Helix Nebula – The Science Cloud

Title: D6.1 Interoperability Requirements Report
Editor: Sergio Andreozzi, Carmela Asero (EGI.eu)
Work Package: WP6 – Interoperability with e-Infrastructures
Submission Date: 21.05.2013
Distribution: Public
Nature: Report

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

Log Table

Issue	Date	Description	Author/Partner
V0.1	19 April 2013	Initial draft	Sergio Andreozzi (EGI.eu) Carmela Asero (EGI.eu)
V0.2	22 April 2013	Provided comments	Rachida Amsaghrou (CERN)
V0.3	26 April 2013	Added comments from dedicated Helix Nebula supplier call and from Robert Jones (CERN)	Sergio Andreozzi (EGI.eu)
V0.4	29 April 2013	Added comments from Michel van Adrichem	Carmela Asero (EGI.eu)
V0.5	6 May 2013	Added comments from Mick Symonds and Jurry de la Mar	Carmela Asero (EGI.eu)
V0.6	16 May 2013	Further refinements based on the provided comments	Sergio Andreozzi (EGI.eu)
V0.7	21 May 2013	Small revision after management team review	Sergio Andreozzi (EGI.eu)

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

Executive Summary

This document reports on the activities carried out in Helix Nebula, during the first year of the project, to investigate, identify and analyse interoperability and integration concerns between publicly funded e-Infrastructures and commercial cloud providers. The report has been structured around interoperability levels as described in the European Interoperability Framework and it also includes a specific analysis concerning different business cases and operational scenarios of interoperating service architectures.

The results of two Helix Nebula workshops addressing interoperability aspects, held in September 2012 and January 2013 plus the outcome of an open session at EGI Community Forum in April 2013 were fed into this document, as well as the work of ad hoc interoperability task forces addressing the different interoperability levels: Political and Legal, Organizational, Technical and Semantic.

Thanks to valuable input from several Helix Nebula stakeholders, the document provides an agreed set of recommendations, it proposes actions for implementing them, and describes motivations and a plan for an interoperability test case.

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

Table of Contents

1 Introduction

Cloud computing offers scientific communities new opportunities to flexibly build their virtual research environments, scale up their activities and optimise the utilisation of resources. The integration and interoperation of resources from publicly funded e-Infrastructures and commercial providers are key elements to improve freedom of choice among users and to enable a virtual single digital market for cloud services that is affordable and can meet almost unlimited demand.

Helix Nebula has brought commercial cloud providers and representatives from publicly funded ICT infrastructures for research to work together and conceive a comprehensive European cloud model and architecture with a potential for exploitation in the context of wider EU public services.

This stands for an unprecedented challenge, in finding the right balance of requirements and constraints, barriers and solutions among participants to the project. This collaboration has particularly been evident in the conception of technical and service architectures as well as in business models for the services to be provided by Helix Nebula.

Identifying possible scenarios for service delivery and the roles of different categories of players (e.g. integrators, brokers, marketplace aggregator) will help Helix Nebula in seizing different opportunities for growth among a range of stakeholders and thus strengthen its future sustainability.

This document analyses the work done under Work Package 6 in the context of the EC-funded Helix Nebula project. It provides a number of high-level recommendations for interoperability and integration of publicly funded infrastructures with commercial providers as well as it opens areas for discussion.

During the second year of the project, it is expected that with the development of a more detailed service and technical architecture, with the maturing of the Blue Box implementations and with the continuous dialogue with publicly funded infrastructures, concrete actions could be agreed on, and implemented.

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

2 Method of work

In order to tackle the interoperation and integration issues between commercial cloud providers and publicly funded e-Infrastructures, the European Interoperability Framework (EIF) for European public services was selected to structure the discussion at different levels of concern [R1]. This framework envisions five levels of interoperability: 1) political context, 2) legal interoperability 3) organisational interoperability, 4) semantic interoperability and 5) technical interoperability (see Figure 1).

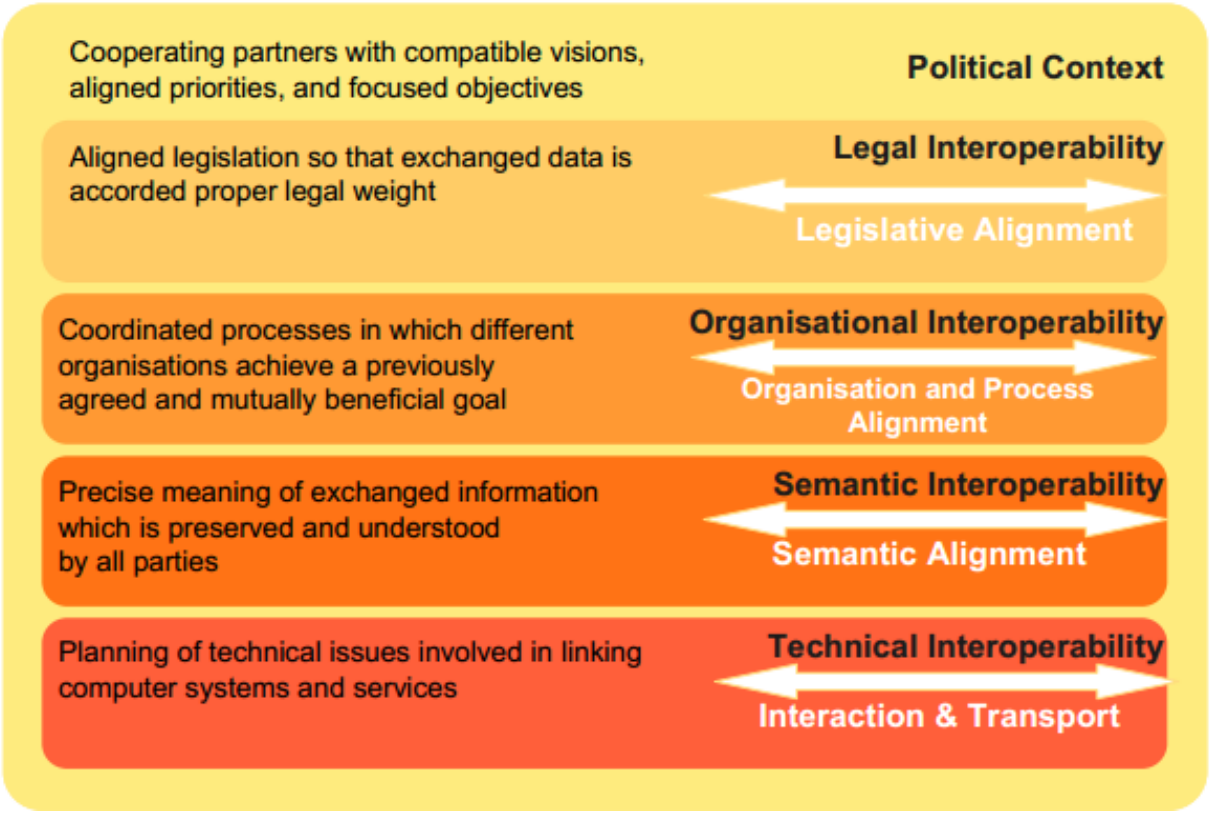

Figure 1 Interoperability Levels

These interoperability levels need to be analysed across the various organisations involved in this context. Integration and interoperation of emerging and mature e-Infrastructures (i.e. commercial clouds, GEANT, EGI, PRACE and EUDAT) required an analysis of the technology and policy interfaces around each infrastructure; in order to identify how an effective integration can be achieved. We can observe that the various infrastructure contexts cut across the various interoperability concerns as depicted in Figure 2.

