strongly in the used hardware. In this case, the license allocation model only favours the continuous actualization of the hardware as a method to reduce execution costs. So, another license metric is need which produces the same results without dependencies in the hardware (or in a Cloud environment, without interferences of another user). Also, the final records must include information which allow third-parties the audit of this usage so final users and telecom providers can solve any discrepancy in the final usage.

At the current state in Helix Nebula PoCs these processes are not implemented and each supplier is offering its own procedures.

### **Service Request and Provisioning**

- Application discovery through the Application Marketplace
- Ability for split service provision in Federated and hybrid provisioning.
- Ability to access available Data sources.

At the current state in Helix Nebula PoCs these processes are not implemented and each supplier is offering its own procedures.

### **Monitoring and Service Level Reporting**

- Monitoring of both application and infrastructure SLAs

At the current state in Helix Nebula PoCs these processes are not implemented and each supplier is offering its own procedures.

### **Auditing**

Providing means of auditing: detecting, analysing and reacting to incidents within the infrastructure and application services,

At the current state in Helix Nebula PoCs these processes are not implemented and each supplier is offering its own procedures.

## Billing and Payment

- Users invoke services which result in charges from the relevant providers. These need to be recognised and collated, for ascribing to the relevant department or budget.
- Payment can take place on an “N-to-M” basis, from any user organisation to any supplier(s) involved, or can be collated on both sides: on the demand side, a funding agency may actually pay the bills on behalf of end users; on the supply side, where there may be a multi-supplier supply chain, the charges may be handled separately or combined. Where combined, consideration needs to be given as to revenue doublebooking, margin dilution, fiscal constraints, etc.

At the current state in Helix Nebula PoCs these processes are not implemented and each supplier is offering its own procedures.

## 6. Helix Nebula Marketplace Catalogue

This section presents a summary view with all Flagships, Services and Identified providers. This information is summarized in a table that presents findings obtained in Flagship analysis performed in Section 4.

Flagship	Services	Atos	CloudSigma	T-Systems	Interoute	The Server Labs
<b>ESA SuperSites Exploitation Platform</b>	VM Provisioning	☑	☑	☑	☒	☒
	Compute Capacity	☑	☑	☑	☒	☒
	Storage	☑	☑	☑	☒	☒
	Network	☐	☐	☐	☒	☒
<b>CERN Atlas High Energy Physics use case</b>	VM Provisioning	☑	☑	☑	☒	☒
	Compute Capacity	☑	☑	☑	☒	☒
	Storage	☑	☑	☑	☒	☒
	Network	☑	☑	☑	☒	☒
<b>EMBL: Genomic Assembly in the Cloud</b>	VM Provisioning	☑	☑	☑	☒	☒
	Compute Capacity	☑	☑	☑	☒	☒
	Storage	☑	☑	☑	☒	☒
	Network	☐	☐	☐	☒	☒
<b>PIC Neuroimaging Center on Cloud</b>	VM Provisioning	☑	☑	☑	☒	☒
	Compute Capacity	☐	☐	☐	☒	☒
	Storage	☐	☑	☐	☒	☒
	Network	☑	☑	☑	☒	☒
<b>Ocean &amp; Coastal Information Supersite</b>	VM Provisioning	☑	☑	☑	☒	☒
	Compute Capacity	☑	☑	☑	☒	☒
	Storage	☑	☑	☑	☒	☒
	Network	☑	☑	☑	☒	☒
<b>Weather Data Information Supersite</b>	VM Provisioning	☑	☑	☑	☒	☒
	Compute Capacity	☑	☑	☑	☒	☒
	Storage	☑	☑	☑	☒	☒
	Network	☑	☑	☑	☒	☒

☑ Covered / ☒ Not covered / ☐ Partially Covered

## 7. Conclusion and Next Steps

This “D4.2 Cloud Provisioning: case histories of decisions taken” deliverable has presented the process and methods implemented for matching requirements to possible supply combinations in each flagship use-case and the workflow for “on-boarding” a case onto a required set of services.

The process followed in the initial wave of Use Cases has demonstrated that a more automated approach is required in the future developments of the Blue Box and Helix Nebula as a whole by further developing the envisaged Broker model described in this document.

In addition, the document has also identified that initial focus of the developments was to determine the technologies needed to establish a working Proof of Concept (PoC). This focus has led to an insight in the technical requirements of the current flagships, which turn out to be quite in line with the initial requirements from the new flagships.

It has been identified that in the future, further attention to the non-technical requirements will be required, in order to address the overall Use Case description described in Section 4. The vision for Helix Nebula is that it evolves into a research and business ecosystem enabled through a combination of capabilities, data and services from several actors. The ambition would be that each actors can use services provided by other actors and contribute specific abilities that can be used by other actors. Consider, for example, the following actors with the services they use and the capability they contribute (both paid and unpaid contributions are possible). The cloud services of the suppliers give the research organizations the infrastructure to host their data and to run their processing. The research organizations could give access to their diverse and enormous data sources. Innovative businesses could use the available data and processing capacity to create (information) services for other businesses. From this network of organizations and the current data explosion new and innovative business models could emerge. New actors need to be introduced in the Helix Nebula ecosystem with specific emphasis on the supplier side. Identified actors at this stage are: Infrastructure as a Service (IaaS) providers, Platform as a Service (PaaS) providers, independent software vendors, Software as a Service (SaaS) providers and technology and solution providers.

## 8. Appendix Literature list

Reference number	Document number	Document name	Version
[1]	-	Helix Nebula Technical architecture <a href="#">Helix Nebula Architecture v1 2</a>	v1 2
[2]		Helix Nebula Service architecture HN_Serv_Arch_MAS_07	v07
[3]	-	Helix Nebula Service Properties catalogue Helix Nebula Service_Properties_0.4	v04
[4]	-	Matching process perspectives HN WP4 Matching process perspectives v0.2 20121206	V02
[5]	-	Helix Nebula Service Architecture Operational Model Helix_Nebula_Service_Architecture_Operational_Model_v0.4	V04
[6]		Report on Cloud Computing to the OSG Steering Committee, <a href="http://www.spec.org/osgcloud/">http://www.spec.org/osgcloud/</a>	
[7]		<a href="http://opennebula.org/cloud:sandboxtestdrive">http://opennebula.org/cloud:sandboxtestdrive</a>	
[8]		Open Data Center Alliance: Master Usage model Compute Infrastructure as a service	v1.0
[9]		Service Descriptions for Cloud Services - The Customer's Perspective; <a href="http://ssrn.com/abstract=2109361">http://ssrn.com/abstract=2109361</a>	
[10]		A Reference Architecture for Multi-Level SLA Management	

## 9. Appendix Services descriptions per supplier

This appendix presents in detail some of the Cloud Services offered by suppliers in Helix Nebula. The ability of the suppliers to fulfill the requirements of the flagships, which is presented in section 4, has been determined on the basis of these descriptions. The detailed services descriptions are presented for the five existing suppliers in the Helix Nebula ecosystem: Atos, T-Systems, Cloudsigma, Interoute and The Server Labs.

