Helix Nebula – The Science Cloud

Title: D9.1: A roadmap of future developments

Editor: EMBL, CERN, T-Systems

Work Package: WP9

Submission Date:

Distribution: PU

Nature: Report

Revision History:

27.02.2014 – First draft by EUTechStrategy distributed internally

03.03.2014 – CERN Feedback

12.03.2014 – EMBL Feedback

17.03.2014 – Second draft by EUTechStrategy

06.04.2014 – T-Systems, EMBL Feedback

14.04.2014 – Third draft by EUTechStrategy

01.05.2014 –EMBL, CERN, T-Systems feedback

06.05.2014 – Fourth draft by EUTechStrategy

08.05.2014 – EMBL, CERN, T-Systems feedback

09.05.2014 – Fifth draft by EUTechStrategy

15.05.2014 - EMBL, CERN, T-Systems, ESA, ATOS feedback

04.06.2014 – Final document by EUTechStrategy

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

Revision History		
Date	**Version**	**Contributor(s)**
18 February 2014	V1.0	Carmela Asero (EUTechStrategy)
11 March 2014	V2.0	Carmela Asero (EUTechStrategy) Rupert Lück (EMBL) Bob Jones, Rachida Amsaghrou (CERN), Jurry de la Mar (T-Systems)
03 April 2014	V3.0	Carmela Asero (EUTechStrategy)
14 April 2014	V3.1	Carmela Asero (EUTechStrategy) Rupert Lück (EMBL) Bob Jones, Rachida Amsaghrou (CERN), Jurry de la Mar (T-Systems)
01 May 2014	V3.2	Carmela Asero (EUTechStrategy) Bob Jones, Rachida Amsaghrou (CERN), Rupert Lück (EMBL), Jurry de la Mar (T-Systems)
06 May 2014	V4.0	Carmela Asero (EUTechStrategy)
09 May 2014	V4.1	Carmela Asero (EUTechStrategy) Rupert Lück (EMBL), Bob Jones, Rachida Amsaghrou (CERN), Wolfgang Lengert (ESA) Jurry de la Mar (T-Systems)
04 June 2014	V4.2	Carmela Asero (EUTechStrategy), Rupert Lück (EMBL),
06 June 2014	V4.3	Carmela Asero (EUTechStrategy), Rupert Lück (EMBL),

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

Contents

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

1. Executive Summary

This document summarizes the main outcomes and the agreed roadmap resulting from the two-year Helix Nebula (HN) project contextualized with the objectives set in the "Strategic Plan for a Scientific Cloud Computing Infrastructure for Europe"[1] published in 2011, which laid the ground for the Helix Nebula Initiative and the related EU coordination and support action.

The first phase of the project was devoted to the investigation and gathering of complex requirements by the scientific organizations on the demand side of the initiative, for the planned deployment of the three pilot use cases.
This analysis of requirements triggered the suppliers toward adapting their capacity and testing of a brokering tool - the Blue Box - to federate resources by commercial suppliers toward a more scalable cloud computing offering at IaaS level.

Another key aspect of the project concerned the federation of publicly funded e-Infrastructures and follow up of the consequent interoperability aspects. [2]
The added value of the Helix Nebula cloud business proposition in comparison with current major cloud vendors lies in its federation and integration characteristics.

In addition, this document reports about the outcome of the activity linked to the definition of a lightweight governance model and its evolution from an initial phase, through an operational phase, with outlook toward the long-term vision of the Helix Nebula governance, including policies for trust and possible funding streams.

A set of proposed business models for the Helix Nebula services is also summarized, to highlight the planned evolution from an initial offering of generic IaaS cloud offering toward an envisaged more complex and added value INFOaaS model.
The business models also take into account results of a cost comparison analysis carried forward for the CERN ATLAS use case.

[1] "Strategic Plan for a Scientific Cloud Computing Infrastructure for Europe" http://www.helix-nebula.eu/index.php/news/30/61/Strategic-Plan-for-a-Scientific-Cloud-Computing-Infrastructure-for-Europe.html

[2] http://ec.europa.eu/isa/documents/isa_annex_ii_eif_en.pdf

A final roadmap, building on project results, presents steps needed to complete achievement of the original goals and the pursuit of further and more ambitious objectives to support the growth of a competitive cloud market in Europe, thus allowing EU based cloud industry to play a leading role in this global sector.

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

2. Introduction

This document highlights the key developments and achievements of the Helix Nebula projects and sets up a roadmap for future activities in line with a revamped strategy for the Helix Nebula initiative outlined in D9.2: Strategic Plan for a Scientific Cloud Computing Infrastructure for Europe: Three years on.
The main outcome of the activities within the Helix Nebula initiative is reported here, with particular focus on the two-year EC funded project supporting it.
The Helix Nebula project and the activities towards establishing a federated pan-European multi-tenant multi-vendor cloud gained enormous visibility. It helped to attract new members and to increase the consortium.

The Helix Nebula initiative expanded remarkably from 20 to 37 members during the last two years including major Europe based cloud providers, world leading research organizations, network operators, pan-European publicly funded e-Infrastructures, technology providers and innovative SMEs.
The initiative has forged a strong and enriching cooperation allowing commercial vendors to benefit from technical expertise of highly qualified staff involved in development and the operation of large e-Infrastructures in the research and education sector.

Likewise, members of publicly funded e-Infrastructures participating in HN activities acquired a wider business perspective by interfacing and analysing complex governance and business model scenarios for a hybrid environment in which different revenues and funding streams, including pay-for-use in publicly funded research, are foreseen.
This document starts by reporting on results achieved then identifies areas of future work spanning beyond the end of the Helix Nebula EC project in order to continue and enhance this partnership, by capitalising on knowledge acquired since the publication of the strategy paper back in 2011.

3. Roadmap: From Infrastructure as a Service to Information as a Service

3.1. From requirement gathering to cloud provisioning

Helix Nebula took up the difficult challenge of investigating requirements across multiple users, and of implementing actions to ensure a successful deployment of the large scale pilot use cases by CERN, ESA and EMBL.
These challenges were faced simultaneously for the broker and the suppliers with the common aim to improve functionality and to create competitive services.

After initial use of a light-weight pilot application form for the proposal of the first three flagships and in view of additional future flagship proposals, the need emerged for a more elaborate requirements gathering methodology run by CloudSigma. It led to production of D3.1 Requirements Definition Template[3], allowing a more structured requirements gathering. This activity set the basis to match the specific requirements for the heterogeneous flagships with the capabilities of the individual providers and it helped to setup the initial testing of flagship deployment on the different clouds of individual providers (Proof of Concept phase).
A key result from the PoC phase was that a means of federated access to the individual and heterogeneous cloud environments was needed. This required the definition and set up of a common access interface (broker) to the Helix Nebula multi-cloud infrastructure. The requirements specification for this interface, the so-called Blue Box is based on the work of the supplier's TechArch team with support provided by the demand-side.