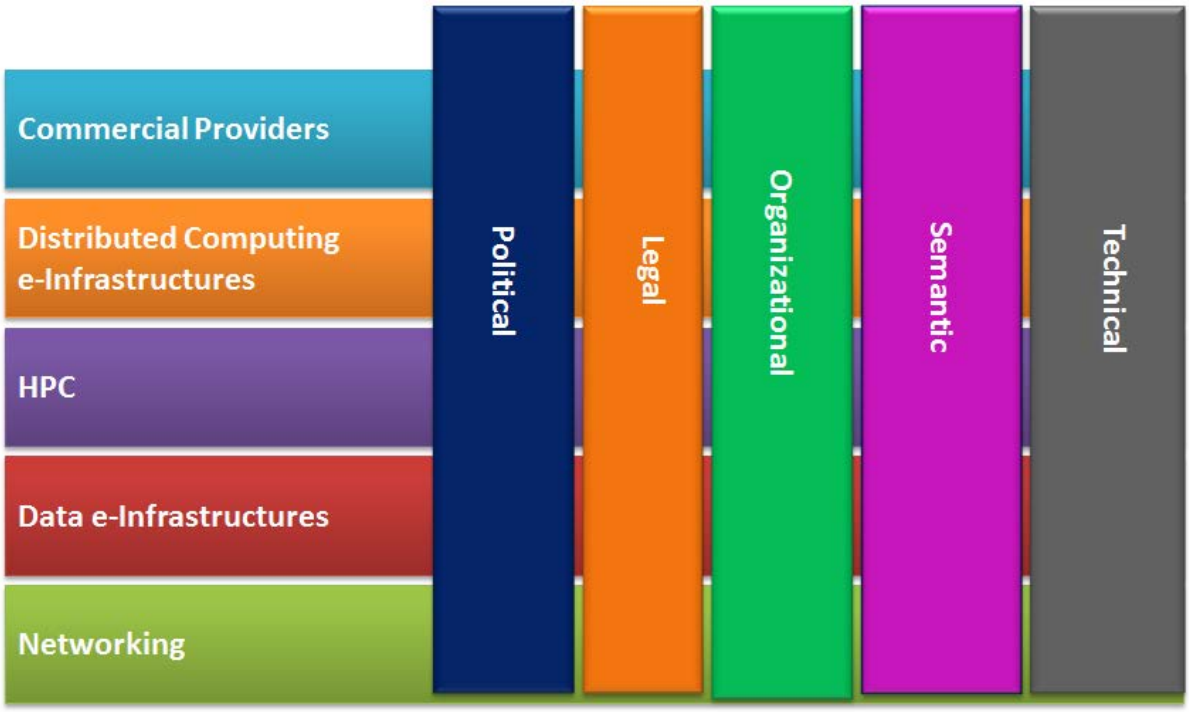

Figure 2 Interoperability levels versus e-Infrastructures layers

Figure 3 shows the relationship between Work Package 6 and the other Work Packages of the Helix Nebula project. So far, WP6 received input from WP3 on identified requirements and from WP4 regarding the validation and assembly of the resources needed to meet those requirements into possible and feasible service delivery.

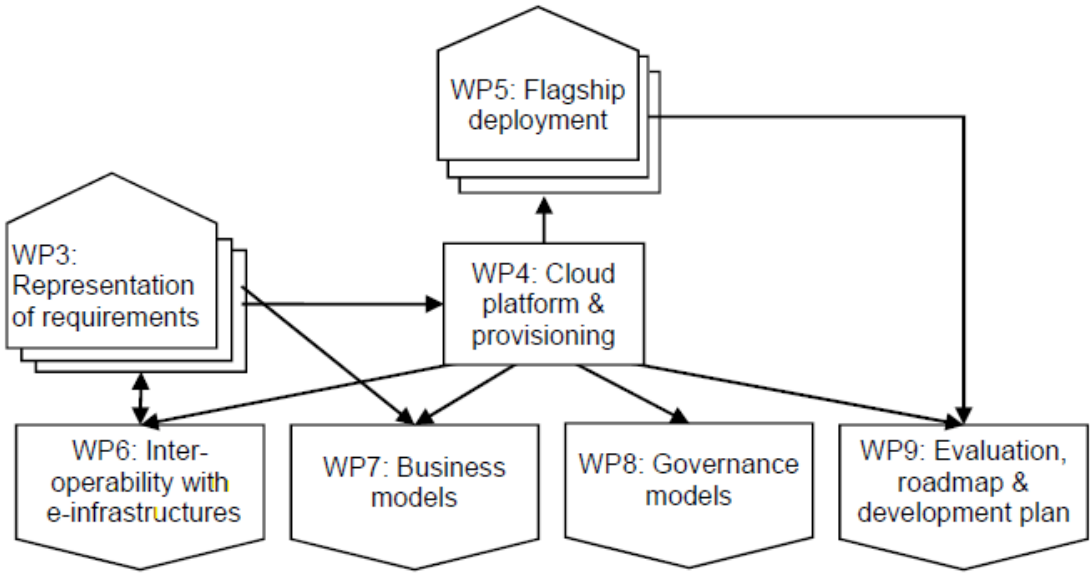

Figure 3 Helix Nebula project structure

Led by WP6, key technology and policy representatives (e.g., developers, operational, legal, management) of each e-Infrastructure gathered at two open workshops [R2, R3] during the first year of the project. These meetings offered the opportunity for information exchange between each group, and areas needing integration/interoperation to be identified. Following the second workshop, WP6 also decided to set up dedicated task forces to focus on discussions with experts [R4]. Three task forces have been set up to cover respectively: 1) political/legal level, 2) organisational/semantic level, 3) technical level.

3 Business Case

Interoperability aspects that need to be investigated depend on the possible business case scenarios for the integration of commercial services with publicly funded e-Infrastructures. The commercial and public sector suppliers within the Helix Nebula Initiative have started working on the conceptual technical architecture which has as its core the Service Enabling Framework, the so-called "Blue Box", a complex component providing API services and a Web Portal that will enable users to interact in a central and transparent manner with and federate the infrastructures of multiple Cloud Providers.

So far, the discussion has focussed on the interoperability between commercial suppliers on the one hand, and EGI & DANTE e-Infrastructures on the other hand. Interaction with

PRACE[1] and EUDAT[2] will be considered during the second year of the project, when this work will have a more solid basis. It should be noted that EGI serves a part of the "big science" research communities. Other "big science" communities and the "long tail" researchers rely on their own e-Infrastructures or are already being served by the private sector. EGI is based on a contributor model where users receive national funding to buy resources and procure them through public-funded resource centers. Through the EGI technology and integration services, they are able to form virtual organizations (VO)[3] to share the acquired resources across organizational boundaries and enable cross-border collaborations.

From the publicly funded resource center[4] viewpoint, identified use cases for interoperability are:

1. A publicly funded resource centre has valuable data that attracts a considerable amount of workload, but does not have enough capacity to serve the needs of their users
2. A publicly funded resource centre wants to offer different types of resources not available in house (e.g., GPU's)
3. A publicly funded resource centre wants to offer different SLA's

From the researcher's viewpoint, the identified use cases for interoperability are:

[1] PRACE: Partnership for Advanced Computing in Europe (http://www.prace-ri.eu/) provides access to world class computing and data management resources and services through a peer review process

[2] EUDAT: European Data Infrastructure (http://www.eudat.eu/) aims at ensuring adequate services for research data management, access and preservation

[3] Virtual Organisation: A group of people (e.g. scientists, researchers) with common interests and requirements, who need to work collaboratively and/or share resources (e.g. data, software, expertise, CPU, storage space) regardless of geographical location. They join a VO in order to access resources to meet these needs, after agreeing to a set of rules and Policies that govern their access and security rights (to users, resources and data) (Source http://www.egi.eu/about/glossary/glossary_V.html)

[4] In EGI, a Resource Centre, is the smallest resource administration domain in an e-Infrastructure. It can be either localised or geographically distributed. It provides a minimum set of local or remote IT Services compliant to well-defined IT Capabilities necessary to make resources accessible to Users. Access is granted by exposing common interfaces to Users (source: http://www.egi.eu/about/glossary/glossary_R.html)

1. A researcher who wants to consume data-sets available as open access in a publicly-funded infrastructure has no computing capacity available but has budget allocated to buy cloud services
2. A researcher who wants to combine usage of different types of resources from both publicly-funded and commercial infrastructures

From the commercial provider viewpoint, the identified use cases for integration are:

- To provide its available resources to researchers without the researchers having to use another way of getting access
- To exploit a market segment with limited costs of sales, that otherwise may not be served
- To benefit from economies-of-scale that some big science use cases might contribute

Figure 4 presents the ex-ante scenario considering EGI, DANTE+NRENs and commercial providers before Helix Nebula. In this scenario, the publicly funded resource centres are connected together through the publicly funded network, while commercial providers are connected to the Internet and/or the commercial network.

