Helix Nebula – The Science Cloud

Title: Consolidated User and Service Requirements Report

Editor: Cloudsigma

Work Package: WP3

Deliverable: 3.2

Submission Date: 20.12.2012

Distribution: Public

Nature: Report

Abstract: This document aims to consolidate the lessons learned from the three proof of concept (PoC) Infrastructure as a Service (IaaS) environments, with the scope of offering compute power to ESA, CERN and EMBL in a utilitarian way making compute power a commodity such as water, electricity and power.

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

Log Table

Date	Version	Contributor(s)
10/09/2012	0.1	Robert Jenkins
26/09/2012	0.2	Plamen Ganchosov, Micheal Higgins, Robert Jenkins, Bernino Lind
12/10/2012	0.3	Bernino Lind
24/10/2012	0.4	Micheal Higgins, Bernino Lind
30/10/2012	0.5	Micheal Higgins, Robert Jenkins
05/11/2012	0.6	Plamen Gantchev, Robert Jenkins
06/11/2012	1.0	Patrick Baillie, Robert Jenkins, Bernino Lind
13/11/2012	1.1	Plamen Ganchosov, Robert Jenkins, Bernino Lind
27/11/2012	1.2	Plamen Ganchosov, Robert Jenkins, Peter Pruijssers, Sjaak Theelen, Michel van Adrichem, Ana Juan Ferrer
13/12/2012	1.3	Bernino Lind, Plamen Ganchosov, Robert Jones, Rachida Amsaghrou
19/12/2012	1.4	Bernino Lind, Plamen Ganchosov, Robert Jones, Rachida Amsaghrou

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

Table of Figures

Table of Tables

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

Introduction

This document aims to consolidate the lessons learned from the three proof of concept (PoC) Infrastructure as a Service (IaaS) environments, with the scope of offering compute power to ESA, CERN and EMBL in a utilitarian way, making compute power a commodity similar to water, electricity and power.

These PoCs were created as part of the Helix Nebula initiative and deployed by T-Systems, Atos and CloudSigma, during the course of 2012.

All three flagship use-cases share a common theme of multi-institutional collaboration, which will be a major factor in the speed of uptake of cloud-based computing. CERN estimates that cloud services for the LHC community could impact as many as 10,000 physicists. EMBL reckons that cloud infrastructure could enhance the relevance of genomics to the broader medical, pharmaceutical and agricultural communities.

Each flagship has been selected for the unique challenges and opportunities it presents, as described below: The ATLAS Distributed Computing environment (ADC) which consists of several pieces

- A distributed Data Management component (DDM)
- A distributed production and analysis system (PANDA), and associated tools as well as the processing and analysis binaries.

This flagship is part of a wider project within ATLAS, to research the applicability of cloud computing to ATLAS and CERN's computing needs.

The second PoC was provided by EMBL and aims to establish a leading bioinformatics pipeline to perform fast and on-demand genomic data analysis and furthermore to provide a basis for future extensions of genomic research using cloud computing Infrastructure as a Service (IaaS).

The third PoC was provided by ESA and aims to achieve an open source, unified infrastructure for solid Earth data, to secure solid Earth data sharing via the cloud and to further enable international collaboration and enlargement of the science user base.

The findings from both the supply side and demand side participants have been considered and included in this document. Additionally, recommendations and consolidated feedback from the technical architecture committee were utilised in formulating the various sections and contents of this document.

Deliverable D3.1 (http://cdsweb.cern.ch/record/1484437) was used as a framework to provide an online form for capturing the features and requirements of future demand side participants. This online framework is subject to continuing review as various work package activities continue to refine the requirements of the Helix Nebula workloads.

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

It is expected that this document will continue to evolve as part of WP3 during the course of the project. Additional requirements around business models and service architecture are expected to be added and refined as they are still subject to intensive discussions and reviews.