### Atos

Provider	Atos
Services	<p>Atos Services in Helix Nebula are delivered by providing an Open Source cloud management stack on top of our normal Cloud Infrastructure Services (CIS) hardware.</p> <p>Atos offers customers a secure and centrally hosted environment that meets the requirements of research software developers and scientific users, yet addresses the needs of IT departments for quality production environments, and could also be used for other purposes, e.g. acceptance testing and training. The capacity provisioned can be adjusted to the customers' needs as they vary over time. The immediate availability of infrastructure resources reduces service provisioning times significantly. The services are targeted at customers who want:</p> <ul style="list-style-type: none"> <li>– Easy self-management and responsive delivery</li> <li>– Based on highly standardized hardware environment to ensure quality</li> <li>– Flexibility in processing hardware, storage and network in terms of how this capacity is used</li> <li>– Flexible adding or reducing capacity</li> <li>– No investment compared to customer-dedicated hardware</li> <li>– Certified, compliant and secure environments eligible to customer needs.</li> </ul> <p>The scope of these Atos services comprises:</p> <ul style="list-style-type: none"> <li>– Provision of datacenter infrastructure</li> <li>– Provision of server capacity on demand</li> <li>– Provision of storage capacity on demand</li> <li>– Provision of network (LAN, SAN) connection</li> <li>– Operation of infrastructure</li> <li>– Online portal and/or API for service requests (ordering and changes)</li> <li>– Online portal provides standard reporting (usage and reserved capacity if applicable) and billing</li> </ul> <p>The services are offered in two different delivery models, Basic (or Uncommitted) and Committed:</p> <p>Basic:</p> <ul style="list-style-type: none"> <li>– Multi-Tenant Private Cloud</li> <li>– Assignment of requested capacity, if available with Atos</li> <li>– Guaranteed capacity after assignment by Atos</li> <li>– 'Pay as you Go' per hour charging model of allocated capacity</li> <li>– Access via internet (default) or dedicated network connection (e.g. via GEANT)</li> <li>– Capacity is provided on Shared hardware</li> </ul> <p>Committed:</p> <ul style="list-style-type: none"> <li>– Multi-Tenant Private Cloud</li> </ul>

- Atos guarantees 75% of the capacity ordered by customer
- Customer is free to vary in his monthly capacity reservation in defined steps. New reservations have to be committed for 1 months
- Customer can use capacity up to ordered volume 100%, if available with Atos.
- Per month charging model of committed capacity
- Access via internet (default) or dedicated network connection
- Capacity is provided on Shared hardware

Both offerings include:

- Customer environments are separated by solution embedded firewall
- Predefined virtual server order items, based on Linux (Redhat, SUSE) and Windows, can be used for the fast rollout of virtual machines. Customers can in addition place own server / application templates to customize their environment
- All management functionality within the environment is possible through the WEB-Portal interface or API. Various rights and roles can be assigned to the operators to build up a customer specific role model.
- The required resources e.g. CPU, memory and storage are automatically allocated and provisioned. The used/committed resources are monitored, reported and invoiced.
- Atos manages the Cloud computing environment in accordance to ISO 9001 (quality), ISO14001 (environmental) and ISO 27001 (security) standards.

Further service variants are possible:

- A single-tenant Private Cloud variant can be established, if required, with similar attributes as the Committed option.
- Microsoft Operating system licenses can be provided, where required, based on Service Provider License Agreement (SPLA).

Other optional but standard service components are available, including:

- Secure connectivity via the internet;
- Additional network security zones;
- Production support for the running environment;
- Backup and restore of customer data;

On-boarding services

Service Class	Provisioning Portal
<b>Basic Startup</b>	
To start using the services, the customer needs access the provisioning portal and/or API. This access will be granted if the customer fills in the authorization form and the information is verified and accepted by Atos.	
<i>Service availability</i> Service Availability: 99,9% Service Availability Window: 7*24 hours (all days) 24 hours a day Maintenance Windows: Twice a year, Saturday 08:00 h - Sunday 24:00 h Service Access Point: Availability is measured at the administration portal and at the virtualization farm.	