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

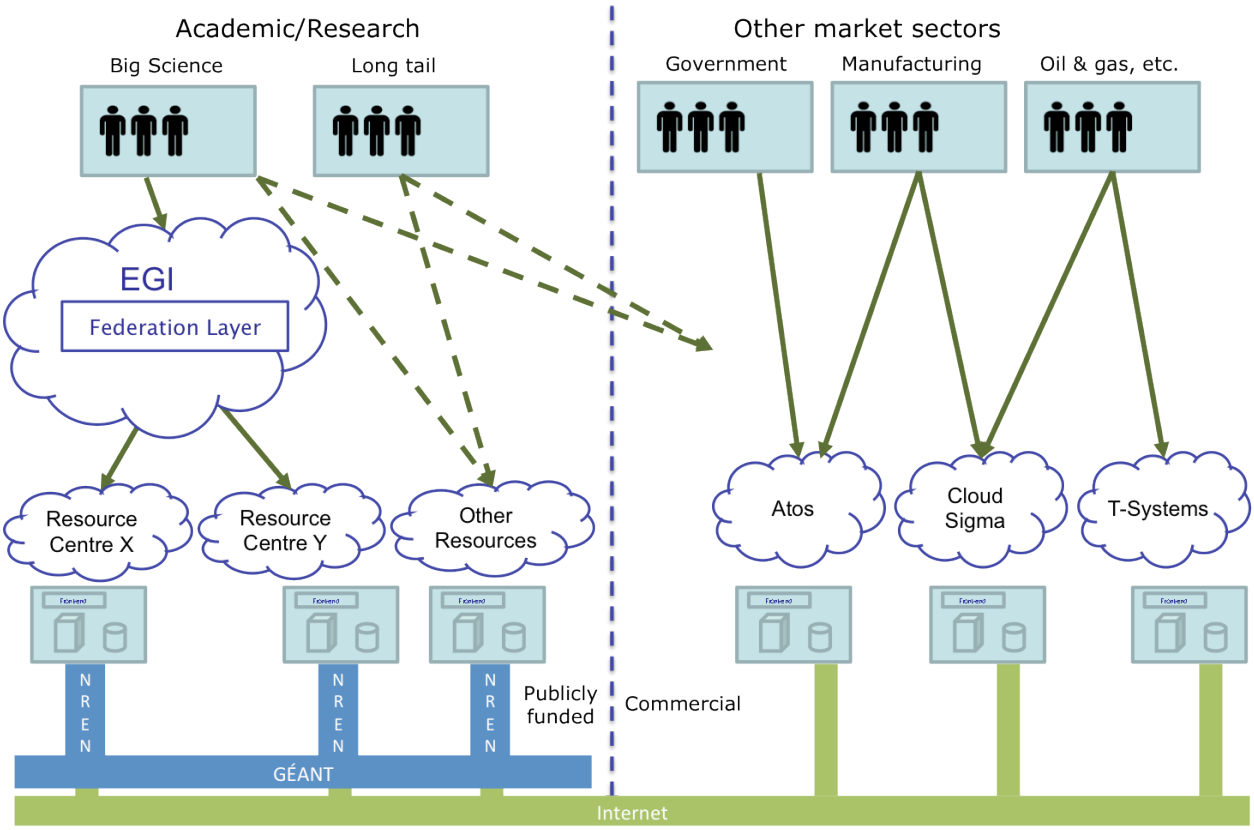

Figure 4 Ex-ante Scenario

In the following sub-sections, we present three possible scenarios describing the relationships between the EGI Federated Cloud, the research and education network (DANTE and the NRENs) and the Helix Nebula Blue Box. At this stage, the analysis does not consider the integration of research communities that are not part of EGI. Such requirements will be addressed during the second year in collaboration with WP7.

3.1 Scenario 1: Federated Infrastructure Bursting

The first scenario is the 'Federated Infrastructure Bursting' (see Figure 5). In this scenario, the EGI Federated Cloud layer is integrated with the Helix Nebula Blue Box (both at the broker level) in a kind of broker peering mode. An academic user can have the requests served by the publicly funded infrastructures or by a commercial provider. This scenario can serve the needs of extra capacity or of increasing freedom of choice to academic users that may want to combine resources available in the different domains. It also raises the

issue of how to combine two different procurement models: 1) for EGI, the current approach is that research communities receive grants to acquire resources to be installed at publicly-funded resource centres (CAPEX); 2) for commercial providers, research communities should have available funding to buy the cloud services (OPEX)..

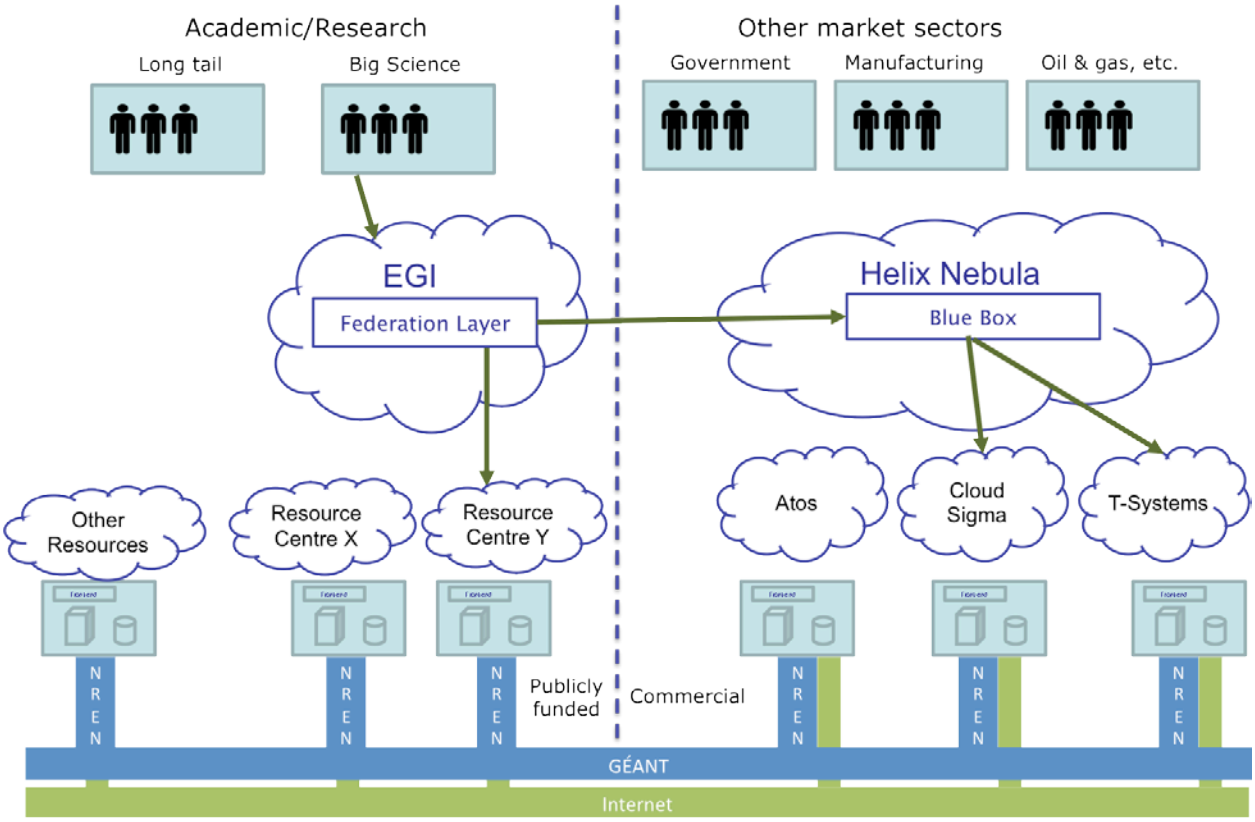

Figure 5 Integration Scenario #1: Federated Provider Bursting

3.2 Scenario 2: Integration with External Broker

In the "Integration with External Broker" scenario, the academic user directly accesses the Blue Box services because they may offer a richer set of capabilities or specialized services not yet available from the EGI Federated Cloud (see Figure 6). In this scenario, an academic user can submit a request for cloud services directly to Helix Nebula, which will assign computing resources from both Commercial and publicly funded resource providers. This presumes that all relevant public resource centres are connected to the HN Blue Box(es).

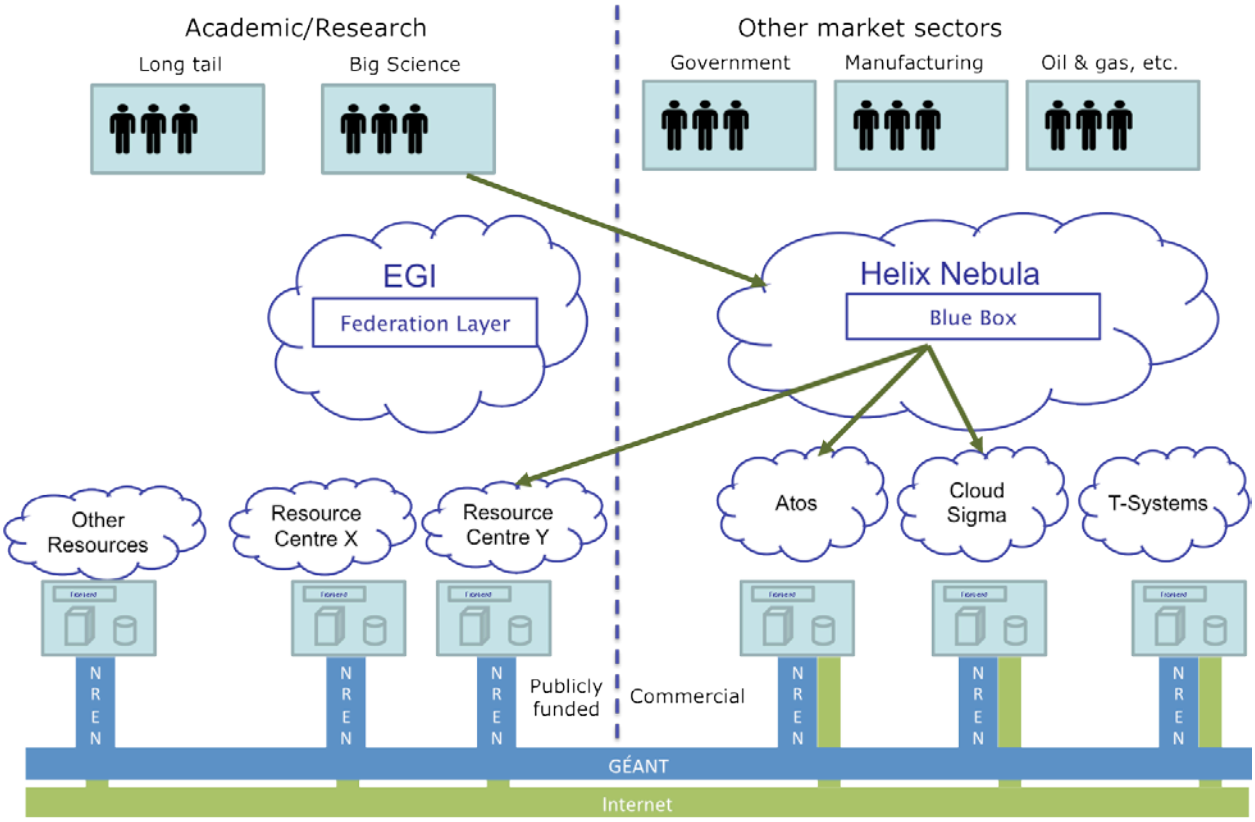

Figure 6 Integration Scenario #2: Integration of External Broker

3.3 Scenario 3: Single Provider Bursting

In the 'Single Provider Bursting' scenario (see Figure 7), the academic users request cloud services from the EGI Federated Cloud infrastructure. An individual resource centre may demand for extra services or computing capacity to commercial providers via the Helix Nebula blue box to expand its capacity.

