Helix Nebula – The Science Cloud

Service Architecture

Abstract: This document outlines the current knowledge of the Helix Nebula Supply-side regarding the need for a structured Service Architecture. It introduces why such an architecture is needed.
Within the Helix Nebula consortium, there are further supporting documents:

- the equivalent if this Introduction, but with much more detailed content;
- a document describing the Service catalogue structure for HN;
- a document describing the operational organisation of HN.

Document Information Summary	
Deliverable title:	*Helix Nebula Service Architecture*
Editor:	*Mick Symonds (Atos)*
Contributing Authors:	*Bernd Schirpke, Jurry de la Mar (T-Systems)* *Giles Hogben (CSA)* *Phil Evans (Logica/CGI)* *Michel van Adrichem (Atos)* *Hervé Caumont (Terradue)*
Distribution:	*Public*
Version/Revision:	*V1.0p*
Keywords:	*ServArch*

Disclaimer

Helix Nebula targets scientific communities, decision makers from SMEs, industry and government as well as European Union policy makers. This document contains information on Helix Nebula core activities, as prepared by the Helix Nebula Supply-side. Any reference to content in this document should clearly indicate the authors, source, organisation and publication date.

This document is work in progress, as are the services it describes. The authors reserve the right to significantly change the underlying assumptions, system and specification described in this document.

Table of Contents

1. Helix Nebula background

The Helix Nebula initiative is building federated cloud capabilities for the European market. Essentially it is capacity building and enabling EU industry to deliver a federated cloud service, initially to address scientific research and associated organisations.

The document "Strategic Plan for a Scientific Cloud Computing infrastructure for Europe"[1], which accompanied the creation of this initiative, describes that it "aims to pave the way for the development and exploitation of a Cloud Computing Infrastructure, initially based on the needs of European IT-intense scientific research organisations, while also allowing the inclusion of other stakeholders' needs (governments, businesses and citizens)".

As part of that overall development, there is an EC FP7 Project[2], which is producing more formal documentation of the process of establishing such capabilities, including creating the intended services structure.

The initiative has established a Service Architecture Working Group (ServArch) in July 2012, with the objective of defining and describing the services that represent and embody what Helix Nebula actually is, and what is its end goal. Key outcomes are interfaces and protocols, at business, service and technical levels, allowing customers and various suppliers to interact, and suppliers in different parts of the value chain to cooperate in delivery of those services. Some scenarios describing those interactions and the related service management aspects are included in this document.

The technical ground work for a federated cloud infrastructure has been defined by the Technical Architecture Working Group (TechArch), defining the technical interfaces between potential customers and providers[3]; the proposed "Blue Box", designed by the TechArch group, may be seen as the reference implementation. From the work it became apparent that there are non-technical aspects which also need to be addressed. Rather than just how something is delivered technically (CPUs, Hard disks, SANs, routers etc…), demand-side requirements should (also) be expressed in terms of the verifiable, comparable service levels (availability, service coverage, maximum change response time, etc., …), which the supply side commits to deliver to the demand side. The actual service delivered should be measurable (via a metering/monitoring infrastructure) and verifiable, as well as comparable and interoperable with services offered by other suppliers.

But there are aspects of federated cloud computing way beyond just enabling the use of Infrastructure as a Service (IaaS), which is anyway usually a means to an end rather than an end in itself. To quote one of our key customers, "our key objective is to create a 'virtual Market Place', where data, scientists, funding bodies, SMEs and downstream

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

[2] See http://cordis.europa.eu/search/index.cfm?fuseaction=proj.document&PJ_RCN=12963846

[3] See http://cdsweb.cern.ch/record/1478364

industry are meeting to work along common interests. This ecosystem should represent a many-to-many relationship, quickly being established, jointly to transform data to valuable information". This "bigger picture" necessitates not only higher-level cloud services themselves (i.e. PaaS and SaaS) but, more importantly, a way for very different organisations to work together seamlessly to deliver those services.

A research and business ecosystem could combine ways to handle and manage the bulk data and processing involved, and provide facilities for all sorts of added-value contributions. These can come from many scales and collaborative combinations of organisations, both large and SME's, who can build services around and on top of the platforms, whether they have their own facilities or not. The benefits could mushroom, while starting from a simple base.

The ServArch work addresses some of those aspects in defining a services architecture. This document represents an introduction to what ServArch is working on, and a general reference on services architectures for the other members of the consortium and those who interact with it.