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

1. Proof of Concept: Lessons Learned

The first round of PoCs showed that suppliers within the Helix Nebula consortium can successfully deploy demand side workloads within their IaaS. Furthermore, clear areas of success and areas of improvement were identified, providing the consortium with a roadmap to establish the upcoming and more advanced round of flagship deployments.

1.1 General Best Practises

Several non-technical areas of feedback were identified in relation to future deployments and flagship candidates, which are at the basis of best practises for Helix Nebula.

First, rather unsurprisingly considering that it was the first round of proof of concepts, the scope and key success criteria of the use cases were identified as requiring further detailing. This tended to cause longer than anticipated deployment phases with the suppliers. Had upfront information been better, additional ad hoc planning would not have been required. However, it should be noted that the underlying problems and solutions are to some extent unknown. A project management methodology, which addresses this issue effectively, is Agile Scrum[1], which could be considered as a framework for future PoCs.

Secondly, operationally-speaking, in all but one case, the suppliers deployed proof of concept environments in series. Taking a serial approach diminished the effort required to provide and acquire valuable feedback and experience from the previous supplier. In reality, a serial approach significantly reduced the combined time across suppliers to achieve the proof of concept aims as well as fostering a much greater spirit of cooperation and higher levels of communication. It is therefore recommended that the deployment model for the future be one of running in series and not in parallel.

Leaving an appropriate gap between a given supplier and the next to deliver a mature exit document saved time for the overall proof of concept deployment phase. Suppliers should take turns in being the first on a particular new flagship in order to share the additional charge of conducting the first deployment. Leaving a conservative gap between the active deployment phase and the next supplier facilitated and improved communication by allowing proper exit documents and deliverables to be created and reviewed by the supply group as a whole prior to starting the next deployment.

1.2 Server Image Conversion

The proof of concept process revealed that suppliers were relying on a wide variety of technologies to deploy their cloud offerings. This resulting in a general lack of commonality in server formats across the suppliers. No

[1] http://jeffsutherland.com/scrumhandbook.pdf

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

formal process existed for accepting server images in the 'correct' format per supplier nor was there any formal methodology for converting server images from one format to another to allow transfer between suppliers. This resulted in additional workload for demand side participants in liaising with the various supply side participants.

Consequently, a key aim for the Helix Nebula consortium is to automate the ingress and egress of server and drive images with respect to any Helix Nebula supply cloud, but also allow the smooth transfer/conversion of drive and server images between suppliers. As a result, the following types of conversions were identified as essential.

Image type	Convert to	Tools&Guides
KVM	Xen, VMDK, OVF	QEMU-img
Xen	Kvm, VMDK, OVF	Xen to KVM
VMDK	Kvm, Xen, OVF	Citrix converter
OVF	Kvm, Xen,VMDK	OVF converter

Table 1 Transfer Conversion of drives/server images

1.3 Networking

The scientific high-performance computing (HPC) profiles of the flagships exhibited some common networking requirements summarised below:

- public internet connectivity was generally required
- high bandwidth connections were preferred to enable large data transfers within short time periods
- bandwidth usage was considered irregular

The above usage profile clearly favours a volume-based pricing model on large capacity networking connections. Inter-supplier networking and networking between demand and supply side partners has been identified as a key area for cooperation and improvement in the future. The low-latency cost effective transfer of large amounts of data whilst maintaining a high data integrity level is a key success criteria for the Helix Nebula consortium. Granularity of this specific requirement could be further decomposed into more concrete SMART technical details (Specific, Measurable, Attainable, Relevant, and Time bound) in the long term process of developing Deliverable 3.3: Finalised User and Service Requirements Report.

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

1.4 Legacy Software Systems

A key finding of the proof of concept stage was the critical role played by legacy software systems used in the scientific processes at the demand side institutions. In a number of cases these required the modification of key interfaces that drove provisioning engines, support for custom server images and specific networking.

It was found that the best way to deal successfully with legacy system requirements was to document them in a detailed way upfront and to allow supply side participants to fully access requirements and impacts ahead of active proof of concept deployments. These additional requirements have been incorporated into the demand side requirements framework.