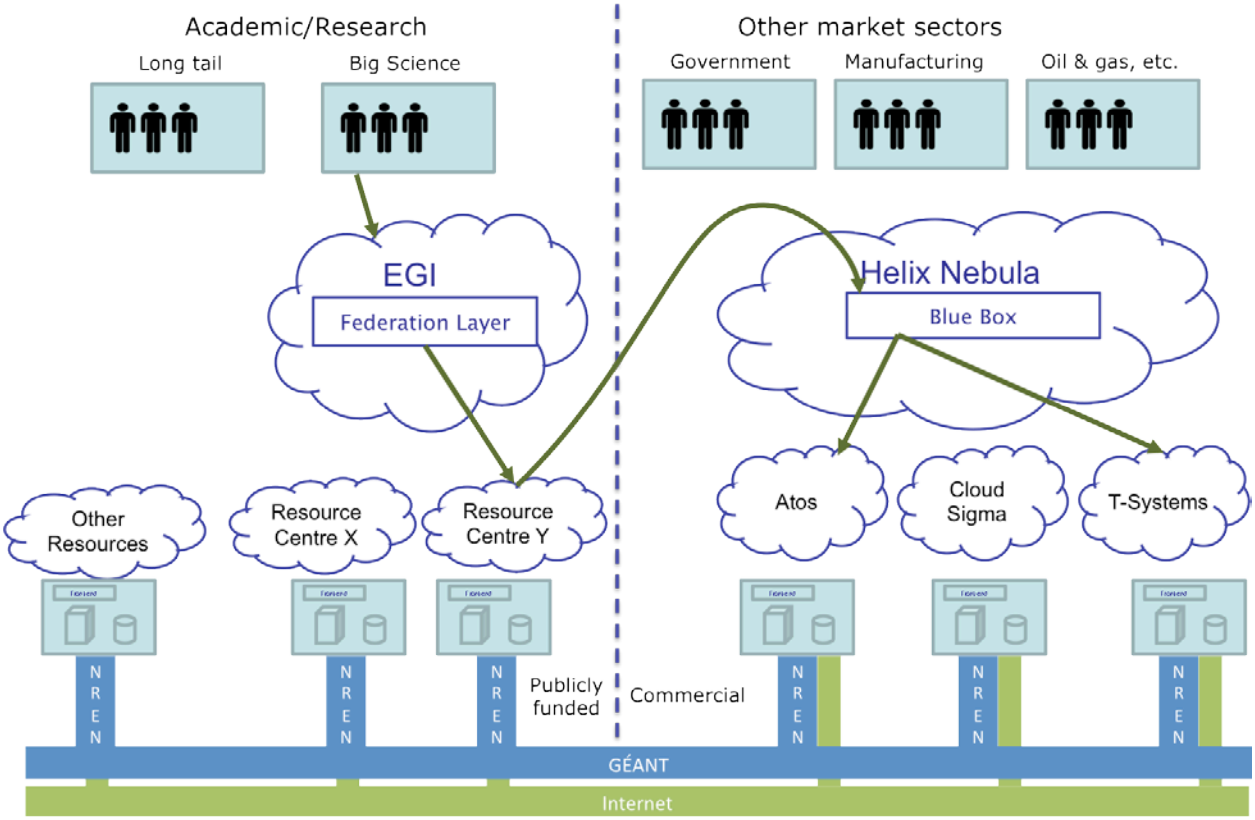

Figure 7 Integration Scenario #3: Single Provider Bursting

The discussion on the most viable scenario to be adopted is still on-going and several could be potentially supported at the same time. The analysis should be further evolved to consider value creation for all involved parties that enable an open ecosystem to develop. However there is no doubt that the Federated Cloud approach is technically viable: from the EGI viewpoint, the Helix Nebula Blue Box would be yet another provider. And as the EGI Federated Cloud is based on open standards, the Helix Nebula Blue Box could easily be integrated with EGI resource centres when it becomes available.

4 Interoperability Requirements

This section describes the interoperability and integration requirements identified so far through the two workshops and the three task forces.

4.1 Political Context

The set up and deployment of pan-European services normally requires the definition and sharing of a common strategy across various types of institutions, e.g., funding agencies, service providers, user communities. Even when new services do not depend on new legislation but are simply developed to improve existing services, an open and effective interaction with the political level for support and sponsorship is often needed to overcome the barriers and cost of integration. In the context of Helix Nebula, it is fundamental that all stakeholders involved hold shared visions, set agreed objectives and align priorities to achieve an effective cooperation.

The Helix Nebula partners have seized opportunities to engage in large cloud events and workshops also involving key policy makers. It is necessary to continue this dialogue at the European and national political level to ensure that national agenda for cloud adoption in research and science have a harmonized approach that can create a favourable and sustainable environment for Helix Nebula and other similar initiatives to provide services at pan-European and international scale.

When we look at the political context around cloud computing services in Europe, the recent cloud strategy and the reformed data protection package to be implemented by the European Commission define the main political framework in which several national regulations are in place, particularly in terms of privacy rules and contractual legislation. In this document, we will focus mainly on regulation at the EU level.

The European Commission Communication "Unleashing the Potential of Cloud Computing in Europe" [R4] identifies three areas hindering uptake of cloud service in Europe and in which policy action is needed:

- **Fragmentation of the digital single market**: the different national legal frameworks in Europe are elements of uncertainties over applicable law, digital content and data location. These aspects have been identified amongst the concerns of potential cloud computing uptake. Cloud service management architectures are

complex and the use of cloud services has implications falling under multiple jurisdictions.

- **Contract issues**: in particular in the area of data access, portability, control and ownership. Some concerns are related to liability for service failures such as downtime or loss of data and possible forms of compensation, user rights in relation to system upgrades decided unilaterally by the provider, ownership of data created in cloud applications or dispute resolution.

- **A jungle of standards**: a myriad of standards is proliferating in cloud computing, however there is lack of certainty as to which standards provide adequate levels of interoperability system, application and data formats to permit portability.

With regards to the creation of a single digital market for cloud services targeted at research communities encompassing both publicly funded infrastructures and commercial provider, it is essential to ensure a level playing field and synergies across the two domains. Further analysis to eradicate the real fear of unfair internal competition between the publicly funded research infrastructures and the commercial actors goes through the identification of appropriate business models and policies.

Some ideas to be considered are: 1) publicly funded resource centres connected to the BlueBox can offer free services to research groups that own resources while they should provide a pricing model following a full cost accounting (FCA) analysis, the provision of services should be limited to non-profit research activities; 2) research groups who have interesting scientific case, but do not have direct funding to buy cloud services could access publicly funded resource centres and pay with "scientific results" to meet the national/EU policy objectives on excellent science; 3) research groups who have a commercially exploitable application must go to the commercial cloud providers; a prototyping phase could be supported by the publicly funded infrastructures if limited budget is available during the incubation phase.; 4) research groups who have funding for cloud services will choose among the publicly funded infrastructures or commercial providers according to the service functionalities, service levels, and pricing models.

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

Recommendation 1: Ensure a level playing field for the various players in the service delivery field and eradicate any potential fear of unfair internal competition among the publicly funded and commercial actors.

Helix Nebula has a clear commitment to adopt open standards to avoid vendor lock-in, to build an open ecosystem and ensure healthy competitiveness. Critical areas in the use of standards are related to managing virtual machines, moving data and Single Sign On (SSO). The European Commission has requested the European Telecommunications Standards Institute (ETSI) to coordinate with stakeholders and identify a detailed map of required standards in areas such as security, interoperability, data portability and reversibility. It is essential to ensure that Helix Nebula will build a common position on priorities to be reported in this context.

Recommendation 2: Both publicly funded and commercial cloud providers should agree on a core set of open standards endorsed by the user communities and liaise with ETSI to ensure that their view on standards selection and road-mapping is considered.

Another important task in this area is under the responsibility of the European Network and Information Security Agency (ENISA). ENISA will support the Commission in the development of an EU-wide voluntary certification schemes in the area of cloud computing (including data protection) and establish a list of such schemes by 2014.

A relevant recommendation in ENISA's Cloud Computing Risk Assessment [R5] report of 2009 is the Information Assurance Framework [R6], a set of assurance criteria designed to assess the risk of adopting cloud services, compare different Cloud Provider offers, obtain assurance from the selected cloud providers, and reduce the assurance burden on cloud providers.