Through this intense work, where the effort was funded by the suppliers in the team and with the support of WP3 within the Helix Nebula EC project, the following important milestones were achieved:

[3] http://cdsweb.cern.ch/record/1484437

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

- The cloud service broker, the so-called Blue Box, was put in place and currently supports:
    - a single common interface for the customer to interact with all suppliers;
    - an image library.
- Service descriptions, service levels, contract terms and conditions and pricing for the use of the cloud services.

From a customer perspective the Blue Box provides a single, common interface to all suppliers. The customer can interact with each supplier through the Blue box without having to cope with specific differences in the interface of each of the supplier clouds. The Blue Box achieves this by the implementation of an abstraction layer and functions that interface with each supplier cloud through APIs.

The broker currently neither has an active role in matching the users' requirements to the services of the suppliers available on the marketplace, nor during the billing and settlement for the delivered infrastructure services. All activities are performed by either the customers or the suppliers. The conclusions to be drawn from the practical experience by the customers and suppliers of the deployment will be used to determine the suitability of this distribution of activities.

Scenarios had been implemented initially based on two blue boxes, the open source software-based SlipStream and the commercial Enstratius solution. Both were connected to multiple providers and tested with the deployment of demand-side flagships.

The flagship deployment results were satisfactory and summarized below:

- **CERN** was able to run simulations previously executed on the Worldwide LHC Computing Grid by deploying the ATLAS experiment flagship application on the Cloud.
- **EMBL** successfully deployed and tested their novel software pipeline for large-scale genomic analysis using real world large genomic data sets.
- **ESA** successfully tested large-scale data processing and dissemination for its radar satellites using different cloud provider infrastructure.

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

Following positive results for the pilot use cases, further actions are needed to allow effective adoption from organizations on the demand side. These actions include the development of a more mature version of the Amazon Web Services EC2 connector, the so-called 'EC2 bridge'. The EC2 bridge is essential for the EMBL flagships and PIC – a new flagship use case illustrated in next paragraph – also considering use of the StarCluster software (http://star.mit.edu/cluster/), which supports the needs for cloud-based elastic high-performance computing of both institutes. A mature EC2 bridge has the potential to offer similar benefits to future Helix Nebula's users who want to employ EC2-based software to use Helix Nebula's cloud capacities. A new version of the EC2 bridge has been made available to the demand-side butneeds further testing.

During the lifespan of the project, the cloud computing infrastructure management platform Enstratius, initially deployed and planned to serve as one of the future Blue Boxes, changed ownership structure following the acquisition by Dell. Blue Box provisioning, mainly offered via SlipStram, as well as cloud integration/configuration had some delays. The overall schedule slipped and required significant additional effort from the suppliers and the demand-side. During the pilot testing phase the level of integration of the blue boxes with the individual cloud platforms still depended on human intervention and involved direct interaction with the cloud providers instead of allowing programmatic and automated access to the different cloud services.

SlipStream is being upgraded in view of its deployment in the Helix Nebula (HNX) marketplace. An agreed service catalogue and a timeline for future updated releases of the SlipStream Blue Box has been prepared. Key development actions to be addressed in the short term future are:

- The introduction of an HNX API on the northbound interface to be based upon a widely used specification (possible candidates are OCCI, CIMI, Amazon EC2, OpenStack, Google)
- Identification of a way to remove the overhead by having to port the server images to the various cloud providers' platforms.

Talks with Enstratius are ongoing and this relationship should be further pursued to consider possible inclusion in the Helix Nebula Marketplace.
Further effort is still required in the following areas:

- Development of a standard set of APIs for suppliers and consumers; this needs to include relevant APIs of current global players offering ease of migration;
- Deployment of a federated identity management system offering a single sign-on facility to access cloud services across multiple-suppliers;
- Tools for automation of cloud management processes to provide on-demand services, resource pooling and rapid elasticity across cloud providers.

Concerning security, Helix Nebula partners have elaborated a Security Management Road Map structured around major milestones to be implemented to achieve the ultimate goal of ISO 27,001 accreditation. Nevertheless, future activity will have to address policies, tools and procedures to increase trust and security of overall cloud infrastructures and services as well as advanced automatic and dynamic deployment and more effective configuration and management of services.

3.2. Deployment of new flagship use cases

The second year of the Helix Nebula project welcomed the PIC Neuroimaging Centre on Cloud as a new member and the PICNICC flagship as an additional use case for the Helix Nebula suppliers.
Medical research in neurodegenerative diseases, both academic and industrial (pharmaceutical), greatly benefits from the results of MRI and PET data processing and needs a user friendly platform for computational services. Yet, doctors and researchers are not interested in computing per-se as they need a service able to turn raw image datasets into feature-rich analysis datasets as quickly as possible.
The scientific objective of the HN-PIC partnership is to improve the speed and quality of research in order to find surrogate biomarkers based on brain images. The technology vehicle for this purpose is PICNIC (PIC Neuroimaging Centre), which is a Web-based platform running on a cluster that allows researchers to easily store, organize and process data with significantly reduced turnaround time (from several months to a few days).
Future challenging use cases may come from ESFRI projects in their implementation phase as well as from exploitation of Copernicus data as reported in section 3.5.

3.3.Cost Analysis

Estimating the true cost of an IT service is very challenging especially for large public organizations. A cost assessment exercise has been run for CERN Atlas use case but constraints in time and effort did not allow an in-depth analysis taking into account a diverse number of elements.

Furthermore some specificities of CERN, such as heavy use of publicly funded internet network and access to very competitive energy rates, did not allow an accurate assessment of CERN's overheads for the specific use case.

Between 2011 and 2013, the e-Fiscal project[4] analysed the costs of the current European dedicated High Throughput and High Performance Computing (HTC/HPC) e-Infrastructures for research, and compared them with the closest equivalent commercial leased or on-demand cloud offerings and evaluated the trade-off between in-house solutions and commercial procurement of cloud services.

Findings from e-Fiscal project found that in-house HPC/HTC e-Infrastructures are cost-effective with high utilisation rates & depreciation rates with personnel accounting for ~50% of total costs and a ratio CAPEX/OPEX=30/70%. However, larger sites have in general less operational manpower per core.

Results from e-Fiscal are also in line with literature (Hawtin et al. 2012), Magellan (2011), and show relevance of economies of scale and utilization rates.

Analysis of cost elements gathered during cost assessment conducted within the Helix Nebula project on the CERN site for two reference services, compared with offering from a generic cloud services such as the one offered by AWS EC2 highlighted competitiveness of services offered by Helix Nebula suppliers.

The most cost-efficient offer for a one year commitment received from Helix Nebula suppliers was approximately 3.4 times as cost-efficient as AWS (3.8 with a three year commitment). This is a very interesting figure, especially for those organizations considering AWS rates are generally the most competitive within the market.

This analysis shows that for large sites such as CERN and services with high and quite stable utilization rate, in-house solutions are still the preferred approach, but that complementary services offered by Helix Nebula suppliers could be more competitive than ones from current market leaders.

[4] http://efiscal.eu/

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

Keeping a competitive edge on cloud service offerings will require continuous effort to achieve a very dynamic, scalable and elastic cloud environment.