1.5 Provisioning and Scaling

The ability to scale capacity and provision resources in an automatic way was identified as a key requirement. The supply side participants used a variety of API interfaces to achieve this functionality, again reflecting differing underlying technologies.

In some cases full API coverage of a supply side platform was not available - which, while acceptable during the proof of concept stage - would need upgrading for any production deployment.

A true cross-cloud common API interface to allow easy multi-cloud deployment and management in the medium term is a requirement and an aim of the Helix Nebula consortium. As potential candidates for achieving this functionality based on the requirements gathering phase, the following were identified as potential solutions:

- Jclouds[2]
- DeltaCloud[3]
- OCCI[4]
- CIMI[5]
- Slipstream/StratusLab[6]

[2] http://www.jclouds.org/
[3] http://deltacloud.apache.org/
[4] http://www.ogf.org/documents/GFD.184.pdf
[5] http://dmtf.org/sites/default/files/standards/documents/DSP0263_1.0.0.pdf
[6] http://indico.cern.ch/conferenceDisplay.py?confId=175046

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

- OpenNebula[7]

1.6 Public versus Private Cloud Deployment Models

Beyond differences in underlying technologies, many of the key findings and areas for further work and cooperation stem from fundamental differences in deployment models between supply side participants. These could be summarized generally into two approaches, those providing fundamentally private cloud deployments and those offering public cloud deployments. The former tend to be built around private connectivity only, providing less dynamism in resource allocation and less automation but conversely offering a private environment. The latter tend to offer public internet connectivity and dynamic resource allocation with APIs for resource management. As such it was found that public cloud deployments tended to be a better initial deployment fit for the operational models required by the demand side flagships. The proof of concepts demonstrated that public cloud deployment models can be used effectively as part of Helix Nebula, the Science Cloud. Here is a simplified summary of key differences between the two models:

Quality	Private	Public
Connectivity	Private, fixed	Public, bursty
Tenancy	Single	Multiple
Service delivery	Bespoke	Standardized
Service model	Supported	Self-service
Capacity planning	With lead time	On-demand

Table 2 Public vs. Private cloud simplified differences

A key challenge for the Helix Nebula consortium is to find workable deployment models able to support both primarily private and primarily public cloud environments in order to maintain the widest possible supplier base for flagship computing requirements.

A hybrid model supporting both private and public clouds will enable current and future demand side participants to benefit from a wider choice of cloud infrastructure and services. Such hybrid deployments can seek to benefit from the strengths of each deployment model and combine them to create a highly effective solution.

[7] http://opennebula.org/

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

Based on the NIST set of definitions[8], the Helix Nebula model would most closely align with the 'community cloud' model, namely:

"The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises."

2. On-going Requirements Gathering Framework

The findings of the proof of concept process have been used to update the requirements gathering framework as reflected in deliverable 3.1 (http://cdsweb.cern.ch/record/1484437).

This has been converted into an online tool allowing the easy capture of future demand requirements and supply side capabilities. The online form captures and stores the information into a consolidated spreadsheet[9] which is then distributed to the relevant Helix Nebula consortium participants.
The form is available at https://sites.google.com/a/cloudsigma.com/helix-nebula/

[8] http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf/

[9] https://docs.google.com/spreadsheet/ccc?key=0AoxxLikF1p7gdDdqMzVNMWxsMGRsOTFJN2tYRnR5WGc#gid=0

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

Flagships Requirements Gathering Template

The goal of the requirements gathering is to create insight for the participating suppliers. Although the participating research organizations can state the current situation and their expectations in this template, it doesn't mean that the resulting Science Cloud will support all of them. The participating suppliers will validate the information collected in this assessment against realistic delivery options (Task 4.2 of Work Package 4) and inform the participating research organizations about the results. Here the iterative approach shows its value. After an iterative step of requirements gathering and validation against realistic delivery options a next iterative step is executed. In this next iteration both requirements and delivery options can be adjusted to create alignment of the two.