<i>Support Availability</i>	
Support availability window: 5*10 hours: Business Days, 08.00 – 18.00 h	
Incident handling window:	
Severity 1:	7*24 hours: All days, 24 hours a day
Severity 2, 3, or 4:	5*10 hours: Business Days, 08.00 – 18.00 h
Support language: English	
Completion of Service Setup for new customer (service level reporting not applicable)	
<b>Customer provisioning portal</b>	
<p>The customer provisioning portal gives the key users of the customer the possibility to execute standard service requests.</p> <p>The portal creates machines based on the values and templates used. Machines are instantly created in a few steps.</p> <p>Atos by default supports the current Operating System release and the previous release:</p> <p>Red Hat Enterprise Linux version 5, Red Hat Enterprise Linux version 6, SUSE Linux, Sun Solaris for Intel Debian and Ubuntu. Microsoft Windows Server 2003 Standard Edition, Microsoft Windows Server 2003 Enterprise Edition, Microsoft Windows Server 2008 Standard Edition, Microsoft Windows Server 2008 Enterprise Edition, Microsoft Windows Server 2008 R2 Standard Edition, Microsoft Windows Server 2008 R2 Enterprise Edition,</p> <p>An OS image will be available within 3 months after availability of the new release of the Operating system. The transition or upgrade of instances based on one release to another is not included.</p> <p>Where Microsoft and other vendors regularly issue a security patch or “service pack”, installation of these service packs into the golden images is covered by this service component . Atos supports the current Service Pack and the previous one for all the supported releases.</p> <p>The service supports management of Network, Firewall, Software catalogs and virtual servers.</p> <p>Atos provides a Shared Software Library. By default it will contain the software that is used in the master images. The customer will be able to create their own customer library to put specific software there, or his own private master images. Customized master images are not subject to the Atos maintenance of master images. Additional software installed by customers on the environment must be licensed conforming to the terms of the vendor of this software.</p>	
<b>Service Class</b>	Compute capacity (processing)
<b>Computer processing</b>	
<p>Computer processing is provided based on an infrastructure located at Atos data centers and includes an access point for the customer.</p> <p>Computer processing consists of:</p> <ul style="list-style-type: none"> <li>– Secure access to the computer processing environment</li> <li>– Management of hypervisor software for virtual machines</li> <li>– Hypervisor software licenses</li> <li>– Storage connectivity equipment and tooling</li> <li>– Storage capacity equipment and tooling</li> <li>– Hardware maintenance and software support</li> <li>– Data centre LAN and SAN equipment and tooling</li> <li>– Shared management infrastructure. Including connections and licenses needed to execute these shared management functions</li> </ul> <p>The computing resources supplied by multiple hosts running a hypervisor are combined in a shared (multiple customer) computer processing farm (depending in the selected delivery model). The instances</p>	

<p>(of several customers) run in this computer processing farm.</p> <p>A single site cluster of hosts running a hypervisor is provided. This cluster has the capacity to overcome a failure of one of the hosts running a hypervisor. Virtual machines will automatically be assigned to another host running a hypervisor within the cluster.</p> <p>Virtualization</p> <p>The virtualization tool that will be used is determined by Atos. Each virtual machine is provisioned from a pre-built template. The image consists of a virtual machine and the associated computing resources.</p> <p>Atos provides images on standard building blocks. Several building block exist, see Customer provisioning portal service for details</p> <p>Data Center</p> <p>The infrastructure is located in an Atos data centre in a secure and robust computer environment. Power, air conditioning, security, fire precautions, racking, cabling, administration, procedures, and access controls are provided.</p> <p>Support language: English</p> <p>Add new virtual machine</p> <p>Started by customer, 95% within 15 minutes after start execution.</p> <p>Downsizing of IaaS</p> <p>Depends on Business Model.</p> <p>Basic – no limitations – it not other specified</p> <p>Committed – after minimum contract period</p> <p>Up scaling</p> <p>Up scaling for Basic vDC is not restricted, but availability of capacity is not guaranteed.</p> <p>For committed or reserved vDC up/down-scaling is available up to maximum size contracted. Change of Committed size will be executed within 14 days.</p> <p>Billing Any effects of a capacity increase or decrease are measured per hour and billed per hour or per month, depending on the delivery model, Uncommitted or Committed. A recurring fee for measured products will be charged.</p> <p>Hardware - virtualization host failure</p> <p>If the hardware failure is related to a single host, the management environment will automatically redeploy the Virtual machine to another virtualization host in the cluster.</p>	
<i>Service Class</i>	<i>Storage.</i>
Storage – Silver Class	
Silver Class storage provides RAID-5 based storage meant for databases and applications requiring moderate storage performance.	
Storage – Gold Class	
Gold Class storage provides RAID based storage meant for databases and applications requiring high performance. Gold Class storage is capable of handling high volumes.	
Backup and Restore	
<p>Atos provides the customer an environment to implement, schedule and restore data. The environment will be provided on request.</p> <p>The standard retention for backups is 14 days and can be changed according to business needs. The storage needed to provide the saved data is charged according to the price list.</p> <p>Customer is handed over an interface for the self-management of backup tasks</p>	

Service Class	Network.
<b>Virtual Security Zones</b>	
Customers are able to create a number of virtual networks to create own security zones within their Virtual Datacenter environment.	
<b>Connectivity – Secure access to Virtual Security Zones networks via Internet</b>	
This connection allows the customer to access his Security Zone over the internet. This connection is VPN based and can be configured for each internal network via the integrated Firewall via Virtual Private Network by the customers. Internet connectivity 2.1.8 is a prerequisite	
<b>Connectivity – Access towards the Internet for Virtual Machines</b>	
Atos can provide access towards the internet for the virtual machines; this access is delivered through a shared DMZ. By default standard service ports of our external firewall are open. The available Internet IP addresses can be mapped by the customer (IP-translation or port forwarding) to a virtual server. Public IP-addresses have to be ordered separately.	
<b>Additional Security Zones</b>	
If Networks outside of the virtual datacenters are needed, customers can order additional networks to reflect their business e.g. to include systems outside the virtual environment. The setup is based on VLANs, Firewalls and a Virtual Private Network (VPN) Route Forwarding (VRF). Physical servers and virtual Machines are placed in this customer-dedicated security zone, depending on the chosen cloud variant. Network types are: SSN screened subnet DMZ a demilitarized zone HSZ high secure zone.	
<b>Connectivity – Dedicated</b>	
The customer can request a connection from his own Wide Area Network connection directly to a secure network zone. If a connection between the customer and Atos already exists, this connection can be used. If a new connection must be set up this will be requested using a Telecom company <sup>4</sup> . The connection is secured using VRF.	
<b>Service availability</b>	
Service Availability: 99,9% Service Availability Window: 7*24 hours (all days) 24 hours a day Maintenance Windows: Twice a year, Saturday 08:00 h - Sunday 24:00 h Service Access Point: Availability is measured at the administration portal and at the virtualization farm. <i>Support Availability</i> Support availability window: 5*10 hours: Business Days, 08.00 – 18.00 h Incident handling window: Severity 1: 7*24 hours: All days, 24 hours a day Severity 2, 3, or 4: 5*10 hours: Business Days, 08.00 – 18.00 h Support language: English	
Service Class	Added value Services
<b>Availability Management</b>	
Availability Management ensures that service availability levels are met by monitoring components of the infrastructure for unexpected events which may disrupt service and, where appropriate, triggering the	

<sup>4</sup> Consider that this may take up to 13 weeks before the connection is delivered.