3.4.Production Platform and Marketplace

At the third Helix Nebula General Assembly held at EMBL in Heidelberg in September 2013, commercial suppliers decided to set up a production platform and launch of the "Helix Nebula Marketplace (HNX)" in 2014.
An intense activity started to:

- Set up the commercial agreements, policies and processes for service management within HNX;
- Build an initial common service catalogue accessible through an user friendly interface and portal;
- Prepare a master plan integrating the necessary building blocks for the first deployment of the production environment expected in March 2014;
- Analyse procurement procedures and rules of the demand side to set up effective appropriate contractual template and billing functions as framework contracts for public organizations require a prime contractor to manage the contracting with a broker and multiple cloud providers.

It was also decided to start the new HNX platform initially based on SlipStream version 2.1 BlueBox. Programmatic access to the HNX resources would need the EC2 bridge to be in place from the beginning (see Fig. 1).
A demo of the production environment was shown at Cloudscape VI on 24-25 February 2014 in Brussels. The first deployment of the marketplace has been planned around the following basic structure.

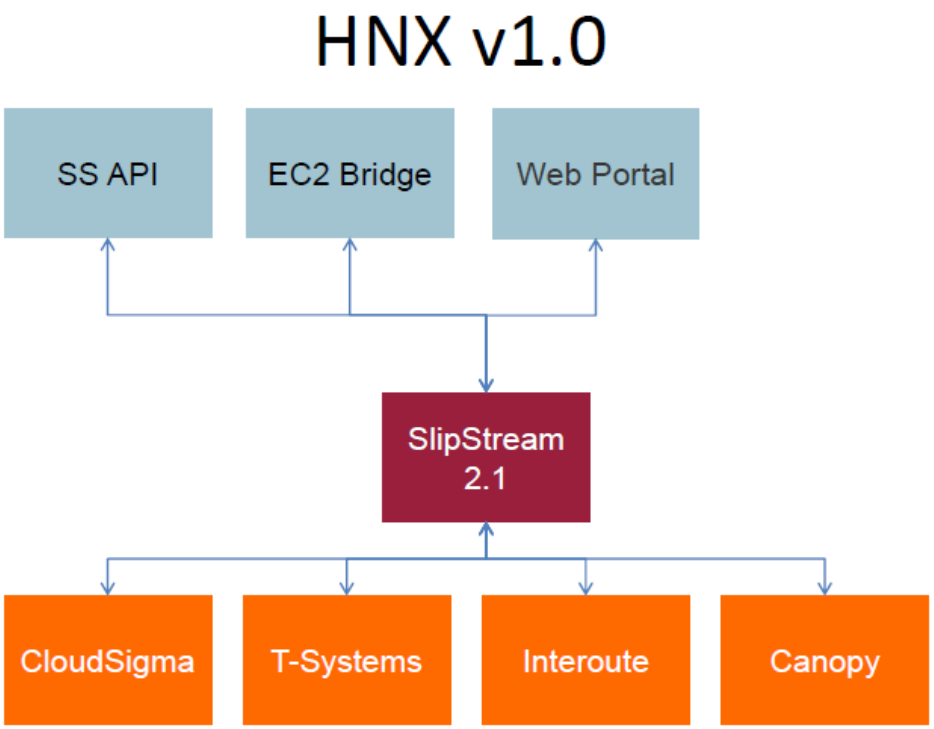

Figure 1 First version of the Helix Nebula Marketplace (HNX)

While the initial HNX platform will be initially be made available to the existing demand-side organisations for further QA testing, it is planned to open the platform to the public later in the year to attract further customers and suppliers (2 new users and 3 more suppliers have already come forward).
At the same time the SlipStream blue box will be upgraded with aricher feature set. A second blue box, such as DELL's Enstratius, will be integrated with HNX. Similar to the EC2 bridge the Slipstream API will be developed further to form an HNX API that will allow programmatic access to the blue box functions (see Fig. 2).

PICNIC on Cloud, to be deployed within HNX, would facilitate the move to a full-cost, commercial charging model, which would benefit the expansion of the use of PICNIC by private industries (mainly pharmaceutical). This would allow:

- Quick turnaround in processing large amount of MRIs or PETs;

- Offer a user-friendly interface to submit jobs to the Cloud without specific knowledge of scripting, OS specificities, etc.

At present, the initial requirements for the PICNIC flagship have been analysed and deployment of HNX use case will continue after the end of the EC project. PIC will initially make use of a sandbox with the StarCluster toolkit for the EC2 bridge, thus being able to profit from the environment of the EMBL use case.

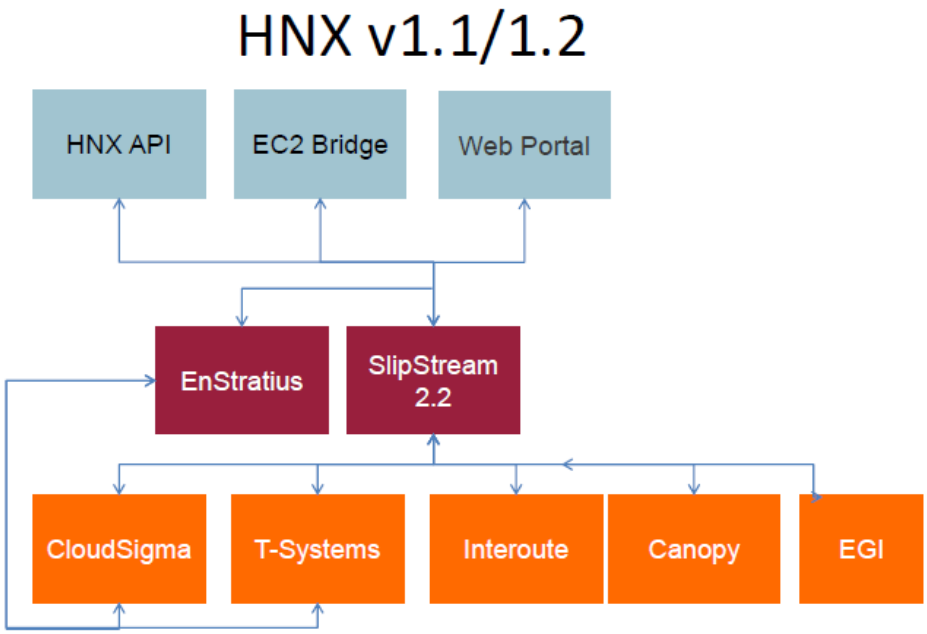

Figure 2 Second version of the Helix Nebula Marketplace (HNX)

3.5. Towards INFOaaS

The establishment of HNX, the first Helix Nebula production environment, will allow the marketplace to seize opportunities emerging by exploitation of attractive data repositories and services that could be integrated to form an initial basis for an emerging Helix Nebula INFOaaS platform.

Transition toward InfoaaS is in a more advanced phase for ESA SPSS platform use case, with roll out SSEP to users foreseen to become operational during 2014.