The internal equivalent of this document includes more supporting information on a number of subjects:

- Security and compliance determination
- Contractual arrangements for service delivery
- Service descriptions and levels
- Processes, organisation and tooling to support service delivery
- Costing and pricing models
- Life-cycle management and release planning
- Other added-value services
- The relationship to infrastructure and technology developments
- Maturity models in service architecture

Separate documents have also been prepared by ServArch to describe:

- The Service structure: how the services will be defined, what service level elements will be delivered, etc.; the "contents":
    - defines the framework for describing the catalogue of Helix-Nebula services in terms of verifiable service attributes;
    - introduces key concepts, of service catalogues, classes, attributes, resource groups, and commitments, reports and triggers;
    - identifies high-level service classes and level packages used within HN;
    - provides extensive tables in which the attributes can be described in a standard way.
- The Organisation and Process structure: how the demand/supply model will be constructed:
    - the Actors involved, and their objectives and requirements: Demand, Supply and possible Broker;
    - processes required to establish a service for any customer and manage its delivery;
    - the organisation structure, including Demand/Supply model;

- tool requirements, to support the above.

2. Business and service scenarios

In its initial development phase, Helix Nebula has started off as a relatively simple M-to-N structure: on the Demand side there are three major research organisations, and the Suppliers consist of a few (currently four) sizable IaaS providers supported by various other SME's with specialist skills and knowledge.

However, it is expected that over time it will develop into a much more complex ecosystem, with a number of further players, at all stages of the supply spectrum. Roles could include:

- Funding agencies, who rather than supply researchers with capital for IT equipment provide them with "points" to access the Helix Nebula facilities;
- Research organisations who do not actually procure facilities themselves, but specify and even facilitate the procurement by associated organisations, SME's, etc.;
- SME's who act as added-value service and information providers, using Open Data from the research organisations, and/or providing processing facilities by brokering them from the Helix Nebula suppliers and incorporating them into a service;
- Pure-play brokers, who survive and add value by finding the best facilities for any one customer and making them more easily available to them, possibly also at a lower cost.

The following scenarios illustrate some of the intentions of Helix Nebula which drive the contents of this document.