<p>incident management process. Atos' processes are certified under ISO 9001 Quality management. The availability of applications is not within the scope.</p> <p>Availability Management takes care of automated availability monitoring, using system management tools to:</p> <ul style="list-style-type: none"> <li>– Prevent and detect incidents</li> <li>– Prevent unplanned outages</li> <li>– Trigger the incident management process automatically when required</li> </ul>
<b>Monitoring</b>
<p>Monitoring consists of:</p> <ul style="list-style-type: none"> <li>– Monitoring of the infrastructure</li> <li>– Comprehensive event management to detect unexpected hardware events relating to memory, processors, disks and network. Hardware events are monitored using vendor-supplied hardware monitoring tooling. Actions regarding maintenance and / or replacement are initiated</li> <li>– Comprehensive event management to detect unexpected infrastructure related events</li> </ul>
<b>Performance Management</b>
<p>Performance management is done for the virtualization farm, not for the individual virtual machines. The customer can initiate the change to upgrade the virtual machine if more performance is needed.</p> <p>Performance Management ensures that the infrastructure has an acceptable performance level. Atos follows the guidelines and rules for settings provided by suppliers. Performance management is executed when there are related monitoring alerts and/or incidents that indicate a need to check the performance. In such a case, reactive analysis is executed for the information files related to hardware, storage or network, and corrective actions are taken.</p> <p>Performance Management of applications that are running on the infrastructure is not included.</p>
<b>Capacity Management</b>
<p>Atos ensures that there are sufficient resources in the virtualization farm to fulfill requests within the limits and time frames defined in this service description.</p>
<b>Security Management</b>
<p>Atos manages the virtualization farm in accordance with ISO 27001 standards, which are a set of administrative security guidelines that help maintain a high level of security at an organizational and technical level.</p> <p>Physical security of the hardware infrastructure is provided based on Atos data center policies. The customer or third parties have no physical access to the infrastructure.</p>
<b>Change management</b>
<p>For changes requested by the customer, we distinguish two types:</p> <ul style="list-style-type: none"> <li>– Standard changes, can be executed by the customer</li> <li>– Non-standard changes (i.e. every change that is not listed as a standard change) – These are billed on T&amp;M basis</li> </ul> <p>For changes related to infrastructure management</p> <ul style="list-style-type: none"> <li>- These are carried out by Atos to comply with the service levels and are for system farm maintenance, update and incident resolution purposes of the virtualization farm. These are part of the basic service and not invoiced separately to the customer</li> </ul>
<b>Incident Management</b>
<p>Incident management covers all actions necessary to ensure that a failed (or failing) service is restored within the service levels described. Incident management is restricted to the virtualization farm.</p> <p>Incident Management as delivered for this Service includes</p> <ul style="list-style-type: none"> <li>– Handling incidents reported by automated monitoring processes</li> </ul>

<ul style="list-style-type: none"> <li>– Handling incidents that affect the service and are raised by the customers</li> <li>– Executing crisis management</li> </ul> <p>Incident Management concentrates on second-line activities. It covers the handling of incidents that cannot be answered by the first-line support organization.</p> <p>Incident management does not include root cause analysis. This is included in a separate paragraph.</p>	
<b>Problem Management</b>	
<p>The goal of problem management is to prevent (re)-occurrence of incidents by eliminating their root cause. Problem Management is restricted to the virtualization farm and includes:</p> <ul style="list-style-type: none"> <li>– Carrying out trend analyses on recurring incidents</li> <li>– Processing error information of relevant hardware manufacturers</li> <li>– Root cause analysis</li> </ul> <p>When after a defined period of time no permanent solution has been found or implemented, a problem is flagged as a Known Error.</p> <p>Problem Management may lead to the implementation of changes when necessary for a permanent fix.</p>	
<b>Query Management</b>	
<p>The goal of Query Management is to answer customer questions. Query Management covers the handling of queries that cannot be answered by the first-line support organization.</p>	
<b>Complaint Management</b>	
<p>Complaint management attempts to resolve expressions of dissatisfaction. A complaint always receives management attention at an appropriate level.</p>	
<b>Configuration Management</b>	
<p>Configuration Management records the virtualization hosts in the farm, the virtual machines which are located on each piece of hardware, their related configuration details and allocation of resources.</p> <p>As this information is volatile, it is gathered in an automated fashion. Atos uses configuration information also for billing and capacity forecasting purposes.</p>	
<b>Service Reporting</b>	
<p>Service Reporting consists of:</p> <ul style="list-style-type: none"> <li>– Reporting on Service Availability as defined per each service</li> <li>– Reporting on the service levels listed per each service</li> </ul>	
<b>Virus Outbreak Management</b>	
<p>Atos is, as part of the basic service, not responsible for the status of Anti-Virus software and definitions neither on running VM's, nor for the sanity and safety of individual VM's.</p> <p>As part of Atos' continuous efforts to protect against virus outbreaks, Atos monitors Network traffic volumes constantly. When a sudden increases in network traffic is detected, giving rise to the belief that a (part of) a customer environment may be infected by malware, Atos will inform the customer directly.</p> <p>Atos may without further communication with or agreement by the customer decide to isolate (part of) the customer environment, if Atos believes that the continuity of other customers is jeopardized, or if continuity of services delivered by Atos to customers is in danger.</p>	
<b>On boarding Services</b>	
<p>On Boarding Service helps the customer to migrated from other environments to the platform</p> <p>The service is charged at Time &amp; Material price.</p>	
<b>Service Class</b>	Support
<b>Basic Support</b>	
<ul style="list-style-type: none"> <li>– Web-Portal to support (1st Level) a Service self-management process (e.g. FAQ, Service documentation, Article and instructions, Tips and Tricks, Service Cloud Community)</li> <li>– Support (2nd and 3rd Level) to the Organization Administrator when the 1st Level is not helpful</li> </ul>	

and the problem is caused by a faulty Cloud environment.

#### Graded Premium Support

Fast support channel engaged with experienced technicians to support customers who are looking for technical help to utilize the environment. (E.g. configuration of VMs, development of customized images, firewall configuration, usage of 3rd Party products, APIs).

Graded Premium Support is limited to 8 hours per month.