* Required

Scientific Organisation(s) sponsoring the flagship: *

Contact person (name, affiliation, email): *

Scientific Objective:

Summarise the scientific objectives for the flagship in laymen's terms.

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

Example Server Type B

Please set parameters of the server type A

CPU

GHZ/core

RAM

GB/per VM

HDD

GB/per VM

SDD

GB/per VM

Network

Gbps/per VM

1.1.6 Virtualization

If you currently use virtualisation technologies and may rely on specific hypervisors please outline your usage currently. Select all applicable options for the flagship legacy deployments.

Select all applicable options for the flagship legacy deployments

	Semi	Full
KVM	○	○
XEN	○	○
VMware	○	○
Hyper-V	○	○
ESX	○	○
Other	○	○

Other

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

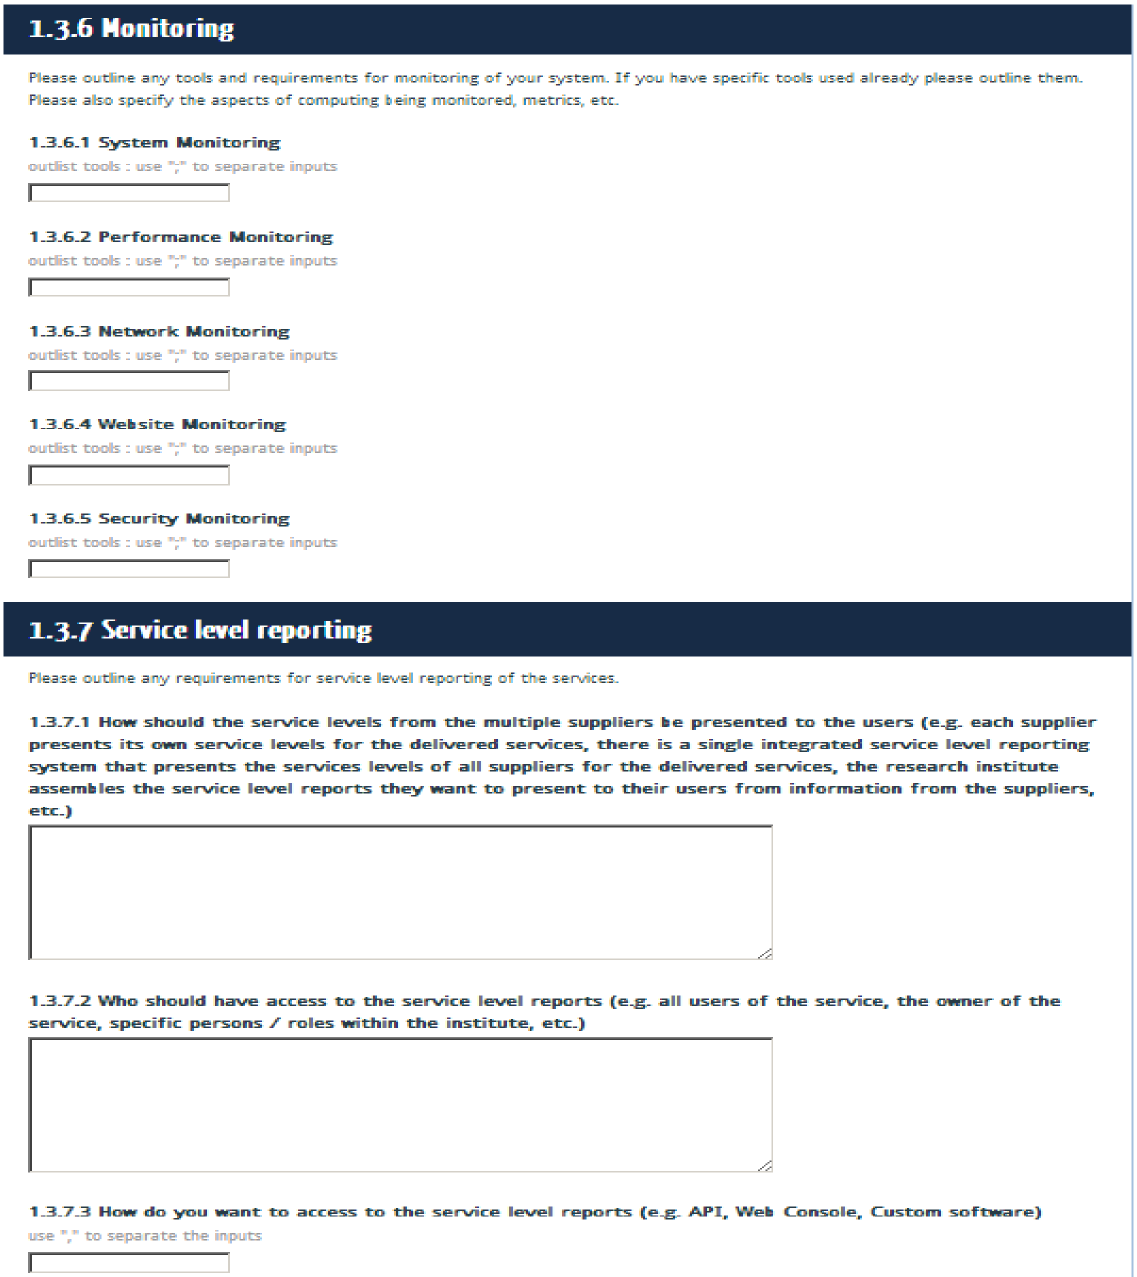

Figure 1: Illustrative screenshots from the online requirements gathering tool

The spreadsheet stores all entries in a structured form for further analysis and review as part of the consideration process in relation to all new proposed flagships.