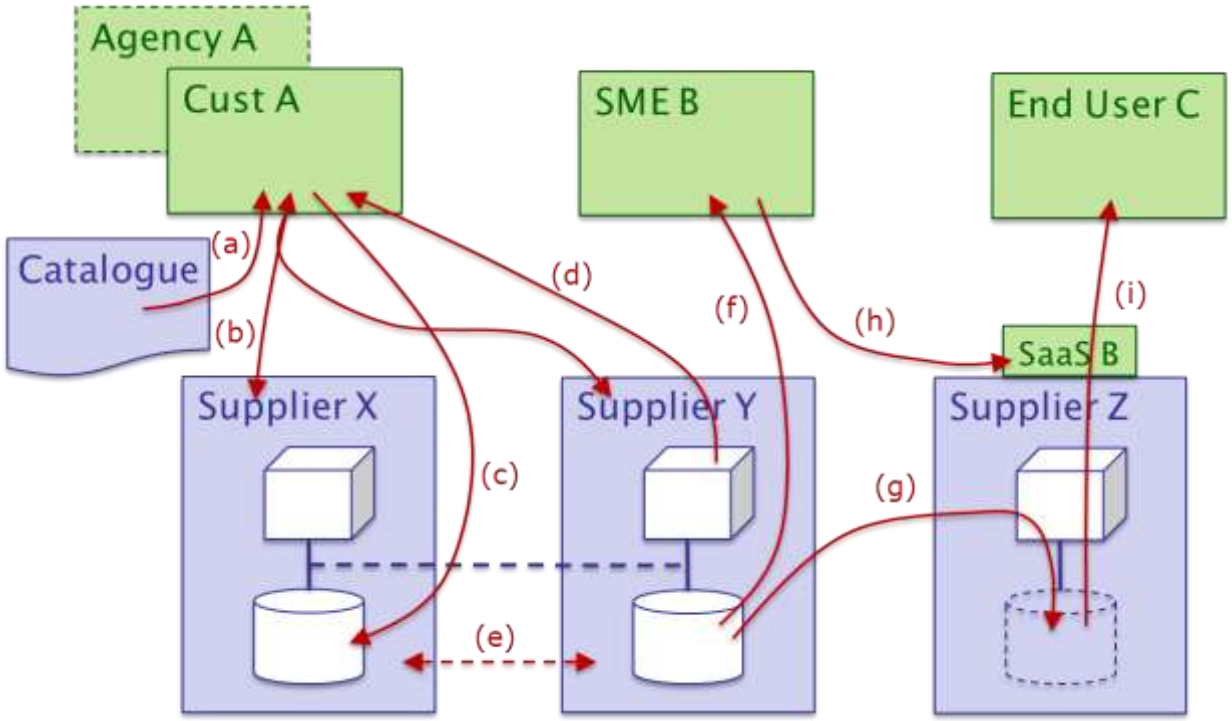

(a) Customer A determines what services are required, and available from the catalogue;
(b) then invokes supply arrangements with Suppliers X and Y;
(c) stores data with Supplier X;
(d) invokes processing capacity from Supplier Y;
(e) Suppliers X and Y manage data availability and access;
(f) SME B determines the location of and access to required (possibly Open) Data;
(g) pre-processes data, with Supplier Z, to produce useful information;
(h) establishes SaaS-based Information service with Supplier Z[4];
(i) End User C invokes Information service.

Meanwhile:

- Support: there are coordinated points of contact between Suppliers X, Y and Z, to provide necessary support facilities to all users;
- Service levels and reporting: to agreed coordinated levels and formats;
- Billing: for storage, processing, information, etc.;
- Payments: from funding Agency and End Users to Suppliers for storage and processing, and/or Customer A for access to their Data and/or SME B for their Information; with back-payments where between them, as appropriate.

Note: the role and position of a Funding Agency "behind" Customer A may seem Helix Nebula-specific, but is akin to the strategic position of a CxO-level payment

[4]See also *Building a 'Big Data' Ecosystem for Science*, Micheál Higgins and Hervé Caumont

authorisation behind many businesses' IT departments. It is further explained within the main document in the section on the Demand/Supply model.

3. Scope

There are differing depths of knowledge of the services architecture required by different parties, as shown by the following triangle:

- At the highest level, customers need to know what services are being offered, what they cost and how to provision and interface with those services;
- Suppliers need to know how to deliver those services to the customers, and how to work together where required;
- There is potential for one or more central "broker(s)" to coordinate services from multiple suppliers;
- It is up to the suppliers themselves to know how they will actually manage their own service delivery in accordance with the interfaces and protocols defined.