ENISA is also involved in key areas such as procurement activities and Critical Information Infrastructures Protection (CIIP). In the recent document "Procure Secure: A guide to monitoring of security service levels in cloud contracts" [R7] ENISA gave guidance to cloud services customers on continuous monitoring of security service levels and governance of

outsourced cloud services. This is achieved through the reporting and alerting of key measurable parameters, as well as a clear understanding of how to manage the customer's own responsibilities for security. The document underlines how both the cloud service provider and the customer should be able to respond to changes in the threat environment on a continuous basis by monitoring the on-going implementation of security controls and the fulfilment of key security objectives. This is also recommended as a priority in the US government's 2010 report on the implementation of the federal information security management act (FISMA) [R8].

ENISA also issued a recent report "Critical Cloud Computing - A CIIP perspective on cloud computing services" [R9] identifying a number of scenarios and relevant threats related to the uptake of cloud computing, large cyber-attacks and disruptions of cloud computing services.

The Helix Nebula partners are committed to tackling security challenges highlighted in ENISA's reports and in aligning with its recommended security requirements. To date, a pragmatic approach has been taken by defining each partner's security policies and procedures as the starting point, when moving towards a federated environment. Since all partners are European and operate for many years under relevant EU and Member State legislation this already builds a framework. However, it is essential that stronger commonalities will be built. And within EGI, the Security Policy Group (SPG) [R16] have been developing security policies for a federated environment for many years.

Recommendation 3: Set up a common task force between commercial cloud providers and publicly funded infrastructures to define common security policies aligned with guidelines issued by ENISA.

4.2 Legal interoperability

With regards to the contracts, many activities are on-going so as to identify the important aspects that should be covered and to define templates that could simplify the negotiation process while providing transparency. It is observed that large providers offering a very strict form of terms of use, usually unidirectional and not negotiable, currently dominate

the cloud market. The terms of use document usually stands for the primary legal document between the cloud provider and the cloud user/customer.

For instance, EuroCloud Austria has published the document "Cloud Contracts", a catalogue of recommended contractual components in General Terms and Conditions of Business (AGB) and Service Level Agreements (SLA) for Cloud Service Providers. One aspect to be monitored is also the intention to extend the Common European Sales Law (coherent set of rules adapted to the distance supply and in particular supply online of digital content and related services) to cloud services.

The harmonization of contractual aspects for cloud computing services is also one of the objectives outlined in the Strategy for Cloud Computing in Europe, yet this process is still in a preliminary phase and different national legislations are still in place.

Important is that users will be able to build services from modular building blocks in a transparent and consistent way. This will require terms and conditions to provide a certain legal interoperability, whereby individual service levels or KPIs might vary. E.g. a user wants to load balance a service built on a commercial and public building block, and maintain a certain overall availability needs to know the availabilities of the individual building blocks and the demarcation points in a consistent way. The same consideration applies for the security aspects, where the overall security depends on the compliance level of the elements. The definition of common KPIs can also improve the understanding and comparison of different cloud offerings.

Recommendation 4: Publicly funded and commercial cloud providers should agree on a set of important elements that customers should consider when agreeing on cloud contracts, both for terms and conditions and SLAs.

In 2012, the European Commission proposed a major reform of the EU legal framework on the protection of personal data. The European Parliament is expected to define its position on the EU's new data protection regime by mid-2013. Although this may not affect the work of Helix Nebula with reference to the requirements applicable to current flagship use cases, such policy developments should be carefully monitored to keep pace with possible

changes that may impact a potential uptake of the Helix Nebula cloud environment to deliver public services.

An interesting initiative on data privacy is being carried on by Cloud Security alliance, who has recently published a "Privacy Level Agreement Outline for the Sale of Cloud Services in the European Union" [R17], offering an interesting input to implement proper privacy and data protection practices in cloud services.

Recommendation 5: Ensure that both publicly funded and commercial cloud providers follow the code of practice developed by Helix Nebula to comply with the data regulations.

Another delicate legal aspect is the applicability of legislation from outside the EU, such as the US Patriot Act and other pieces of legislation impacting cloud service provisioning. The recent study from the University of Amsterdam "Cloud Computing in Higher Education and Research Institutions and the USA Patriot Act" made clear what follows: "It is a persistent misconception that U.S. jurisdiction does not apply if the data are not stored on U.S. territory. The key criterion in this respect is whether the cloud provider conducts systematic business in the United States, for example because it is based there or is a subsidiary of a U.S.-based company that controls the data in question" [R13].

Apart from data sensitivity, there is also a generally less known continuity of service issue. In the USA, a service can be stopped based on the Patriot Act. The Helix Nebula providers are aware of such aspects and of the need to ensure that data are stored in the EU. However, it is important to consider two aspects: where data is stored and where data is processed; as long as a provider processes and stores data in an EU Member States, the Helix Nebula services can assure full compliance to EU law, but further investigation is needed in the case of a European provider using a datacentre owned and located on the US territory.

In view of possible future exploitation of the Helix Nebula services beyond science and toward more general public services, the collaboration started by Helix Nebula with G-Cloud could be beneficial for both initiatives to study legal constraints on storage and processing of data in the cloud.

Recommendation 6: Publicly funded and commercial cloud providers should develop a common understanding on the impact of an extra-EU legislation on the provision of their services to consumers inside and outside of their local legal jurisdiction.

The WP6 analysis of legal interoperability also addresses the management of IPR policy within Helix Nebula, which was discussed but this topic still requires further analysis following the development of business model scenarios and growth of the Helix Nebula partnership.

At this stage, there is no IPR policy in Helix Nebula covering the activities that potential users/customers can perform on the provided services. EGI has a policy stating that resource providers do not retain any intellectual property rights on the software, information and data provided to the services by its users.

Recommendation 7: Publicly funded and commercial cloud providers should agree on a common policy that protects users' IPR on the provided software, information and data.

4.3 Organisational interoperability

One relevant aspect to be considered in this area is IT service management in federated environments. Publicly funded infrastructures such as EGI are organised in a partnership where multiple suppliers provide federated services with the support of a shared service centre (EGI.eu). Helix Nebula also is a partnership with one or more shared service centres to emerge (the ones managing Blue Box instances). It is important that services consumed within the partners' infrastructures and by the users are managed according to agreed processes and criteria. In order to assess them, their maturity should be defined.

The challenge of shift to a service management perspective is particularly relevant for federated cloud environments which are cross organizational, multi-disciplinary, and cutting across national boundaries with *considerably less* control over interaction between participants than a normal service contract or commercial relationships.

The EC-funded FedSM project has defined FitSM [R18], a standard for lightweight service management in federated IT infrastructures based on ISO/EIC 20000 standard. The approach is meant to be pragmatic and to identify IT processes, their requirements and an assessment framework that can be used by resource centres participating in a federated service provision such as in EGI. This approach could be re-used by the Helix Nebula initiative since most candidate commercial providers already comply or are certified to the ISO standards and manage the services according to ITIL.

For Helix Nebula, the scope of the minimal set of requirements for service management should be restricted to the areas where resource centres interact with each other. The interaction between resource centres improves when the interaction processes align with each other through compliance to a common set of requirements. However each resource centre should be free in determining the set-up of service management processes that do not play a part in the interaction between resource centres. This limitation in scope gives resource centres the ability to implement strictly internal service management processes in a way that fits its services description and service levels.

Recommendation 8: Publicly funded and commercial cloud providers should agree on a minimal set of requirements for IT service management and a related maturity assessment framework that should be adopted by all members of Helix Nebula to evaluate the alignment of their service management practice.

Entering into a partnership to provide multi-supplier services with the potential roles of mediators that decouple the interaction between customers and resource centres require that a common ground should be set in terms of services that are needed to be provided.

Recommendation 9: Publicly funded and commercial cloud providers should agree on the common set of service management structures that are needed to participate in the federation.

In order to provide the needed quality of service for academic users to run their services in the commercial data centres, a proper connectivity is needed to reach the data stored in the publicly-funded infrastructures. The research community benefits from advanced

networking infrastructures provided by national organisations (NREN) and connected together through the European backbone (GÉANT) managed by DANTE. While it is technically possible to connect the commercial cloud providers to the publicly funded network, policy issues exist around the type of traffic that could be exchanged. Through the first workshop [R2] it was clarified that commercial cloud providers can connect to GÉANT through the local NRENs and that only research traffic is allowed regardless of if it originates from public or private organisations.

Following the workshop, DANTE joined the Helix Nebula partnership and set up a fruitful collaboration with commercial cloud providers of Helix Nebula to carry out the traffic of the three flagships use cases. The commercial providers involved have set up agreements with the local NRENs or directly with DANTE (see Figure 8).