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

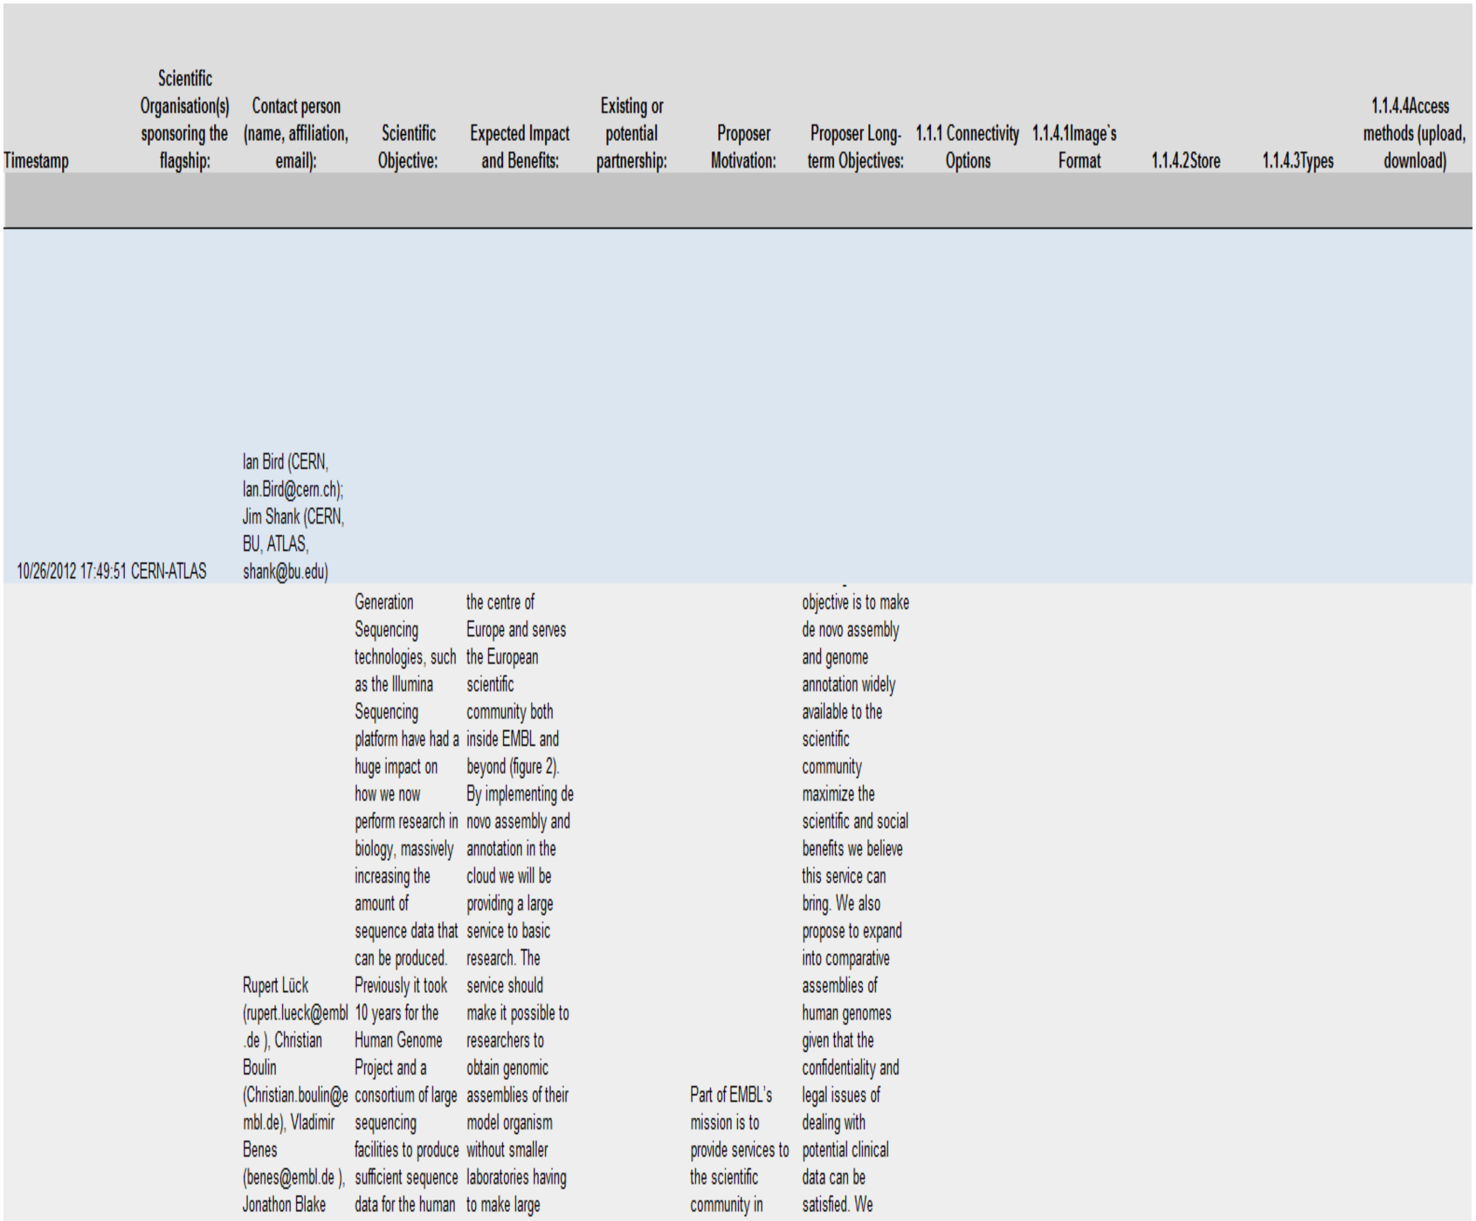

Timestamp	Scientific Organisation(s) sponsoring the flagship:	Contact person (name, affiliation, email):	Scientific Objective:	Expected Impact and Benefits:	Existing or potential partnership:	Proposer Motivation:	Proposer Long-term Objectives:	1.1.1 Connectivity Options	1.1.4.1Image`s Format	1.1.4.2Store	1.1.4.3Types	1.1.4.4Access methods (upload, download)
10/26/2012 17:49:51	CERN-ATLAS	Ian Bird (CERN, Ian.Bird@cern.ch); Jim Shank (CERN, BU, ATLAS, shank@bu.edu)										
		Rupert Lück (rupert.lueck@embl .de), Christian Boulin (Christian.boulin@e mbl.de), Vladimir Benes (benes@embl.de), Jonathon Blake	Generation Sequencing technologies, such as the Illumina Sequencing platform have had a huge impact on how we now perform research in biology, massively increasing the amount of sequence data that can be produced. Previously it took 10 years for the Human Genome Project and a consortium of large sequencing facilities to produce sufficient sequence data for the human	the centre of Europe and serves the European scientific community both inside EMBL and beyond (figure 2). By implementing de novo assembly and annotation in the cloud we will be providing a large service to basic research. The service should make it possible to researchers to obtain genomic assemblies of their model organism without smaller laboratories having to make large		Part of EMBL's mission