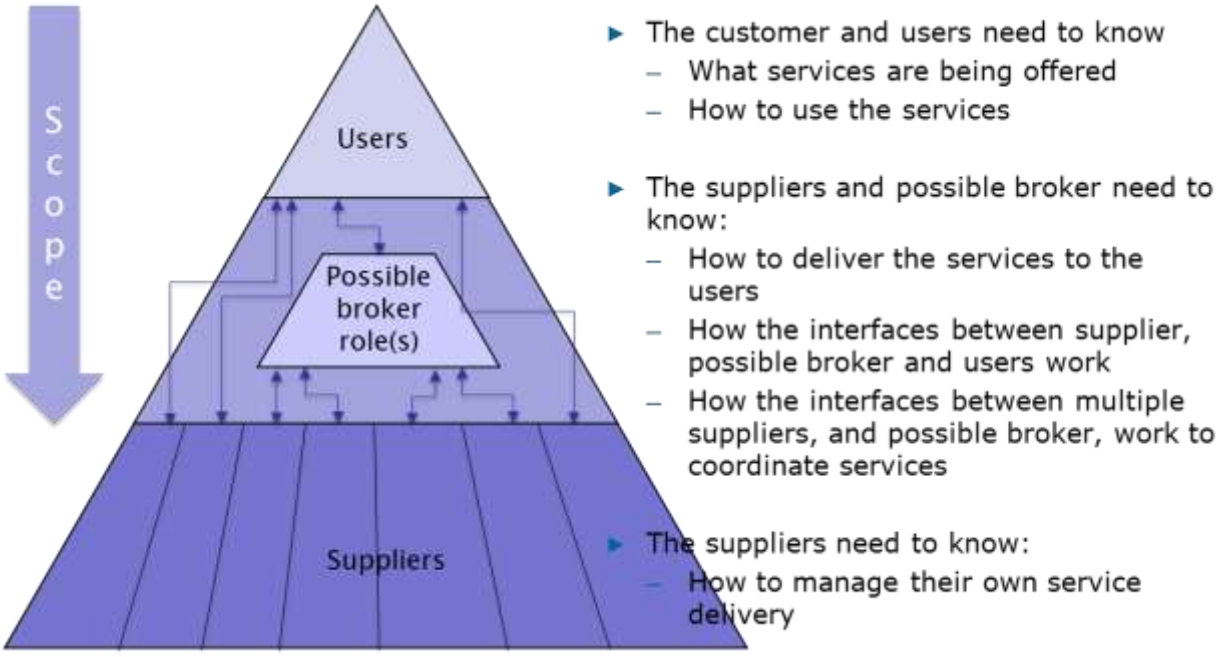

4. Understanding customers' requirements

One of the exercises undertaken within Helix Nebula, first under the aegis of the TechArch group and then under the EC Project's WP3, was to capture the requirements of the customers. Whilst this provides some of the service-related requirements, it is mainly oriented around their technical requirements, rather than those service aspects.

The good news is that the generic requirements of the majority of IT users are by now quite well understood by suppliers, and more work is being done to collate those by bodies such as the ODCA[5]. That is just as well, because the nature of cloud services is

[5] http://www.opendatacenteralliance.org/

that they are pre-configured to commoditised standards, so that they can be assembled as required, rather than being built and deployed in any previous build-to-order model. A cafeteria model (like McDonalds) applies: the supplier has to decide what they will deliver, and then fine-tune it based on take-up. The subscriber selects the best fitting service from the available options provided by the various suppliers, and accepts any limitations of the selected service.

The requirements thus have to be understood in two different cycles:

- To establish services: generic requirements need to be understood as input to service building plans;
- To invoke services: specific needs for any particular case usage, to match to the services then available.

An additional requirement that arises within a body such as Helix Nebula, which relies on collaboration for its success, is for the suppliers involved to subscribe to a standard interface for service and discovery, service level commitment and monitoring, and potential redress, to achieve the required results.

This means that not only do they address the customer in a common way, but that they can themselves cooperate where required in order to deliver multi-supplier services (see triangle).

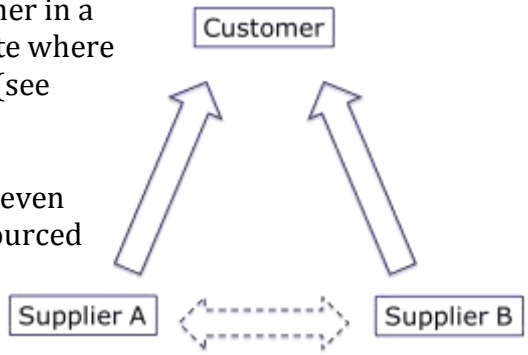

The good news is that most suppliers have years, or even decades, of experience of co-existing within multi-sourced outsourcing contracts. Industry standards such as ITIL[6] can be used to provide common terminology and common processes, so that services can be delivered and support can be provided across those multiple suppliers.

5. Defining the services

The services available need to be defined so that users can understand quantitatively what they are getting, at what quality and price. That needs to be done in a consistent way:

- Firstly, so that they can recognise comparable services from multiple suppliers, and "compare apples with apples" in terms of service quality; and
- Secondly, so that consistent supplier service offers can be used during the matching to any specific user request.

Again, industry bodies such as the ODCA provide some guidance on the format of service catalogues, so that they can be compared and collated between suppliers[7].

The service levels also need to be defined as committed quality levels. That can be done in terms of:

[6] http://www.itil-officialsite.com/

[7] See ODCA Master Usage Model: Service Orchestration rev 1.0

- inputs (the service is to "do something", e.g. we will react when something happens), and/or
- outputs (the service is to "deliver something", e.g. availability levels) and/or
- outcomes (the service is to "achieve something", e.g. user satisfaction),

although a combination of all of those is usually required for differing purposes. This is explained further later in this document.

This definition of levels is often done using bands, such as Bronze, Silver, Gold and Platinum, where Bronze represents the bare minimal level and Platinum the highest quality of service[8]. It will be appreciated that these are predictions of service quality, and some sort of assurance may be required that those quality levels will actually be met and/or penalties if they are not. That, in turn, dictates that a means is required of recording and reporting on actual service levels delivered.

Services have costs, usually defined for each service level: higher service levels costing more. A common structure for the pricing model allows easier comparison to be made between competing offers. In order to avoid any potential cartel, it will be evident that Helix Nebula may only recommend pricing structures, but that pricing levels will always be subject of an exclusive confidential discussion between an individual customer and supplier.