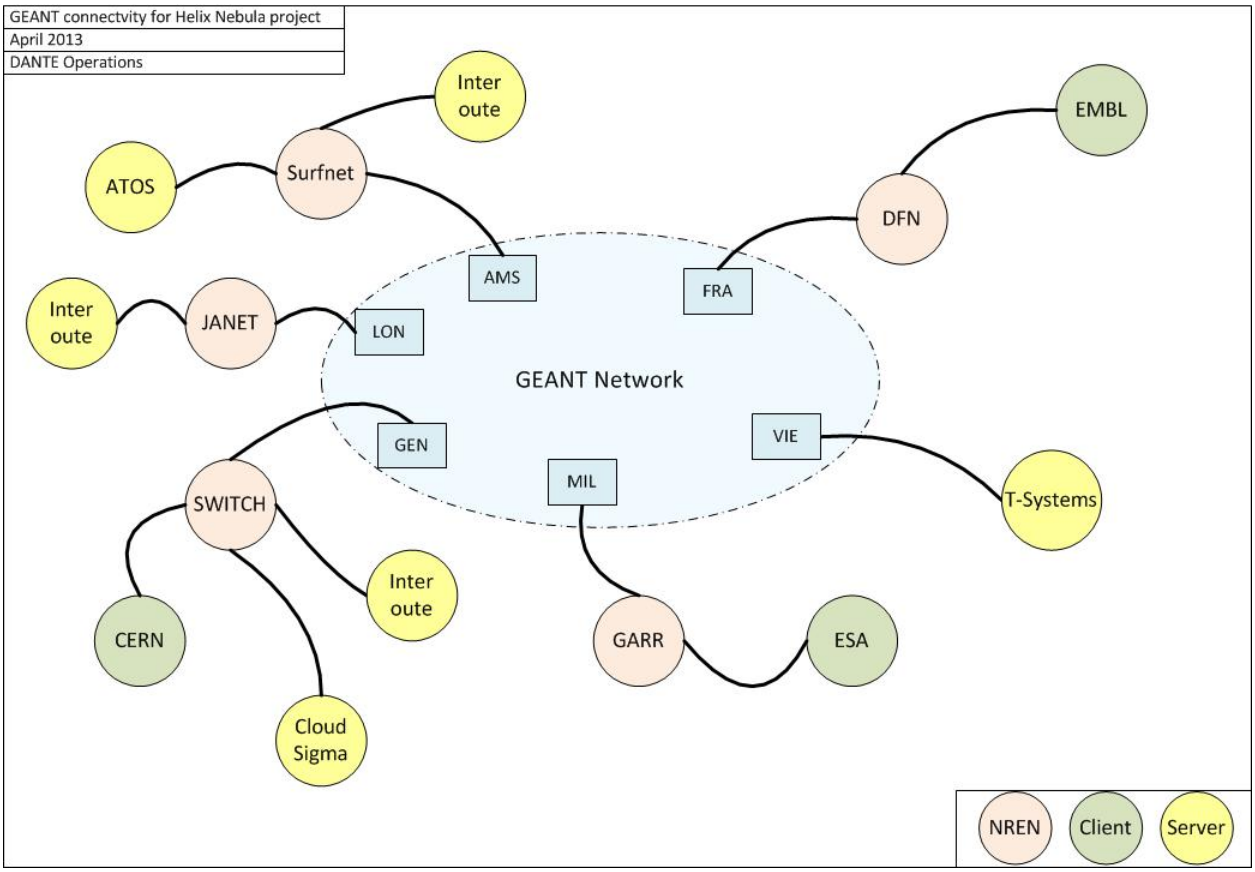

Figure 8 Inter-connection of Helix Nebula partners to GÉANT

The European Commission is encouraging NRENs to connect more entities than just Research centres, following recommendations from the GÉANT Expert Group report [R12]. It is not completely clear how costs should be covered for the production phase, but many stakeholders in Helix Nebula suggest that NRENs should bill the cloud provider for the connectivity service used.

Recommendation 10: Commercial cloud providers should be able to connect to GÉANT, on the basis of an agreed business model, to ensure that users can have the same connectivity level available in publicly funded e-Infrastructures to transfer data.

4.4 Semantic interoperability

One aspect that needs harmonisation on the semantic level is a compatible service catalogue structure to enable customers to easily compare and select services across different service providers that represent the best fit at any given time. For this reason, a common service catalogue template or scheme should be defined.

Recommendation 11: Both public funded and commercial cloud providers should agree on a compatible scheme to describe elements of a service catalogue, which eases service selection across different providers.

Similarly, customers would benefit from uniform accounting and billing data. However, there are many billing issues that need to be clarified before a compatible scheme can be defined and single integrated billing and payment is possible. An example of such an issue is that the taxation on the service depends on the country from which the service is delivered. If a service is delivered by two or more providers that reside in different countries, then having a single bill for the customer with settlement of costs between the providers leads to complicated taxation issues for all parties involved.

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

Recommendation 12: The issues concerning accounting, billing, payment and settlement should be analyzed to determine the possibilities and restrictions for both public funded and commercial cloud providers to agree on compatible accounting and billing parameters, cross settlements between providers and single integrated billing towards the customer.

4.5 Technical interoperability

The core of Helix Nebula is the Service Enabling Framework, the so-called "Blue Box", a complex component providing API services and a Web Portal that will enable users to interact in a central and transparent manner with all the Cloud Providers. Figure 9 presents a high-level view of the Blue Box role and functions.

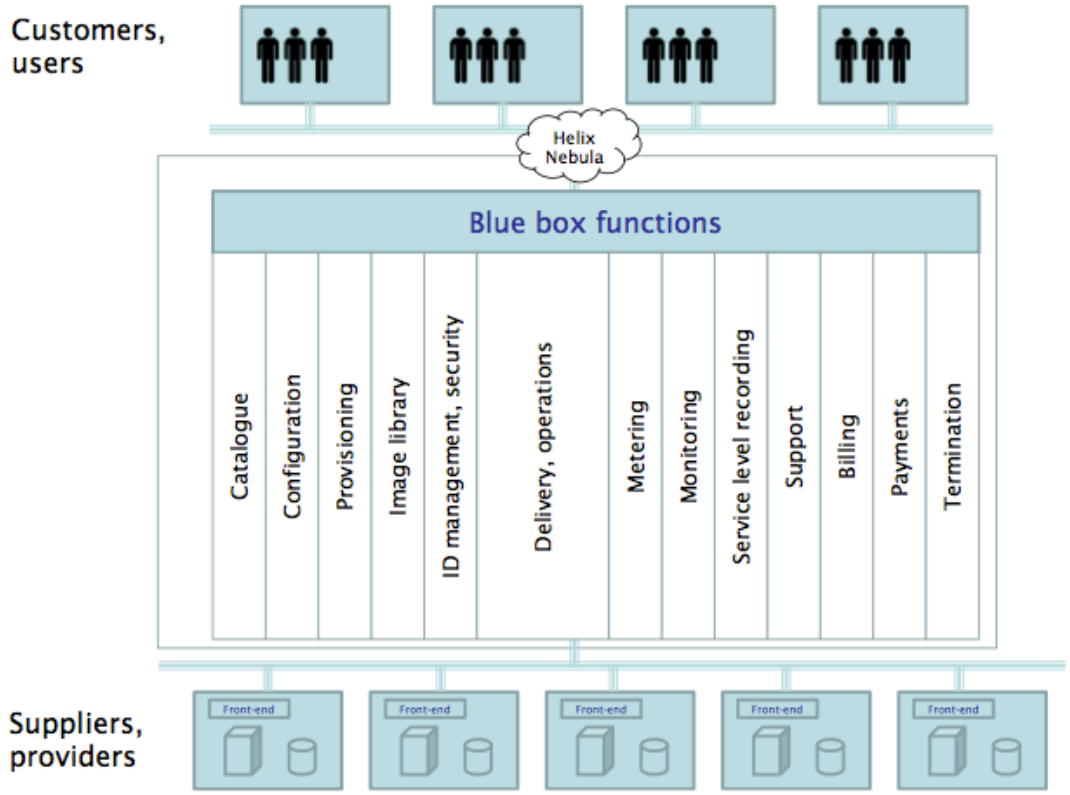

Figure 9 Federated architecture for Helix Nebula

The EGI Federated Cloud task force has developed a pilot test bed, moving into production, that integrates IaaS cloud services managed with different technologies (e.g., OpenNebula, Open Stack) into a federated service provision using open standards (see Figure 10).