## CloudSigma

Provider	CloudSigma
Services	
<b>Description:</b>	<ul style="list-style-type: none"> <li>Resources at CloudSigma are unbundled and purchased independently which means users are able to build virtual server instances with the exact combination of CPU, RAM, storage and bandwidth required.</li> <li>Users are given the option of full API access with all account actions available allowing complete automation and remote monitoring.</li> <li>Users are also given the option of a feature-rich yet intuitive web browser based GUI. It has been designed to allow easy resource and project management via a web browser. With many advanced and unique features it saves users time and money.</li> <li>Resource usage is calculated using rolling 5 minute billing periods so you only pay for what you use. Subscription discounts are available: 3% for 3 months, 10% for 6 months, 25% for 1 year, 35% for 2 years and 45% for 3 years. Visit CloudSigma's pricing page<sup>5</sup> for up-to-date pricing.</li> </ul>
Service Class	VM provisioning
Guaranteed CPU, RAM and Storage Allocation backed by a 100% Service Level Agreement x50 credit.	
VM provisioning	
<b>Description:</b>	<ul style="list-style-type: none"> <li>Any x86/x64 operating system will run on the CloudSigma platform as long as it is compatible with standard Intel/AMD architecture.</li> <li>Users can upload drive images via our API, web console or via an FTP client. All three methods allow pausing and resuming of uploads. Every account has direct access via secure encrypted FTPS allowing users not only to upload new drive images with ease but also to download all their account data storage at any time.</li> <li>CloudSigma provides an extensive library of ready to use public drive images that can be deployed instantly.</li> <li>An interface for exporting drives in RAW format via the Gui is provided. Third party conversion tools can also be used.</li> <li>Our API is implemented in a REST-style, using the URL to specify an object and an action, HTTPS GET to read state and HTTPS POST to change state</li> </ul>

<sup>5</sup> <http://www.cloudsigma.com/en/pricing/price-schedules/?affid=1155&offid=2>

Service Class	VM image transformation
VM image transformation	
<b>Description:</b> <ul style="list-style-type: none"> <li>– The minimum and maximum range of CPU allocation is as follows: 1 Core 250Mhz minimum and 20 Core 40Ghz maximum.</li> <li>– CloudSigma offers a minimum CPU availability proportional to reserved size. The total CPU available on a machine is shared pro-rata between virtual server instances subject to minimum allocation. At times when the CPU is not fully allocated, virtual server instances on our cloud will actually have more CPU allocated to them than paid for as there is no reason to limit CPU capacity if it is available.</li> <li>– With CloudSigma you can manually override our system and specify the exact number of CPU cores you want any particular server instance to use. In this way each virtual server instance can be optimised for the particular task it is undertaking by the user. The minimum granularity is 1GHz per core.</li> <li>– Our current CPU pricing based on a non-discounted 1 month subscription period is 0.0175EUR Core-GHz/hour.</li> </ul>	
Service Class	Compute capacity (processing)
SigSTORE	
<b>Description:</b> <ul style="list-style-type: none"> <li>– The minimum storage configuration is 1GB and the maximum is 8TB per volume. It is a fully redundant system.</li> <li>– All storage in CloudSigma's cloud is persistent. Storage is available in the form of drives which are created by the user. A drives data is persistent whether associated or not with a server and whether or not a drive is mounted. Storage data is only removed from an account when the account owner deletes a drive.</li> <li>– All user virtual drives are encrypted using a 256bit AES-XTS encryption cascade automatically. Encryption of drives prevents data leakage of deleted drives at a later date. This measure protects 'data at rest' within our cloud.</li> <li>– CloudSigma can offer a tiered storage system and provide automation software for managing the process of moving data between storage media. The number of storage tiers required and the type of media is determined by pre-defined service levels.</li> <li>– Our current SSD storage pricing based on a non-discounted 1 month subscription period is 0.35EUR GB/per month.</li> </ul>	
Service Class	Storage
Storage	
<b>Description:</b> <ul style="list-style-type: none"> <li>– Global connectivity is achieved by leasing physical infrastructure from service providers and operating layer 3 network on top of that. This allows for direct peering connection with other ISPs and guarantees performance and low latency.</li> <li>– CloudSigma offers dual 10GigE networking as standard with free, unlimited incoming bandwidth for all accounts. All users can easily upload as much data as they require to our cloud using our API and FTP.</li> </ul>	

- We don't apply rate limits to incoming or outgoing traffic. Both public and private networks.
- All virtual drives are encrypted with a unique 256bit AES-XTS key that ensures a high level of protection. Traffic is isolated to a private network that is not accessible or visible to other users. This protects critical 'data in transit' and forms part of our comprehensive cloud security framework.
- All virtual server instances within the CloudSigma cloud have IP addresses assigned to them via DHCP. Users can also choose to assign a static or multiple static IP addresses to a server or turn off any external IP allocation altogether.
- If a user already owns their own PI IP addresses, we are able to incorporate these into our BGP session with our upstream providers and make them available for allocation by that user to their virtual server instances within our cloud.
- Users can create private networks within the CloudSigma cloud and configure as standard VLANs on the physical infrastructure. They operate without any restrictions and can be named by the user. Servers are easily added to them in an integrated way. An account can have an unlimited number of VLANs. All network traffic between virtual server instances in a VLAN is not billed.
- Under the advanced settings for each virtual server instance the CloudSigma platform allows the user to specify both the public network card and private network card they prefer to use from a standard list. More specialised and advanced users can optimise the hardware simulation choice to best fit their networking usage and needs.
- Our current pricing for network resources based on a non-discounted 1 month subscription period are 3.5000EUR per month for Static IP Admin and 7.0000EUR per month for VLAN.
- Our current pricing for data transfer based on a non-discounted 1 month subscription period is **free** for Incoming Bandwidth and 0.0455EUR for Outgoing Bandwidth per/GB.

Service Class	Network
---------------	---------

**Service Level Offered:**

- 100% guarantee on network availability in any given month.
- A network latency of 1ms or less for data packets between servers within CloudSigma's service and network.

Network
---------

**Description:**

- CloudSigma uses the latest generation of server-class 1600GHz DDR3 EEC RAM.
- Our current RAM pricing based on a non-discounted 1 month subscription period is 0.0228EUR GB/hour.