6. Security and compliance

Whilst cloud is usually characterized as not knowing where your data is held, that is a concern within the context of European data privacy and protection legislation. Thus, it may be necessary to hold data in known locations and confine specific instances of it to one country. More work is being done on identifying the impact of such restrictions and the effects of those measures on the cloud services.

These days, all users of IT are subject to compliance with various laws and regulations, whether from their national government and/or industry regulators. Suppliers need to be prepared to adhere to a superset of these, for all of their customers, which is why the concept of Community Cloud has emerged: a cloud environment which is designed and built to contain the necessary measures for a given set of users. In this context, Helix Nebula at least begins as a community cloud, established to fulfil the common needs of the scientific research sector in Europe with a common set of security policies and compliance regimes.

Security requirements are also expressible within the framework of a service architecture, as has been shown by work from ENISA[9] and CSA[10]. Security requirements can themselves be expressed as service level attributes (availability, incident response time, key length, etc.), and can be linked to certification (providing a point-in-time

[8] As is embodied in the ODCA document on Standard Units of Measure.

[9] http://www.enisa.europa.eu/activities/application-security/test/procure-secure-a-guide-to-monitoring-of-security-service-levels-in-cloud-contracts

[10] https://cloudsecurityalliance.org/research/ctp/

assessment of requirements) and security monitoring capabilities to provide continuous assurance.

That compliance may also extend in the future towards a need to identify the environmental impact of given services, for instance in terms of the carbon generated by their use[11] for Corporate Social Responsibility (CSR) purposes.

7. The services organization

Although as much as possible of the service delivery is automated, there still needs to be a human organization around it, to establish it, maintain it, and intervene if things go wrong. There must be a way into that organization, to find someone to handle queries or request support.

Again, within a support structure, indications have to be given of availability and committed responsiveness: how quickly will an initial response be given, will a problem be fixed overnight, or will it have to wait for the next working day, for instance? Will updates be given of progress towards resolution?

A matching and balanced structure needs to be put in place in both the demand and supply side of the relationship, including hierarchies from strategic relationships down to operations. These allow different timescales and perspectives, such as from how do we strategically plan future services down to getting operational problems fixed. In a multi-supplier environment, these need to be coordinated, to prevent conflicting service developments and also to prevent problems being "bounced" between providers.

The federated nature of the Helix Nebula cloud service means that several suppliers will deliver a part of the cloud service for a flagship to an Institute. This leads to a demand to coordinate certain activities between the demand and supply. The various roles with possible added value to the Helix Nebula partners are:

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

[1111] See the ODCA document on Carbon Footprints

with integrated service reporting and billing. This role is also known as Service Integrator.

It will be seen that there is some overlap and the potential for complementarity between these roles, but that they need not be unified. Further, it is possible for multiple such brokers, etc. to coexist.