ESA is actively working on involving communities in SSEP, affiliated to ESA data and interested in joining thematic exploitation (eg. using ESA/Copernicus data as an entry point).

The Copernicus data could be explored as an attractive option to launch an InfoaaS model, by building on collaboration with ongoing FP7 projects such as DORIS, EC funded Copernicus downstream service, and LAMPRE.

ESA's service named "Doppler Orbitography and Radio-positioning Integrated by Satellite" (DORIS) is a concrete starting point for the adoption of an ecosystem. DORIS is offered by the ESA Earth Observation (EO) Mission and "it is an advanced downstream service for the detection, mapping, monitoring and forecasting of ground deformations, at different temporal and spatial scales and in various physiographic and climatic environments" (Guzzetti 2013).

DORIS integrates traditional and innovative Earth Observation (EO) and ground based (non-EO) data and technologies to improve our understanding of the complex phenomena that result in ground deformations, and to foster the ability of Environmental and Civil Protection authorities to manage the risks posed by ground deformations. DORIS, similarly to Helix Nebula, is led by a consortium including a partnership of leading research institutes and commercial providers, and is expected to stimulate European competitiveness and sustainable development.

Complementary to DORIS is the FP7 project LAMPRE, which proposes to execute innovative research and technological developments to increase GMES limited operational capacity to cope with triggered landslide events and their consequences, in Europe and elsewhere. LAMPRE will enhance landslide risk mitigation/preparedness efforts and post-event-landslide recovery and reconstruction activities, in highly vulnerable geographic and geologic regions. The project improves the ability to detect/map landslides, assess/forecast the impact of triggered landslide events on vulnerable elements, and model landscape changes caused by slope failures.

ESA plans to develop a SuperSites Exploitation Platform (SSEP) "Science Cloud" building on a federation of contributors in order to establish an Earth Observation market place, to be integrated in Helix Nebula, which now includes the Iceland supersite and CEMS as well as

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

US data centres. In this vision, networking connectivity via GÉANT is crucial for reliable service delivery as well as definition of an attractive pay for use model for SSEP running on Helix Nebula

Shift to an InfoaaS model would be more easily be achieved by:

- The integration of further key data sets could support industry and research through aggregating data from currently less related fields. Such key data sets could for example be contributed by an integration of the Embassy cloud offering access to world-class biomedical data repositories at EMBL-EBI. These data are of interest to the wider life science communities as well as to research in the medical field as well as in pharmaceutical and agricultural industries.
- The historical data sets from the medium-range weather forecast such as the one that is hosted by the ECWMCF
- Observation data from the IOC/UNESCO Ocean & Coastal Information Supersite

Such services would release their full potential from cloud services by a quick evolution from a "generic cloud offering" business model - planned for the initial operation of the Helix Nebula Marketplace - toward an INFOaaS model. However, this transition requires a traceable & comparable service catalogue, transparent invoicing, automatic service discovery and selection, as well as open source middleware.
To facilitate the users' uptake and cross boundary cooperation, LAMPRE will test products and services in a range of physiographical and geographical regions, and will use the advice of a specific international Stakeholder User Group (SUG).

Such large applications are expected to attract a critical mass of both users and new suppliers to the marketplace and to grow the offering toward different communities and typology of users, in order to rapidly achieve a large critical mass to leverage investment from the industry and complementary support from EU and national funding agencies, thus boosting the cloud computing sector in Europe.

The exploration of challenges and opportunities of evolution toward a wider "cloud of public services" concept should also be included in the activity roadmap for the Helix Nebula initiative, in order to attract more players building on the Helix Nebula infrastructures. These players would also be able to develop PaaS and SaaS solutions, providing a collection of public services & building blocks which can be:

- ✓ offered in an open & interoperable way
- ✓ re-used and combined by public administrations and third parties

The Helix Nebula marketplace also aims at becoming the first choice for cloud services burst needed by ESFRI projects and ESFRI cluster initiatives. In the next few years a number of projects in the ESFRI roadmap will start their implementation phase to serve a wide and heterogeneous base of user communities, who will largely need reliable and scalable cloud services interoperable with existing e-Infrastructures within an envisioned e-infrastructures Commons Marketplace.
The Helix Nebula roadmap foresees use of federated cloud offered by HNX as an underlying IaaS infrastructure, with SaaS and PaaS layers built on top of it to support:

- InfoaaS
- Application Crowd, with focus at ICT SMEs(European Industry)
- Open City Services for Smart City apps (public sector)

To achieve these objectives, the Helix Nebula partners should focus on engagement of data sources from the weather, environment, earth observation systems, sciences, bioinformatics, healthcare and statistics.

3.6.Interoperability with e-Infrastructures and related actions

During the Helix Nebula project, an intense activity was devoted to analyse and follow up interoperability aspects framed around five different levels (political context, legal, organisational, technical and semantic interoperability) following the structure proposed by the European Interoperability Framework 2.0.
In its first year, the Helix Nebula project identified 17 interoperability recommendations, and then translated into specific actions to be carried forward in the second year and beyond the lifespan of the project, according to agreed levels of priority.
A good number of such actions have shown relevant progress, especially in the area of organizational interoperability, identification of level playing field for e-Infrastructures and commercial providers, as well as service management. For political and legal interoperability, Helix Nebula is progressing in analysis and resolution processes on contractual aspects involving both commercial providers and e-Infrastructures.

Helix Nebula suppliers, public and commercial, are also exchanging view on alignment with political development in the EU in area such as data protection and security.

The outcomes of activities on interoperability were highly satisfactory and initially marked by the use of GÉANT to connect commercial providers in running successfully the Helix Nebula flagship use cases. In the second year, the deployment of the ESA SSEP use case on the EGI Federated Cloud, via GÉANT, following previous successful proof of concept on commercial resources, opened the way for the next phase: the test of the CERN ATLAS use case in a hybrid environment combining publicly funded and commercial cloud resources. These results demonstrated the added value of an interoperable hybrid cloud offering, not only for research communities but also for a wide range of public cloud based services.

EGI is advancing a service production phase for the EGI Federated Cloud, with foreseen launch in spring 2014, while GÉANT is developing use cases for the delivery of Cloud services to campus-based users within its Open Cloud eXchange (OCX). Collaboration is also in progress with GÉANT to have the eduGAIN[5] inter-federation identity management service working in HNX marketplace. In order to ensure progress in interoperability, it is crucial to secure channels and models to exchange information and to collaborate on common issues, beyond the end of the current Helix Nebula EC project.

To create the conditions for a fruitful continuation of work on interoperability, it is suggested to:

- Ensure information exchanges and set up of task forces among representatives of commercial suppliers and publicly funded e-Infrastructures in areas such as security, support, incident management. Such activities should be foreseen in the individual roadmaps of HNX, EGI and GÉANT.

- Engage with main e-Infrastructures for an agreed roadmap toward the establishment of an e-Infrastructures Commons Marketplace to boost the European Research Area.