Service Class	Random-Access Memory
---------------	----------------------

Random-Access Memory
----------------------

**Description:**

Support is divided into three tiers; Silver, Gold and Platinum.

Item	Silver	Gold	Platinum
Uptime guarantee	100%	100%	100%
Chat and Email Customer Service - 24x7x365	Yes	Yes	Yes

Web based ticket system	Yes	Yes	Yes
Online Support Forums	Yes	Yes	Yes
Documentation, White Papers, Best Practice Guides	Yes	Yes	Yes
Access to Technical Support	Yes	Yes	Yes
Named Contacts	2	5	Unlimited
Max. Response Time	12 hours	2 hours	30 minutes
Dedicated Resolution Engineer	No	No	Yes & Continuous Effort
SLA credits for qualifying downtime after 15min	x50	x50	x100
Enterprise Architecture Solutions Support	No	No	Yes
Best Practice Guidance & Webinars	No	Yes	Yes
Direct Routing to Senior Engineers	No	No	Yes
Cell phone Customer Service - 24x7x365	No	No	Yes
Price	Free	CHF 500 /month	CHF 1,500 /month

\*Enterprise Solutions Architecture Support has a response time of 48 hours and is capped at 2 hours per month. Further use is paid as consulting by the hour of 250 CHF.

**The escalation policy:**

	Silver	Gold	Platinum
First level Customer Relation	Entry	Initial review, resolution during working hours. Escalation after 24 hours.	3 hours continuous effort then escalation
Second level - Engineer	Final escalation level	1 working day then escalation	6 hours continuous effort then escalation
Manager		Final escalation level	12 hours continuous effort then escalation

D- and C- level executives			Final escalation level
Service Class: Support			
The availability guarantee is 100% and the SLA credits of x50 and x100 kick in if that has been violated after 5 minutes downtime.			

## T-Systems

Provider	T-Systems International							
Services								
Dynamic (DSI vCloud)	Services	for	Infrastructure	with	VMware vCloud	Datacenter	Services	
<p>Dynamic Services for Infrastructure with VMware vCloud Data Center Services, expands T-Systems DSI private cloud family with a secure hybrid cloud solution. Fully aligned with the VMware vCloud Datacenter Provider certification programme, it has a build-in Internet connection combined with exposure of the market recognized VMware vCloud APIs &amp; User Interface to the customers. The primary use-case addresses the needs for <b>secure hybrid cloud</b> scenarios of customer's production grade workloads. The core functions of the IaaS services consist of:</p> <ul style="list-style-type: none"><li>– Virtual Datacenter (vDC) resource boundary defined by ordered compute capacity (RAM &amp; Performance Units (PU)) where the customer workloads (VMs or vApps) can be flexibly deployed in by the customer in self-service mode. Measuring and charging is by vDC usage, not individual VM usage.</li><li>– Virtual Organisation (vOrg) User&amp;Access management boundary where the customer manages his Users, Groups and their privileges to various objects in the cloud (e.g. VMs/vApps/Catalogs). This includes the optional ability to integrate to customer existing Identity &amp; User Management solution e.g. via SAML or LDAP.</li><li>– Data storage with two different performance characteristics and dispersed locations to support backup use-cases.</li><li>– Environment designed and operated in fulltime 24x7 mode to fit the needs of enterprise customers with both business critical Production Tier2/3 and Test/Dev workloads. The operation includes attended service by providers own personnel and ability to open an incident via email.</li><li>– Seamless integration with the existing private VMware-based virtualized environments / clouds of the customer.</li><li>– VLANs for secure network separation, mainly between customers.</li><li>– Virtual Networks (VXLANs) for secure separation, mainly between workloads (VMs or vApps) of the customer - pooled on demand and fully managed by customer via the self-service portal or API.</li><li>– Virtual Networking &amp; Security services incl. Load Balancing, Firewalling, Routing and other</li></ul>								

capabilities usable in flexible self-service mode. This includes the ability to deploy isolated networks with overlapping IP Addresses (Prod --> Test) while virtual Firewalling/NAT ensures that they are able to communicate with the surrounding environments.

- Secure Internet Access for customer workloads (VMs or vApps) to establish a high security DMZ-like environment (e.g. for Webservers). This includes redundant high bandwidth internet connection, implemented via multiple upstream providers, protected by multiple security levels of independent security technology and firewalled/encrypted (SSL/VPN protected) access routes.
- Access to the vCloud Portal from Internet and/or Intranet.
- Access to the market recognized VMware vCloud APIs to enable customers to use their own Service Orchestration & Management / Shopping Portal / Cloud Broker / SaaS services etc. based on the T-Systems IaaS services.
- Ability to import workloads (VMs or vApps) to the DSI-vCloud Datacenter environment and export them back via standardized open virtualisation formats (OVF is released via DMTF, no vendor or provider lock-in).
- Catalog based functionality, incl. the ability to pre-configure the setup of whole application landscape by definition of multiple workloads (VMs) stored as encapsulated package reusable once or many times.
- Easy to use Snapshot & Catalog based backup combined with rapid VM deployment functionality, which support Test / Dev / Training / Backup & Restore and similar use-cases, where very fast 'prepare once' (create new vApp or copy existing Production vApp), 're-deploy once' or 'deploy many times' is needed.
- Provider engineered and pre-configured VM templates of most common operating systems, stored in the vCloud "public catalog" to help customer jumpstart his deployment in the cloud. (Windows, Linux).
- Ability to define "Leases" incl. notifications in case of automated expiry of objects to improve the cost/usage control of the customer (e.g. for scheduled test-runs).
- Highly secured environment with physical and logical protection levels, hosted exclusively in Tier 3+ availability level Datacenters (mainly twin core DCs, starting in Munich/Germany, will be subsequently extended to other locations).
- Contractual Framework with high SLA guaranties (99.9%) and short cancelation period (3months), which is based on German legal framework and data privacy regulations.
- Consumption based charging model, supporting both, stable production workloads and rapidly changing workloads - all based on OpEx instead of CapEx.
- Support of complex customer driven security & compliance audit needs compliant with the VMware vCloud Datacenter regulations with access to various activity logs.
- Easy access to the network backbone of T-Systems and Deutsche Telekom, with extensive telecommunication experience and state of the art MPLS/VPN/WAN-Acceleration/Application-Performance-Monitoring and other products.
- Optionally Cloud Readiness Services of the TSI System Integration division can be ordered and help the customer to pick the right workloads to be transferred to the cloud, modernize applications to take full advantage of hybrid cloud operational and commercial models, and help to improve the customer cloud compatibility.