8. Co-existence with other existing facilities

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

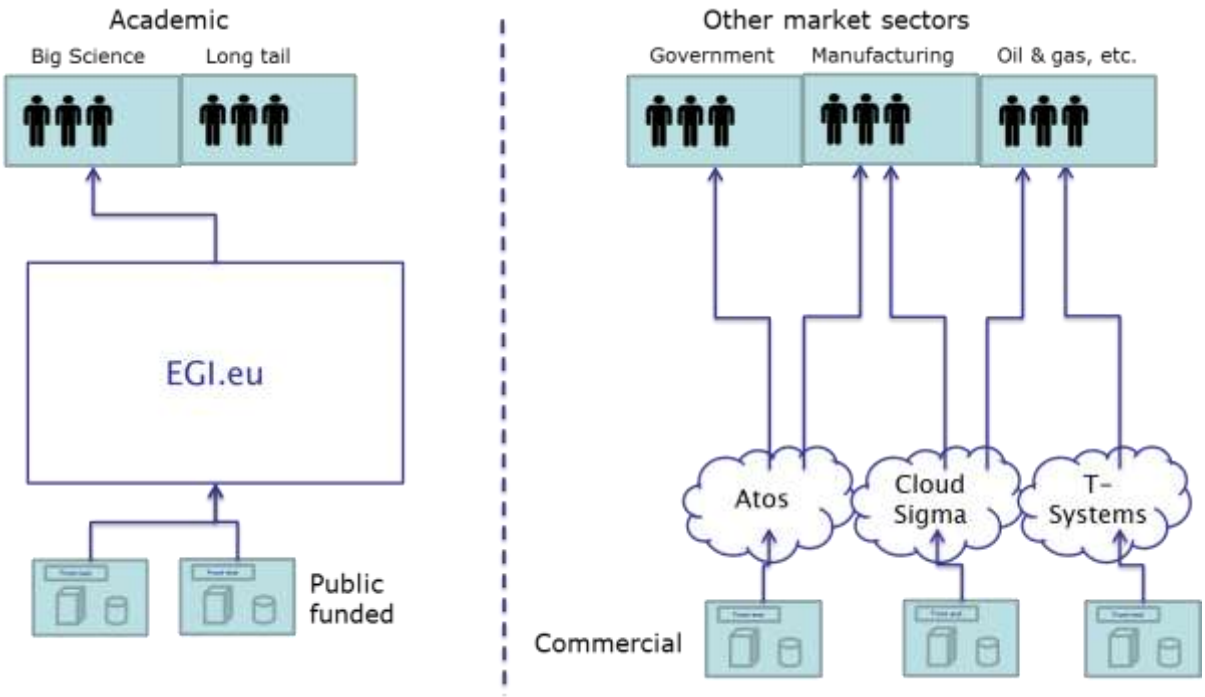

We can envisage a number of developments taking place in the coming period:

- EGI will themselves develop more sophisticated service management mechanisms, via their own Federated Cloud Task Force and the Federated Service Management (FedSM) project of the EC;
- In doing so, they may position themselves better to address the "long tail" of small research organisations, as well as the large research institutions;
- The commercial service providers are aligning their services within Helix Nebula by the development and deployment of a "Blue Box", to provide a common front-end to and inter-connection between them;
- The commercial service providers will be in a position to offer higher-level services than just commodity infrastructure to this market, e.g. PaaS and SaaS facilities;

- The transparency and sustainability of the market would need to be assured by EGI becoming a fully-costed/charged service organisation, not surviving on a marginal costing basis;
- Compatible service management architectures and interfaces would need to be put in place so that, for instance, changes and trouble tickets could be shared and coordinated between all parties;
- Then each group, EGI and Helix Nebula, could offer services to/via the other, whether of a commodity or specialist nature.