[5] http://www.geant.net/service/eduGAIN/Pages/home.aspx

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

- Identify main e-Infrastructure events, involving a large and wide range of user communities during the next years, to discuss further alignment and integration of respective development plans according to users' need.

- Investigate economies and synergies deriving by sharing horizontal legal, economic, and policy expertise, common training and communication activities, when such aspects invest the overall Helix Nebula initiative and are of common interest to a number of its members.

In the second year of the Helix Nebula project, talks also started with EUDAT and PRACE aiming at better integration of e-Infrastructures offering through HNX.

Discussions are ongoing with EUDAT to determine if the data services being developed by EUDAT can be made available within HNX..

Concerning HPC, a constructive dialogue was initiated with PRACE.
In March 2014 the management team of Helix Nebula initiative met PRACE management board representatives to discuss possible areas of collaboration.
The Helix Nebula environment already includes the HPC site of ATOS Spain, and more HPC facilities are likely to join Helix Nebula initiative in the future, so training activities for use of HPC machines was identified as an area for joint programmes and initiatives.

Current PRACE trainings are open to anyone in Europe and not limited to proposals having passed successfully a PRACE peer review call, and PRACE trainings are similar to the ones needed for use of the HPC environment in Helix Nebula.
Data management was also addressed as a topic of interest of PRACE users. PRACE offers data storage, especially for analysis in Tier-1 machines, for a limited time, after which the users have to take back and manage their data.

Applicants who are not successful at a PRACE call to obtain allocation of HPC machines and PRACE users needing data management services could benefit from Helix Nebula's offering and from its integration with PRACE.
Identity management and data staging could be common activities, in the context of an e-Infrastructures commons marketplace including major public e-Infrastructures and commercial providers.

Helix Nebula has contributed to the democratisation of access to HPC, notably for SMEs and also PRACE is working on facilitating access to HPC and cloud services to innovative SMEs in science and applied research. In this domain PRACE suggested Helix Nebula to engage with the EU project Fortissimo, aiming at enabling European manufacturing industries to benefit from the efficiency and competitive advantages from use of simulation on HPC based cloud infrastructure.

3.7.Governance

To organise an initiative made of public and private partners, structured around a demand and supply side with different roles and involvement, the Helix Nebula partners decided in 2012 to start with a "light-weight" governance structure as shown in figure 4. The initial objectives were to ensure:

- A common strategy and mission
- Governing structure including demand and supply-side boards and a management team
- Decision-making procedures
- Rules for adding and excluding partners
- Rules for information exchange

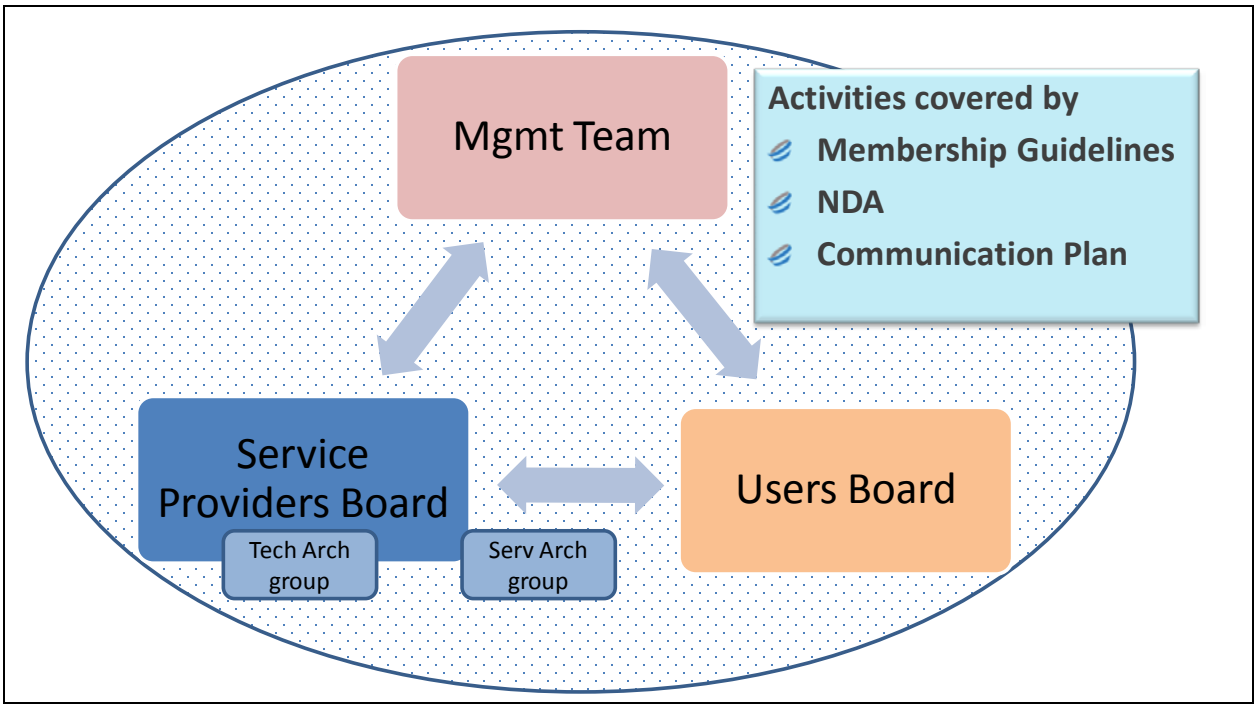

Figure 4 Initial Helix Nebula governance structure

These objectives were documented and confirmed in April 2012 by all founding parties, including CERN, EMBL and ESA from the demand-side, and Atos, Cap Gemini, CloudSigma, CSA EMEA, EGI, Interoute, Logica, Orange Business Services, SAP, Sixsq, Telefonica, Terradue, Thales, The Server Labs, Trust-IT, T-Systems and Universidad de Computense Madrid from the supply-side.

Since then, the Initiative has considerably grown. The following new members have joined: The following new members have joined: CNES, CNR-IREA, DLR, ECMWF, ESO, PIC and UNESCO from the demand-side. AWST, CNRS, DANTE, Emergence Tech Ltd., Ifremer, Nextworks, Switch, VisioTerra, Memset, Ultimum Technologies, Yandex, Nephos Technologies, INDRA, and CloudEO for the supply-side.

Relations among all parties involved in developing Helix Nebula are currently governed by basic legal instruments, namely a non-disclosure agreement (NDA) and a membership agreement that all partners sign to join the Helix Nebula consortium. 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 new type of users of Helix Nebula services, in the coming future, questions arose about what should be required from a customer (a research institute or governmental organization) when they want to use the Helix Nebula services. The following guidelines were then adopted:

- 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'.