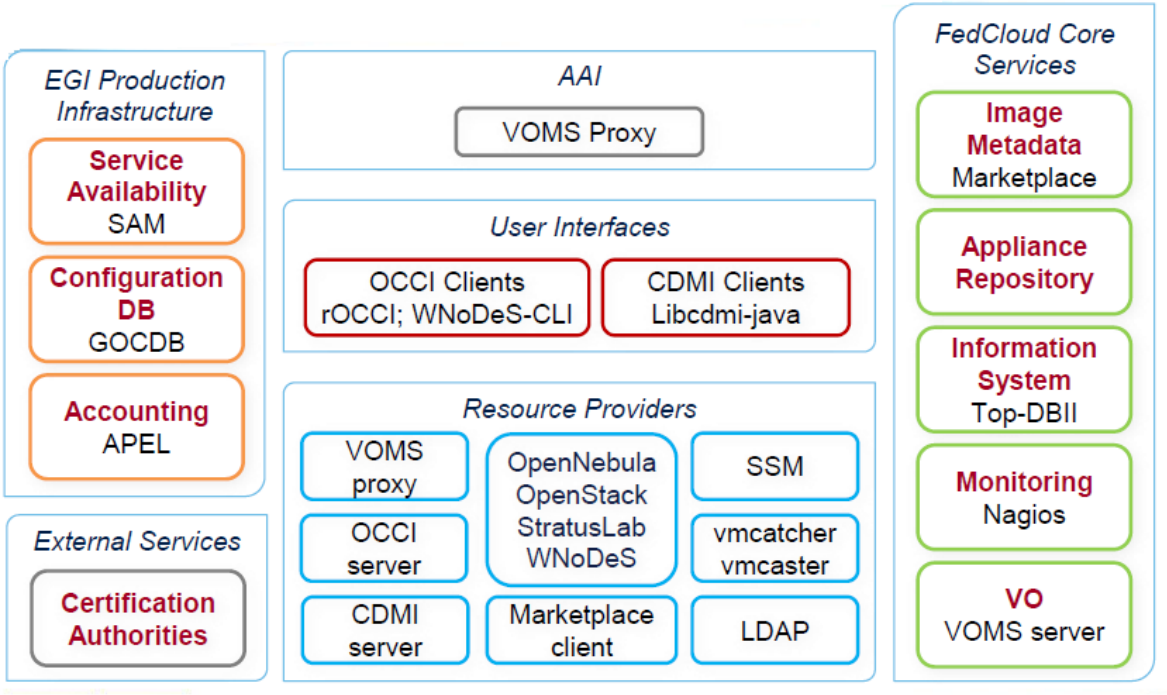

Figure 10 EGI Federation test bed composition – April 2013

An initial requirement to be considered is Single Sign-On (SSO) to enable users to authenticate across the various service providers using a single credential. EGI relies on X.509 certificate and proxy certificates for the authentication while the networking community has developed a federated authentication service based on the SAML protocol. Helix Nebula has defined SSO as a mandatory function of its Blue Box and the current version being piloted includes an SSO-function. However, SSO-functionality very much depends on the authentication mechanisms implemented by individual providers and the level of control they enable for 3rd party solutions. From the user perspective an SSO-function should not limit the control of the underlying infrastructures.

Recommendation 13: Users should be provided with a Single Sign-On (SSO) mechanism that enables them to authenticate on and have full control of the cloud services available in the federation. Reuse of existing sources of authentication should be considered.

Another aspect that emerged from the users related to the quality of experience and to avoid lock-in, is the possibility to easily port VM images where they deploy their applications or services across different service providers without the need to directly re-generate them. Three possible strategies have been identified:

1. The user specifies the characteristics of the VM starting from generic ones and an intermediate service can automatically generate different VM formats for different providers,
2. The user generates its own VM using a standard format (i.e., OVF) that is supported by the service providers,
3. The user generates its own VM and a translator service is able to map it into different formats supported by the providers.

Within Helix Nebula, some service providers favour the solution 1) to provide a good and stable service experience and limit the need for complex error-handling. The EGI Federated Cloud task force has currently adopted solution 2). A solution for interoperability has not yet been agreed on.

Recommendation 14: Users should be able to port virtual machine images across different providers without the need for complex translators.

Users should be exposed to a uniform interface for all services that need to be directly accessed using a programmatic approach. Three main strategies are possible: 1) common client-side API with adapters for different server-side API from resource centres, 2) common server-side API supported by a mediator/broker translating the requests to the different APIs supported by the resource centres, 3) common server-side API adopted by all resource centres.

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

With regards to the management of virtual machines, the EGI Federated Cloud has adopted solution 3) with the decision to support the OCCI open standard for all resource centres involved in the federation. The Helix Nebula Blue Box approach is a combination of solution 1) and 2) and provides the users 3rd party access to the API.

Recommendation 15: Users should be able to manage their virtual machines using a uniform interface across the different cloud providers (both publicly funded and commercial). Interfaces should be based on open standards and should not limit the available service functionality. Additional de-facto standard interfaces may be exposed to lower entry barriers and simplify the transition to open standards.

As for creating/retrieving/updating/deleting data objects, the EGI Federated Cloud has adopted solution 3) with the decision to support the CDMI open standard for all resource centres involved in the federation. The Helix Nebula Blue Box will support different APIs and it is expected to provide a uniform interface to the users.

Recommendation 16: Users should be able to create, retrieve, update and delete data elements using a uniform interface across the different cloud providers (both publicly funded and commercial). Interfaces should be based on open standards and should not limit the available service functionality. Additional de-facto standard interfaces may be exposed to lower entry barriers and simplify the transition to open standards.

Connectivity aspects are key to the definition and deployment of cloud services. In this area the requirements analysed include: internal/private and external/public connectivity, availability requirements, private and public networking needs, DNS requirements and an overall expected deployment topology. During 2012, the Helix Nebula providers successfully peered directly with respective NRENs, thus getting direct connectivity to the GÉANT network to carry the Helix Nebula traffic. However, the definition and configuration of connectivity in a large-scale federated and secure infrastructure showed to be complex and time consuming. Therefore, in collaboration with GÉANT, Helix Nebula started to investigate the adoption of Software Defined Network (SDN) functionalities towards Networking as a Service (NaaS).

Recommendation 17: EGI, commercial providers and DANTE should investigate how SDN can benefit a federated cloud computing infrastructure

5 Use Case for Interoperability Testing

This section describes a test case to be run during the second year of the project to demonstrate the interoperability among publicly funded infrastructures and commercial providers. Given the limited resources available both from the demand-side and from the supply-side, a pragmatic approach is proposed to maximise the results and to help refining the final roadmap to be proposed at the end of the project. The strategy is to re-use as much as possible the work that has been done for the deployment of the three flagships use cases.

In the current scenario:

- Two BlueBox implementations are being tested, one based on the open-source Slipstream solution and one based on the closed-source Enstratius
- In the deployment activity on the commercial suppliers, all three flagships use cases are using Enstratius (CERN, EMBL and ESA), while two flagships use cases are being deployed using SlipStream (CERN and ESA); the evaluation of the deployment will be performed using the criteria defined by Work Package 5 (see Helix Nebula Deliverable D5.1 "Evaluation of initial flagship deployment")
- EGI and DANTE are the publicly funded infrastructures directly involved in Helix Nebula (with EGI being the only one offering cloud services through the EGI Federated Cloud)

Being the BlueBox acting as mediator layer between the demand side and the supply side, the proposed idea is to re-use one or more flagship use cases to be deployed on the EGI Federated Cloud and to use the same evaluation criteria.

The deployment of the flagships use cases on EGI requires the availability from the demand side to invest further manpower to run the test case, the availability of EGI affiliated resource centres to provide resources and expertise and the availability to integrate the EGI Federated Cloud on the southbound interface of the BlueBox(es).

The EGI Federated Cloud has selected OCCI as a standard interface to be adopted by all resource providers participating in the federation. To date, none of the two BlueBox implementations have connectors for such an interface, and therefore work needs to be done in this direction. EGI has recently launched a mini-project to develop the OCCI connector for the open-source SlipStream and this should become available in the next 6 months [R19]. On the other side, no planned activity no available budget is available to develop a similar connector for the closed source Enstratius.

Concerning the availability of resources, it should be noted that the ATLAS and CMS experiments which are participating in the CERN flagship deployment own IT resources installed in EGI resource centres that have being used over the last years to support the large-scale computing needed by the affiliated researchers.

Given this scenario and the limited availability of the budget, the proposal for a test case is to deploy the CERN flagship use case on some EGI resource centres through the SlipStream as soon as the OCCI connector will become available. The deployment will be evaluated according to the same criteria used for the commercial providers.

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

6 Roadmap for Implementation

This section provides a summary of the identified recommendations, maps them to the defined scenarios, proposes a number of actions for implementing the recommendations and identifies owner of the recommendation as well as contributors. Concerning the proposed actions, we provide an estimation of the degree of difficulty for performing them according to the following classification:

- E for Easy: the action can be implemented with little effort and this can be achieved in the lifetime of the Helix Nebula project
- M for Medium: the action can be implemented with some dedicated effort and this is likely to be achieved in the lifetime of the Helix Nebula project
- D for Difficult: the action needs a high effort or there are external dependencies that cannot be controlled, therefore it is likely not being completed in the lifetime of the Helix Nebula project.

Since the implementation of the actions relies substantially on external collaborations or unfunded effort, the implementation plan proposed in Table 1 should be taken as tentative and serves to steer the activities towards common goals. The execution will be monitored and updated during the second year of the project and a report will be provided in Deliverable D6.2.