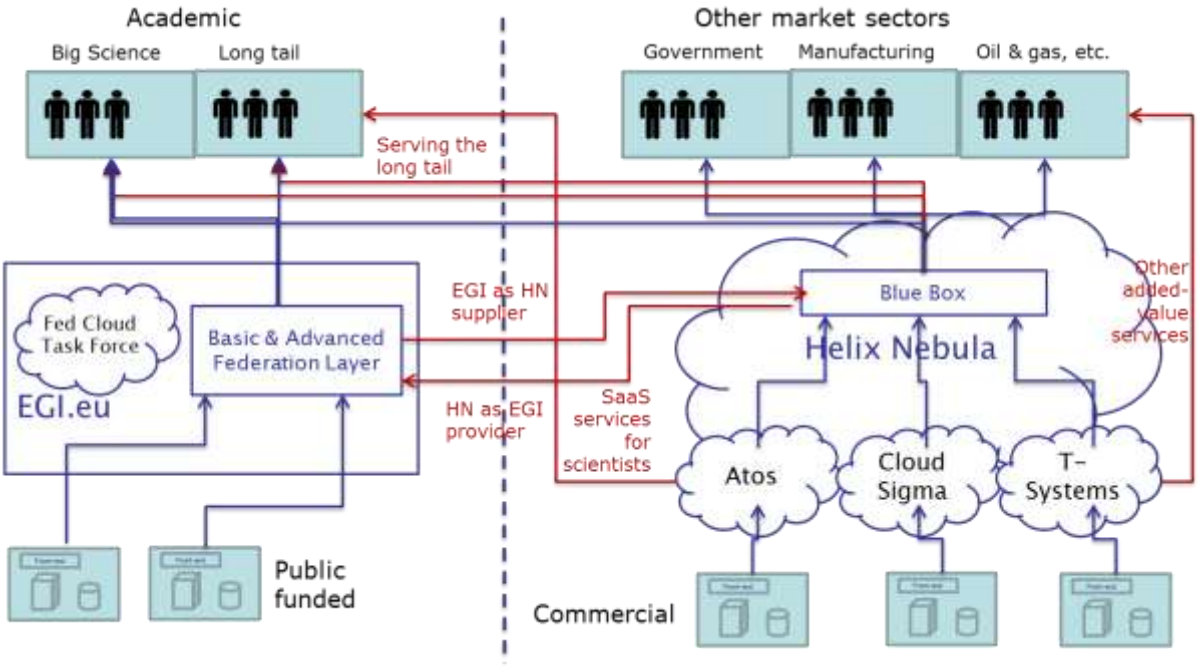

9. Service development and maturity

Two of the major concepts which are well understood in service delivery circles are those of Demand/Supply models and Maturity Models. As explained, the former outlines how a complex, multi-dimensional customer organisation interacts with a similarly complex supplier organisation. The latter shows how many aspects of service delivery can be evaluated as to how well-implemented they are, and that they all need to be developed and matured together in order for service delivery to be effective. Both are further explained in the internal HN Service Architecture document.

The two aspects come together in a demand/supply maturity model.

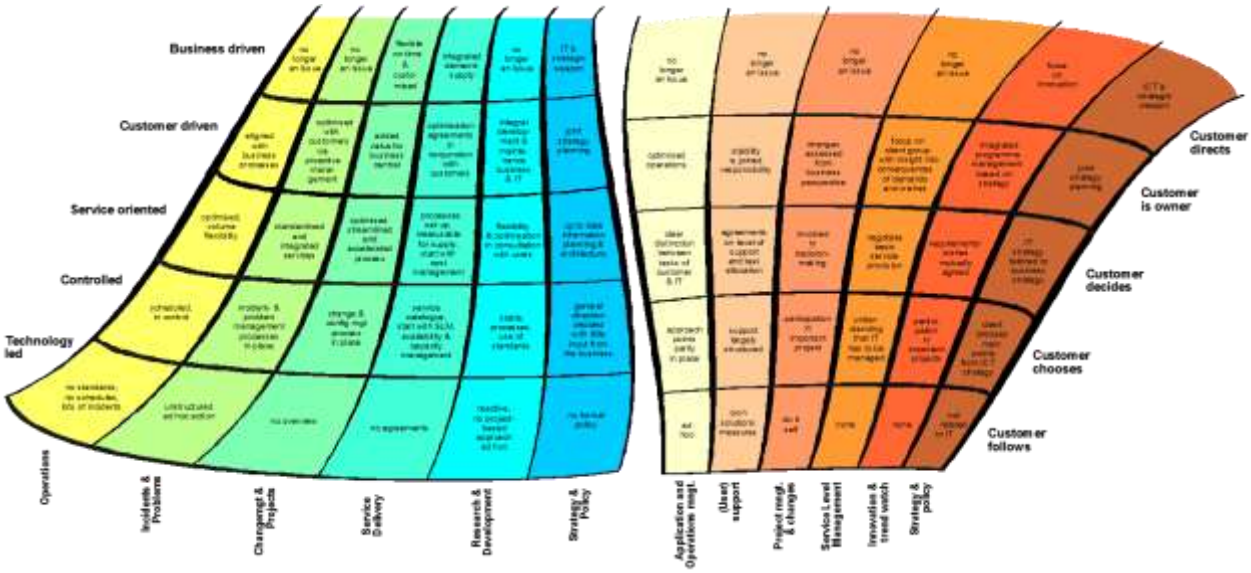

This maps out characteristics of the service delivery as it matures from being technology-led to business-driven. It also illustrates a few key aspects:

- The maturity of the Demand and Supply sides need to be in line with each other;
- When combining two organisations together, you typically drop to the maturity level of the lower, or beyond, by doing so;
- Experience tells us that it takes around one year to develop from one level of maturity to the next.

The relevance for Helix Nebula is that we currently have very different levels of awareness of service delivery on both sides of the supply spectrum, not just in one of them. So suppliers, themselves of varying service maturity, are delivering services to customers of varying maturity.

We can envisage an environment where Helix Nebula can develop over time in a number of dimensions, e.g.:

- From infrastructure provider to research and business ecosystem;
- In Big Data management;
- From IaaS to PaaS and SaaS;
- From simple generic infrastructures towards the provision of other variations of hardware and infrastructure services, such as High Performance Computing (HPC), large memories, etc.

The Service Architecture documents cannot map out in one move how we can implement something to fulfil all of our needs overnight. But they can and do map out some of the key aspects, and give some road signs for further development in the future.