Following the launch of a commercial production platform, governance aspects become even more critical for the successful deployment of the Helix Nebula marketplace, its effectiveness, sustainability and ability to deliver added value service to customers and end users.
Some governance scenarios with different assignments of roles and responsibilities among participating partners have been investigated as well as related mechanisms for decision-making and implementation of consequent actions. These scenarios are documented in deliverable D8.1 "A story of governance models for public-private cloud partnerships"

To enable a smooth exchange of necessary information while avoiding risks of anti-competitive behaviour, Helix Nebula partners committed to follow a set of Antitrust Guidelines governing exchange of information and extent of collaboration, to avoid potential issues by applicable competition law, which has been a major concern for all the players since many of them are major IT players in Europe. Helix Nebula has adopted these Antitrust Guidelines for itself, its members and all partners, customers and other

participants, as guidance in connection with participation in Helix Nebula activities, which are subject to strict compliance with the "Helix Nebula's Antitrust Guidelines".

As integration of Helix Nebula with existing e-Infrastructures is one of the key objectives, the governance structure and decision-making process have steered towards first achievements with e.g. the joint work performed with EGI.eu on the Federated Cloud services and the setup of a pilot communication infrastructure over GÉANT to exchange large scientific data-sets between public-funded and commercial cloud data centres.

Therefore, governance for Helix Nebula needs to provide a well-balanced framework of rules and processes to manage expectations, provide strong leadership and enable good decision-making, by taking quite a diverse user and stakeholder community into account. In particular, the building of a public-private ecosystem with stakeholders from the public sector and business, each having differing requirements towards governance calls for new approaches not yet common in Europe. Examples in the USA, parts of Asia and other European Industry Sectors e.g. Aviation show that good governance amongst public-private stakeholders can be achieved and lead to innovation. A typical business approach to address such complex requirements is to keep a strong pragmatic focus on the benefits that can be established for users, stakeholders, public sector and businesses and create WIN-WIN scenarios.

In D8.1 "A story of governance models for public-private cloud partnerships", it is proposed to initially pilot two Brokers, one operated by a commercial entity with focus on short term availability and one by an independent neutral entity for best possible uptake in the public sector.
Effective delivery of brokerage services and (electronic) processes known from retail industry will be offered while maintaining an open competition and adopting an innovative approach towards management of specific requirements by users and service providers.
The figure (fig. 5) below outlines the requirements for the future Helix Nebula governance:

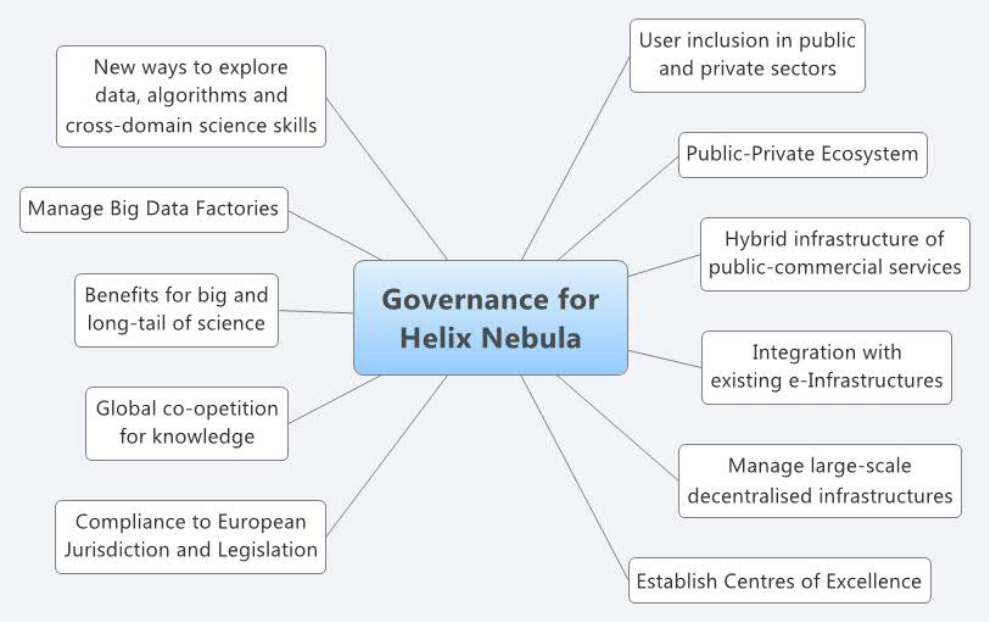

Figure 5 Governance requirements for Helix Nebula Initiative

The proposed governance elements and their relations are shown in the next figure and will be described in more detail in the following sections.

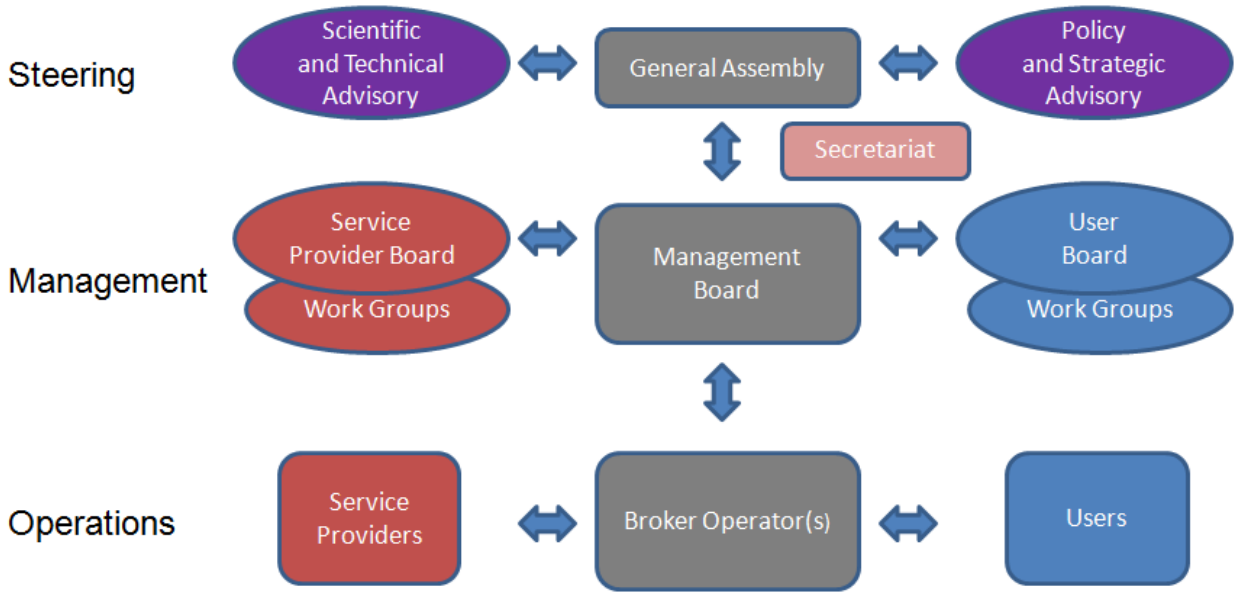

Figure 6 Proposed Governance Structure for Helix Nebula

3.8.Business models

The Helix Nebula partnership has demonstrated that it possesses all the expertise and technical blocks required to offer generic cloud platform at the IaaS level, but the shift to a production service must be fast and effective in attracting the needed critical mass to trigger the related strong network effects.