Table 1 Suggested Recommendations Implementation Plan

#	Description	Scenarios	Action and Easy of Doing (E/M/D)	Owner	Contributors
1	Ensure a level playing field for the various players in the service delivery field and eradicate any potential fear of unfair internal competition among the publicly funded and commercial actors	1,2,3	Agree on appropriate policies/business models that ensure level playing field when accessing resource through the BlueBox (M)	EGI	HN WP6, WP7, WP8, EGI FedCloud,
2	Both publicly funded and commercial cloud providers should agree on a core set of open	1,2,3	Agree on a common standards profile (M)	HN TechArch	HN TechArch/Ser

	standards endorsed by the user communities and liaise with ETSI to ensure that their view on standards selection and road-mapping is considered		Liaise with ETSI and promote the identified standards profile (E)	HN TechArch	vArch & EGI FedCloud
3	Set up a common task force between commercial cloud providers and publicly funded infrastructures to define common security policies aligned with guidelines issued by ENISA	1,2,3	Define terms of reference for a joint task force and kick off activities (E)	HN ServArch	HN TechArch/ServArch & EGI Security Policy Group & DANTE
4	Publicly funded and commercial cloud providers should agree on a set of important elements that customers should consider when agreeing on cloud contracts, both for terms and conditions and SLAs	1,2	Identify the important elements that customer should consider when evaluating cloud contracts and SLAs and define KPIs for comparing them (M)	HN ServArch	HN ServArch & EGI FedCloud
5	Ensure that both publicly funded and commercial cloud providers comply with the EC data protection regulation, once this gets approved	1,2,3	Engage with organisations performing similar study and liaise with them (M)	HN WP6	HN ServArch + WP6
			Obtain a statement of adoption of the code of practice by each provider (M)	HN ServArch	
6	Publicly funded and commercial cloud providers should develop a common understanding on the impact of extra-EU legislation on the provision of their services to consumers outside of their local legal jurisdiction	1,2,3	Engage with organisations performing similar study and liaise with them (M)	HN WP6	HN ServArch + WP6
			Update the code of practice to match relevant extra-EU legislations (M)	HN ServArch	
7	Publicly funded and commercial cloud providers should agree on a common policy that protects users IPR on the provided	1,2,3	This could be addressed by the joint task force related to recommendation 2 (E)	HN ServArch	HN ServArch + EGI.eu

	software, information and data				
8	Publicly funded and commercial cloud providers should agree on a minimal set of requirements for IT service management and a related maturity assessment framework that should be adopted by all members of Helix Nebula to evaluate the alignment of their service management practice	1,2,3	Agree on requirements and assessment framework (M)	HN ServArch	HN ServArch + FedSM + EGI.eu
9	Publicly funded and commercial cloud providers should agree on the common set of services needed to participate in the federation	1,2	Identify set of core services to be offered by each cloud provider to participate in the federation (E) Define the technical specification (M)	HN TechArch	HN ServArch/TechArch + EGI FedCloud + WP7
			Define SLA (M)	HN ServArch	
10	Commercial cloud providers should be able to connect to GÉANT, on the basis of an agreed business model, to ensure that users can have the same connectivity level available in publicly funded e-Infrastructures to transfer data	1,2,3	Discuss and identify a suitable business model to connect Helix Nebula's resource providers to GÉANT after pilot phase (M)	HN WP7	DANTE + HN ServArch + HN WP7
11	Both publicly funded and commercial cloud providers should agree on a compatible scheme to describe elements of a service catalogue to ease service selection across	1,2	Define the service catalogue elements that should be exposed in a compatible way by the various resource centres (M)	HN ServArch	HN ServArch + EGI FedCloud

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

	different providers				
12	The issues concerning accounting, billing, payment and settlement should be analysed to determine the possibilities and restrictions for both public funded and commercial cloud providers to agree on compatible accounting and billing parameters, cross settlements between providers and single integrated billing towards the customer.	1,2	Gap analysis of current accounting information available from the various resource centres (E)	HN ServArch	HN ServArch + EGI FedCloud
			Definition of a pragmatic set of parameters for accounting and billing that should be available in a compatible way from the various resource centres (M)	HN ServArch	
13	Users should be provided with a Single Sign-On (SSO) mechanism that enables them to authenticate on and have full control of the cloud services available in the federation. Reuse of existing sources of authentication should be considered.	1,2	Compare the Single Sign-On (SSO) solutions in use by suppliers and evaluate support from BlueBoxes (E) Ask a statement to BlueBox implementers about the possibility to integrate the missing SSOs (E) Offers seamless SSO functionality from the Bluebox (D)	HN TechArch	HN TechArch + EGI Federated Cloud
14	Users should be able to port virtual machine images across different providers without the need for complex translators	1,2, 3	Analyse the results of the flagship deployment and document the gaps (E) Implement solutions to close the gaps (D)	HN TechArch	HN TechArch + EGI FedCloud

15	Users should be able to manage their virtual machines using a uniform interface across the different cloud providers (both publicly funded and commercial). Interfaces should be based on open standards and should not limit the available service functionality. Additional de-facto standard interfaces may be exposed to lower entry barriers and simplify the transition to open standards	1,2,3	Evaluate the functionalities exposed by the BlueBox implementations and identify unsupported features with regards to the resource providers interfaces (M) Extend the BlueBox interface to fill the gaps while using an open standard interface (D)	HN TechArch	HN TechArch + EGI FedCloud
16	Users should be able to create, retrieve, update and delete data elements using a uniform interface across the different cloud providers (both publicly funded and commercial). Interfaces should be based on open standards and should not limit the available service functionality. Additional de-facto standard interfaces may be exposed to lower entry barriers and simplify the transition to open standards	1,2,3	Evaluate the functionalities exposed by the BlueBox implementations and identify unsupported features with regards to the resource providers interfaces (M) Extend the BlueBox interface to fill the gaps while using an open standard interface (D)	HN TechArch	HN TechArch + EGI FedCloud
17	EGI, commercial providers and DANTE should investigate how SDN can benefit a federated cloud computing infrastructure	1,2,3	Create a dedicated task force with a clear mandate and duration (E) Execute the work plan (M)	DANTE	HN TechArch + EGI FedCloud + DANTE

7 Conclusions and Next Steps

This document provides an analysis and recommendations towards the interoperability and integration of commercial cloud providers with publicly funded infrastructures for offering a seamless European federated cloud to the research community.

This work is the result of the first year of activity from work package 6 in collaboration with the various Helix Nebula project partners and other participating stakeholders. A first concrete result was that DANTE joined the Helix Nebula consortium and supported the integration of commercial cloud providers with the publicly funded research and education network. Two workshops have been organised as well as ad hoc task forces have been defined and launched.

The discussion on interoperability aspects has been structured according to the level of concerns identified in the European Interoperability Framework. Integration scenarios have been explored and proposed. A number of high-level recommendations have been identified and agreed across the various stakeholders. Actions for implementing them together with an estimate of the feasibility within the lifetime of the Helix Nebula project have been also provided. Finally, motivations and plan for an interoperability test case have been defined.

During the second year, the activity will focus on the execution of the test case plan and of some of the identified actions for implementing the recommendations. The discussion around requirements for interoperability and integration will be further evolved. The third and last interoperability workshop will be also organised in co-location with the EGI Technical Forum (September 2013). All the work will be fed into the final deliverable to document the achievements and the actions to be carried forward in follow-up initiatives.

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

8 References

R1	European Interoperability Framework (EIF) for European public services http://ec.europa.eu/isa/documents/isa_annex_ii_eif_en.pdf
R2	MS14 Kick-off meeting engaging public and commercial resource providers https://cds.cern.ch/record/1484470
R3	MS15 Technical Workshop (co-located with GA2) https://cds.cern.ch/record/1523745
R4	Helix Nebula Interoperability Task Forces – wiki page https://wiki.egi.eu/wiki/HN_Interoperability
R4	Unleashing the Potential of Cloud Computing in Europe http://ec.europa.eu/information_society/activities/cloudcomputing/docs/com/com_cloud.pdf
R5	Cloud Computing Risk Assessment http://www.enisa.europa.eu/activities/risk-management/files/deliverables/cloud-computing-risk-assessment
R6	Cloud Computing Information Assurance Framework http://www.enisa.europa.eu/activities/risk-management/files/deliverables/cloud-computing-information-assurance-framework
R7	Procure Secure http://www.enisa.europa.eu/activities/Resilience-and-CIIP/cloud-computing/procure-secure-a-guide-to-monitoring-of-security-service-levels-in-cloud-contracts/at_download/fullReport
R8	Federal Information Security Management Act (FISMA) Implementation Project http://csrc.nist.gov/groups/SMA/fisma/index.html
R9	Critical Cloud Computing-A CIIP perspective on cloud computing services http://www.enisa.europa.eu/activities/Resilience-and-CIIP/cloud-computing/critical-cloud-computing
R10	Regulatory framework for electronic communications http://europa.eu/legislation_summaries/information_society/legislative_framework/l24216a_en.htm
R11	Cloud Contracts http://www.eurocloud.at/fileadmin/userdaten/dokumente/the-cloud-contracts-en.pdf

R12	Knowledge without Borders, GÉANT 2020 as the European Communications Commons http://cordis.europa.eu/fp7/ict/e-infrastructure/docs/geg-report.pdf
R13	http://www.ivir.nl/publications/vanhoboken/Cloud_Computing_Patriot_Act_2012.pdf
R14	D3.2 - Minimum requirements for service management in Federated e-Infrastructures http://www.fedsm.eu/downloads
R15	Dassault Systemes EXALEAD: http://www.3ds.com/products/exalead
R16	EGI Security Policy Group: https://wiki.egi.eu/wiki/Security_Policy_Group
R17	Privacy Level Agreement (PLA) Outline for the Sale of Cloud Services in the European Union: https://cloudsecurityalliance.org/research/pla/
R18	FitSM: Standards for lightweight service management in IT infrastructures http://www.fedsm.eu/downloads
R19	EGI Mini Project on the Dynamic Deployments for OCCI Compliant Clouds https://wiki.egi.eu/wiki/Overview_of_Funded_Virtual_Team_projects