- Scalable solution which includes pay-per-use commercial elements
- Default Allocation model with 75% reservation of compute (CPU+RAM) resources and 100% for Storage, which ensure that the customers' business critical application get enough resources (within the vDC assigned capacity). 25% on demand capacity for load spikes. Other models are possible with varying levels of reservation on request
- Multi-Language support for the Portal User Interface (at present 9 languages, automatically adjusted to end-user browser language settings).

The ability to encapsulate whole application landscape(s) with multiple VMs in one vApp, including the vApp-Network with detailed network configuration and network services a specific application needs, and subsequently deploy it easily in the cloud or move back to customer premises environment, is very powerful source of flexibility and enabler of further business agility. This functionality is as well the core of the Software Defined Datacenter / Software Defined Networking everybody speaks of, and the DSI-vCloud Datacenter platform already delivers.

The service is accessible via Internet and/or intranet (private network), for Portal/API/Console as well as for workloads (VMs or vApps).

The primary way to interconnect customer workloads between multiple locations is firewalled and implemented with VPN/MPLS or similar networking services, where T-Systems can ensure network quality.

#### Service Class

- VM provisioning
- VM image transformation
- Compute capacity (processing)
- Storage.
- Network.
- Random-Access Memory
- Virtual Networking & Security services incl. Load Balancing, Firewalling, Routing and other capabilities
- Secure Internet Access for customer workloads to establish a high security DMZ-like environment
- Provisioning of provider engineered and pre-configured VM templates
- Catalog based functionality
- Snapshot & Catalog based backup
- Access to the vCloud Portal and VMware vCloud APIs from Internet and/or Intranet

#### Service Level "Gold"

- Not yet available, planned availability 2013

#### Service Level "Silver"

- Operating hours: 24 x 7
- Attended operation time: 24 x 7
- Planned maintenance windows: 1 / month
- Max. downtime p.m.: 0.7 hours
- Response time: 12 hours
- Availability of the data center infrastructure: 99.9 %

#### Service Level "Bronze"

<ul style="list-style-type: none"> <li>– Not yet available, planned availability 2013</li> </ul>	
<b>Dynamic Services for Infrastructure (DSI) Central</b>	
<p>The "Dynamic Services for Infrastructure" offer concentrates on IaaS services within a private cloud environment (virtual machines are referred to in this context as "deployments" below). The core functions of the IaaS services consist of:</p> <ul style="list-style-type: none"> <li>– Virtual machines (VM)</li> <li>– Managed operating systems</li> <li>– Operating systems licenses</li> <li>– Data storage</li> <li>– VLANs</li> <li>– Internet Access for Deployments</li> </ul> <p>The customer's networks must be connected to the T-Systems data centre. All connections shall be routed in the data centre via a firewall into the network infrastructure. The precise implementation shall be planned and, if necessary, billed as part of project planning.</p>	
<b>Service Class</b>	<ul style="list-style-type: none"> <li>– VM provisioning</li> <li>– VM image transformation</li> <li>– Compute capacity (processing)</li> <li>– Storage.</li> <li>– Network.</li> <li>– Random-Access Memory</li> <li>– Managed OS provisioning (including backup, monitoring, security, patch and release management)</li> <li>– Backup of customer vmdk files</li> <li>– Provisioning of provider engineered and pre-configured VM templates</li> <li>– Secure Internet Access for customer workloads to establish a high security DMZ-like environment</li> <li>– Access to the Zimory Portal and Zimory APIs from Intranet</li> <li>– Access to the VMs from Intranet, optionally via reverse and forward proxy from Internet</li> </ul>
<p><b>Service Level "Gold"</b></p> <ul style="list-style-type: none"> <li>– Not yet available, planned availability 2013</li> </ul> <p><b>Service Level "Silver"</b></p> <ul style="list-style-type: none"> <li>– Operating hours: 24 x 7</li> <li>– Attended operation time: 24 x 7</li> <li>– Planned maintenance windows: 4 / year</li> <li>– Max. downtime p.m.: 4 hours</li> <li>– Response time: 12 hours</li> <li>– Availability of the data center infrastructure: 99.5 %</li> <li>– Availability of the managed operating system: 99.5 %</li> </ul> <p><b>Service Level "Bronze"</b></p> <ul style="list-style-type: none"> <li>– Not yet available, planned availability 2013</li> </ul>	

## Interoute

Provider	Interoute Communications Limited
Services	
Virtual Data Centre Service	<p>Interoute's product portfolio covers Connectivity, Communications and Computing. Computing includes traditional Colocation and Managed Hosting. VDC is the IaaS offering; as a pan European, true cloud service it perfectly aligns with the requirements of the Helix Nebula initiative.</p> <p><b>The Interoute VDC platform is the first cloud computing solution that can be deployed to meet private, public and hybrid cloud demands. It offers a flexible and a cost effective way to build your ICT infrastructure using Europe's largest ICT services platform.</b></p> <p>The Interoute VDC has all the characteristics of a public cloud; the ability to burst, pay as you go pricing and real-time deployment. It is the integration of Interoute's MPLS/IP network that enables VDC to be a cloud delivered using the Internet, or a part of your corporate ICT infrastructure eliminating the cost and resource constraints of a private cloud. MPLS has been trusted for the past 10 years by the world's largest corporations as a technology that provides "Internet any-to any access" but with label separated security. MPLS is routinely used in the most sensitive of environments given its absolute separation from the public Internet. By combining scalable elastic computing with the most trusted network technology Interoute has created a truly unique enterprise-computing platform, the first without compromises.</p> <p><b>Public Cloud Simplicity, Private Cloud Security</b></p> <p>The Virtual Data Centre is designed for the demands of Enterprise ICT and addresses the critical factors necessary for success when choosing whether to build a private or public cloud.</p> <ul style="list-style-type: none"> <li>– Security – built into the fabric of Interoute's MPLS network and secured within Interoute's Certified European Data Centres, Interoute's ownership of the physical network, combined with MPLS technology ensures a secure VDC cloud environment.</li> <li>– Privacy and Compliance – The Data Centre and network platform that supports VDC is Europe's largest. Owned and operated by Interoute this ensures location and access control can be selected to suit European businesses legislation.</li> <li>– Integration – VDC is built into the Interoute network so it easily integrates into customer connectivity and networking solutions.</li> <li>– Open architecture – Interoute VDC is hypervisor independent and open standards based, ensuring as technology changes VDC can change with your business.</li> <li>– Connectivity included – Interoute is able to offer the connectivity other public cloud providers charge for, at no extra cost. Connectivity between VDC environments in Geneva, Amsterdam, London, Berlin and Paris and connectivity between any Interoute customer's VPN site and VDC is free. The Interoute Virtual Data Centre replaces the need to buy equipment, power, co-location, network and people.</li> </ul> <p>Virtual Data Centre (VDC) is Interoute's Enterprise class Infrastructure as a Service (IaaS) solution. Providing on-demand computing, storage and applications integrated into the heart of your IT infrastructure.</p> <ul style="list-style-type: none"> <li>– On demand self-service: Via the Interoute customer self-service portal, customers can provision entire ICT solutions in their own Virtual Data Centre environment. The automated provisioning of the compute, storage and network resources are all achieved at the click of a button.</li> <li>– Easy access: Europe's most trusted network platform allows access to the virtual environment via both IP VPN, Ethernet and Public Internet access.</li> <li>– Resource pooling: The VDC service is based on a multiple partitioned platform, managed by software based allocation business rules and policies, meaning the IaaS resource pools created by</li> </ul>