The service is foreseen to start with existing public organisations on the demand side as they have the necessary organisational maturity and collected requirements to start their procurement processes. These would be independent, but based on common technical specifications & contract terms and conditions for this initial generic cloud platform.
Once the basic production system is in service, it can also act as a platform for innovation with other actors (SMEs, downstream industry, etc.) using it to develop new services.
In addition, the Research Accelerator Hubs, as described in the EIRO forum e-infrastructure implementation plan, would allow operators such as CERN to perform aggregated procurement on behalf of their research partners, thereby simplifying the overall procurement model.

Having put "Generic Cloud Computing for European Science" business models in production for 2014, work must progress in parallel with the supply-side and demand-side to develop higher-level services and expand towards the "Information as a Service" business model, ideally to be achieved in 2015 with broadened scope and engagement.

The business model "Generic Cloud Computing for European Science" (see figures 7 & 8 below) builds on the successful proof of concept run during the two years pilot phase of Helix Nebula and it is based on the needs of European big-science and brings together leading IaaS providers and research centres like CERN, EMBL and ESA in order to provide data capture and processing that elastically meet big science's growing demand for computing power.

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

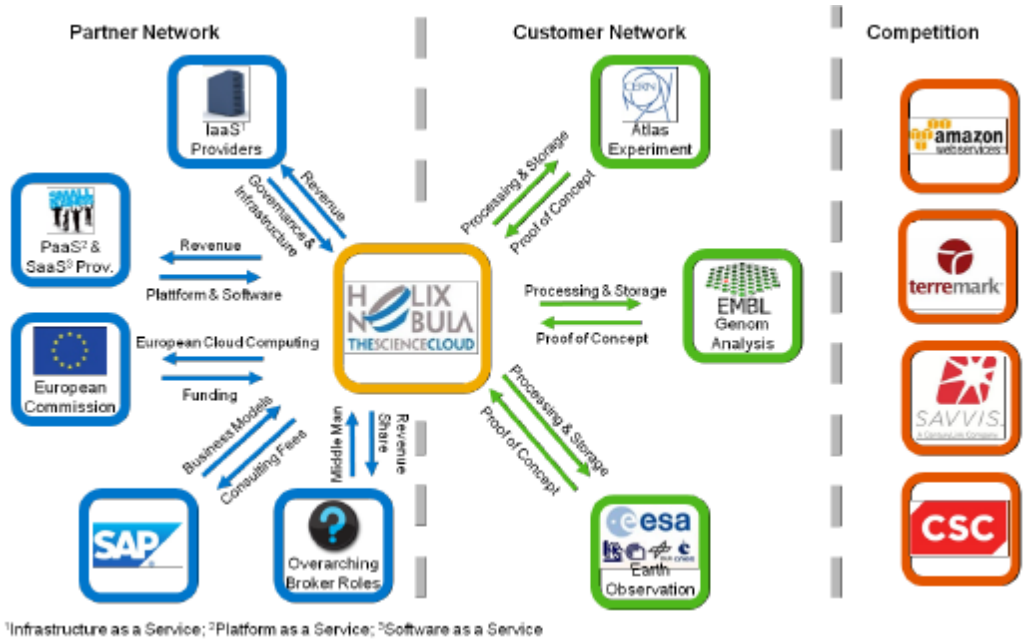

Figure 7 Generic Cloud Computing for European Science (Network View)

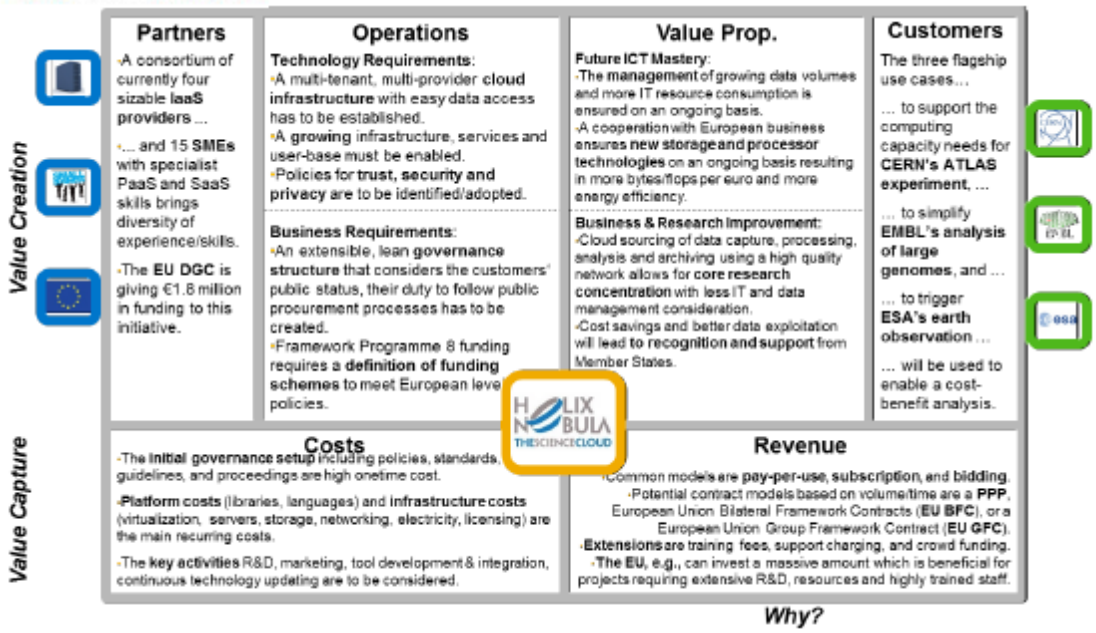

Figure 8 Generic Cloud Computing for European Science (Enterprise View)

This model is planned to be a stepping-stone Helix Nebula, in order to get to the business model “Information as a Service”.

The Information-as-a-Service Business Model significantly extends the Generic Cloud Computing Business Model. Infrastructure as a Service becomes just a means and prerequisite to capture, process, analyse and archive highly attractive data from scientific organisations, thus creating an opportunity to cooperate with further data providers in order to enrich the data in its context. The selling of resulting data sets and knowledge is evaluated as the most promising business model in terms of market need, impact on critical mass, differentiation, and thought leadership.

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

To reach such objectives, the Helix Nebula project identified a four-step roadmap to implement the business model.

The first stage is the implementation which requires a sophisticated added value analysis. In the first phase, awareness and gaining adoption by the initial set of customers will be a critical component of success, and therefore a challenge.

When, as a second step, operations start after the implementation, members are required to join and to be informed that this is a new service. So they will perform more scrutiny, before adopting a new service.

The third phase leads toward maturity as the demands on the service environment will grow. Demands will be based on volumes and usage, but also based on new requirements. So the challenge becomes capacity of scaling. The investment to create a very scalable, trustable, proven infrastructure will be critical. If the system environment fails even once, users can lose confidence. Creating a strong brand, high quality, and hiring highly experienced employees will have higher costs than revenue, but it will be critical because a platform like this cannot afford to ever fail.