the virtualization of the physical underlying infrastructure can be shared by multiple users, or dedicated to single enterprise.

- Elastic: The compute, storage and network resources used to create your ICT service within the VDC environment can all be used in a truly elastic way. Use only what you need and retain the flexibility to grow your VDC to an almost limitless level.
- Measured service: The VDC control centre on the Interoute customer web portal allows you to monitor the usage of the VDC service, reporting and controlling the allocation of resources.
- Customise your machine – Each Virtual Machine can be customised with varied, CPU, RAM and block storage to allow bespoke system builds and reduce over-spending on unused resources commonly found in fixed-specification cloud servers.
- Real not emulated IT. Transferring your ICT infrastructure to VDC doesn't require re- writing applications or compromising them for the cloud. Interoute's VDC allows you to create a cloud network to match your physical network.

#### **Built for Business**

- Interoute knows what it means to be trusted by an Enterprise so operates a dual centre 24/7/365 operations model.
- No additional costs to get to the cloud Computing, storage and applications into the heart of your IT infrastructure avoiding additional bandwidth charges.
- Block level, mountable, persistent storage in locations that you chose – complete data transparency means that you can create cloud based services and stay compliant to regional data regulations.
- Complete control of your appliances; root access, the ability to access storage directly, mount storage, reboot, power on and off the box and delete it.
- Size from one just server for one hour to 100,000 machines across multiple countries. Virtual Data Centre can be sized to precisely as small or as large as you need. Or you can simply let it grow and shrink as your business changes.
- Create your own Appliances or use prebuilt templates. Interoute has created the most popular appliances to get you started or you can simply create or upload your own and build a service catalogue for your business.
- End to End SLA Through our end-to-end service ownership and the ability to offer and integrate end-to-end connectivity, Interoute is able to offer a complete end-to-end SLA for customers, avoiding complex disputes and troubleshooting.
- Role Based Administration IT managers can use their own controlled builds creating and controlling cost of a service catalogue – your community can get up and running fast.

Interoute offers two purchasing models: Utility and Commit

- Utility is a zero commit model based on pay-as-you-go charging; each virtual resource has a charge per hour of allocation. The total amount of resource allocated is chargeable at the end of each month.
- Commit is a fixed monthly charge based on an assigned set amount of resources committed solely for each Commit customer's use. Customers are still able to burst beyond their committed resource as needs demand. Usage beyond these levels is charged according to the associated rate card. Customers can upgrade their committed resource level at any point.

Service Class	<ul style="list-style-type: none"> <li>– VM provisioning</li> <li>– VM image import and export</li> <li>– VM image transformation</li> <li>– Compute capacity (processing)</li> <li>– Storage. <ul style="list-style-type: none"> <li>○ Persistent</li> </ul> </li> </ul>
---------------	---

	<ul style="list-style-type: none"> <li>○ Storage level snapshots</li> <li>– Network. <ul style="list-style-type: none"> <li>○ Public Internet</li> <li>○ VPN (MPLS or IPSEC etc)</li> </ul> </li> <li>– Random-Access Memory\</li> <li>– RESTful API</li> <li>– Catalogue based functionality</li> <li>– Snapshot capability</li> <li>– Chose location <ul style="list-style-type: none"> <li>○ GBR, NLD, FRA, DEU, CHE</li> </ul> </li> </ul>
<p>VDC is a managed Infrastructure As A Service. The virtual resources, compute, network and storage are fully managed and supported by Interoute's operations centre. By default the customer themselves provision and manage their own VDC environments, Virtual Machines and appliances.</p> <p>Schedule 2n details the Terms and Conditions applicable to the VDC service. The Service Level Agreement for the Virtual Data Centre Service includes the availability of the underlying infrastructure and hypervisors. The design and management of the platform has enabled Interoute to guarantee 99.99% service availability. Availability is calculated by measuring the platform's response to ICMP Pings. Pings are made to each element of the platform. Service Credits offered in the form of additional VDC RAM, CPU and Storage resources. These will be credited to the account for future use. A Service Credit will be applicable and issued only if the credit amount for the applicable monthly billing cycle is greater than one Euro (€).</p> <p>The Customer Contact Centre (CCC) is responsible for your "live" services and is the single point of contact for all Interoute customers, providing first-line support 24 hours a day, 7 days a week.</p> <p>Interoute's Customers can contact the Customer Contact Centre (CCC) via email, telephone or the Interoute Hub (our customer web portal).</p>	

## The Server Labs

Provider	The Server Labs
Services	<p>Cloud Strategy consulting and Cloud assessment services.</p> <p>Cloud planning, design and implementation services.</p> <p>Cloud support and Cloud administration.</p>