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

Fourthly, as services become more common, as with Amazon Web Services, a proper governance structure will be needed to review the business model to address both competition and innovation, and this can be challenging. The life cycle management of those three phases and putting a strategy in place is crucial.

INFOaaS is subject to network effects as its value increases for each user as the number of users grows. The more data providers share data, the more complete the services will become. INFOaaS is a multisided platform ecosystem. This kind of ecosystem faces the lifecycle challenge of (1) fast implementation; (2) adoption, (3) scaling, and (4) competition (cf. Figure 15). Those steps are now applied to INFOaaS in order to provide a roadmap for lifecycle management.

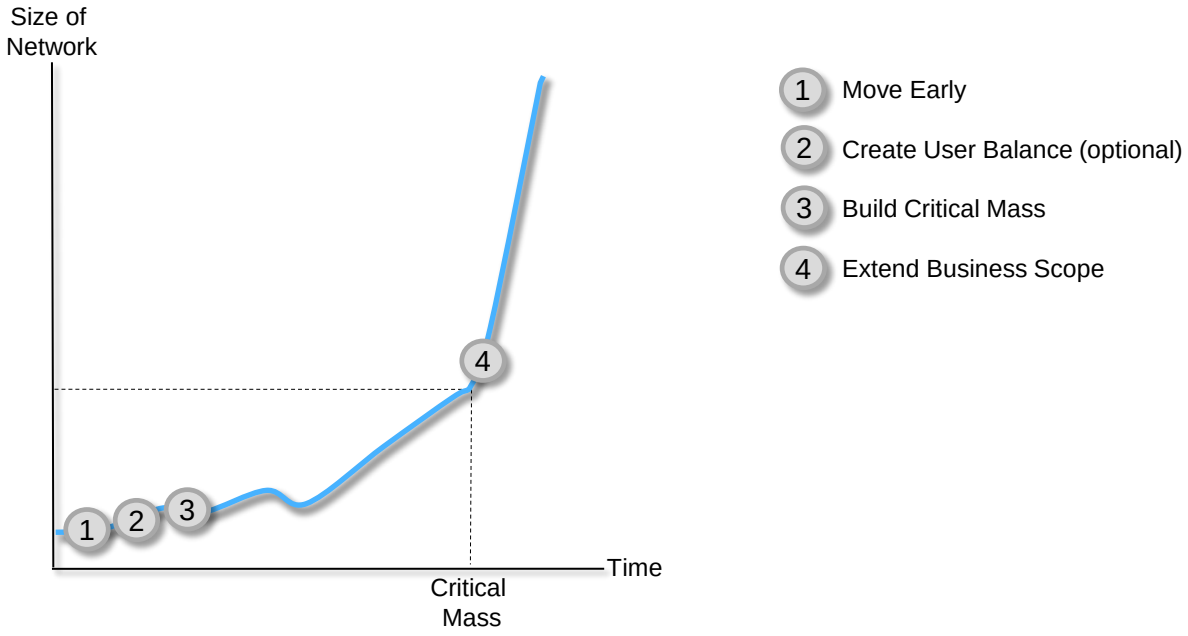

Figure 9 Strategic Steps towards Market Tipping

4. Conclusion

The intense and fruitful work on the deployment of the flagship use cases, the interoperability challenges following the positive engagement of publicly funded e-Infrastructures such as GÉANT and EGI, the effort in identifying suitable business models for this unprecedented public private partnership federating heterogeneous resources from a variety of suppliers with the key contributions of brokering partner and technology providers, including innovative SMEs, have demonstrated the validity of the initial vision and strategic plan.

Current offerings by large, mainly US based vendors still appear inadequate in satisfying European and global users' demands, especially in terms of interoperability, applicable law, data portability, privacy and confidentiality.[6]
Furthermore The Washington based Information Technology & Innovation Foundation (ITIF) issued a report last December that suggested that US cloud computing suppliers could lose $35 billion by 2016 due to this lack of trust. [7]

This scenario offers huge opportunities for European cloud players to leapfrog their role in this global market.
In this context, the European Cloud Partnership has provided a useful discussion forum to enhance collaboration and encourage the economic boost that cloud could provide for Europe; however attention must be focused on exploiting opportunities, rather than just avoiding risks.

The activity carried forward with enthusiasm by the Helix Nebula partners during the pilot phase enriched the initiative with valuable experience which must be retained and exploited to the fullest by addressing the following actions[8]:

- Monitoring evolution of political context and assess opportunities/issues on level playing fields of Helix Nebula's commercial and publicly funded providers also in light of opportunities deriving from pay-for-use schemes for e-Infastructures

[6] http://ec.europa.eu/information_society/activities/cloudcomputing/docs/quantitative_estimates.pdf
[7] http://www2.itif.org/2013-cloud-computing-costs.pdf
[8] This list of actions is neither exhaustive nor presented in order of priority.

services to be made available to scientific communities, pre-commercial procurement (PCP) and Public Procurement for Innovation (PPI).

- Continuing engagement of Helix Nebula initiative with organizations in charge of leading work in EU in areas such as security, certification, standardization, data protection and management for implementation of effective interoperable solutions.

- Pursuing cooperation with other geo-hazard related EC projects (DORIS, LAMPRE) and Iceland Supersite to seize opportunities arising from downstream Copernicus services.

- Moving towards advancement of the Helix Nebula platform through adoption of user friendly common interface allowing automated discovery of services in a transparent way and allowing comparison of contractual terms and conditions, SLAs, pricing, KPIs, QoS.

- Establishing agreed and interoperable service management policies, process and procedures for cloud offering in a hybrid federated environment as well as the correct mechanism to assess their effectiveness and implement corrective measures. The adoption of the FitSM standard developed by FedSM project is a first step in this direction and other service management schemes are also considered for implementation.

- Continuing close collaboration with publicly funded e-Infrastructures such as GÉANT, EGI, EUDAT and PRACE, in a more integrated scenario including an e-Infrastructure Commons Marketplace, to support excellent science in Europe and the achievement of an integrated European Research Area.

- Closely following up and continuing involvement in activities of the European Cloud Partnership, including participation to thematic working groups and collaboration with Cloud for Europe and similar initiatives, including Connecting Europe Facility (CEF) Digital.

- Planning a long term vision for the Helix Nebula initiative and for the development of the Helix Nebula cloud infrastructure also considering opportunities offered by Horizon 2020.

- Focusing on better targeting communication and dissemination activities, in different EU languages, to involve more players and potential users to the marketplace. Achievement of a relevant critical mass would also allow advocating for the establishment of a possible new cPPP for Cloud[9].

- Implementing a strategy for the next three years to achieve large scaling capacity by:

  - ✓ Focusing effort on attracting more big players from European cloud players, including developers of PaaS and SaaS solutions,
  - ✓ Providing a collection of public services building blocks which can be offered in an open and interoperable way, and reused & combined by public administrations and third parties, allowing evolution toward a "cloud of public services" concept.

[9] http://europa.eu/rapid/press-release_MEMO-13-1159_en.htm

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301