is to provide services to the scientific community in	objective is to make de novo assembly and genome annotation widely available to the scientific community maximize the scientific and social benefits we believe this service can bring. We also propose to expand into comparative assemblies of human genomes given that the confidentiality and legal issues of dealing with potential clinical data can be satisfied. We					

Figure 2: Sample screen captures of the aggregated data - 1

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

	EH	EI	EJ	EK	EL	EM	EN	EO	EP	EQ	ER	ES	ET	EU
1	1.6.0.4 Usage Profile	1.6.1.1Block Device Storage	1.6.1.2Volume Storage	1.6.1.3Object Storage	1.6.2.1Geographical data storage requirements	1.6.2.2Large Volume Data Transfer	any preferred or required billing and purchasing models, for example 'burst' based resource	organisational requirements regarding reporting, invoicing, account use-age	1.8.1SLAs	1.8.2Terms of Use	1.8.3Privacy Policy	1.8.4Acceptable Use Policy	1.8.5Intellectual Property Management	1.8.6Vendor Neutrality
2														
3	http://wlcg.web.cern.ch/collaboration/reporting/reliability-availability				Preferably in Europe		The contract price shall be net, firm and inclusive of all costs relating to the performance of the contractor's obligations under the contract and take into account CERN's exoneration from VAT and import duties	Subscription with warning if upper limit reached	Depends on cloud service					Yes, need to be able to migrate code + data.
		for every new EMBL genomic assembly and annotation (steps 1-5 above in overall architecture description) a new GlusterFS based fast shared file system will be build. GlusterFS will aggregate ~4 VM instances and the attached storage (~2.5 TB of total block based storage) to create a shared file system of ~1.2 TB net capacity) In addition to the GlusterFS needs and on top of what is specified below,	If high performance		particular customer uploads for an individual genome analysis (e.g. 0.8TB Next Generation Sequencing Data) needs to be in the same data centre as the compute capacity (e.g. 50-100 compute VMs) and connected via 10Gbit (optimal) or 1Gbit (feasible). However, different individual genome analyses can operate in different data centres after initial setup. The initial setup involves the installation of a	Large data transfer will occur when customers (such as EMBL or other labs) will upload Next Generation Sequencing data (e.g. 0.5-1.0 TB) for	For the initial setup				It will be required that any data which will be sent to and from the Helix Nebula cloud infrastrcuture as	It will be required that any data which will be sent to and from the Helix Nebula cloud infrastrcuture as	It will be required that any data which will be sent to and from the Helix Nebula cloud infrastrcuture as	

Figure 3: Sample screen captures of the aggregated data - 2

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

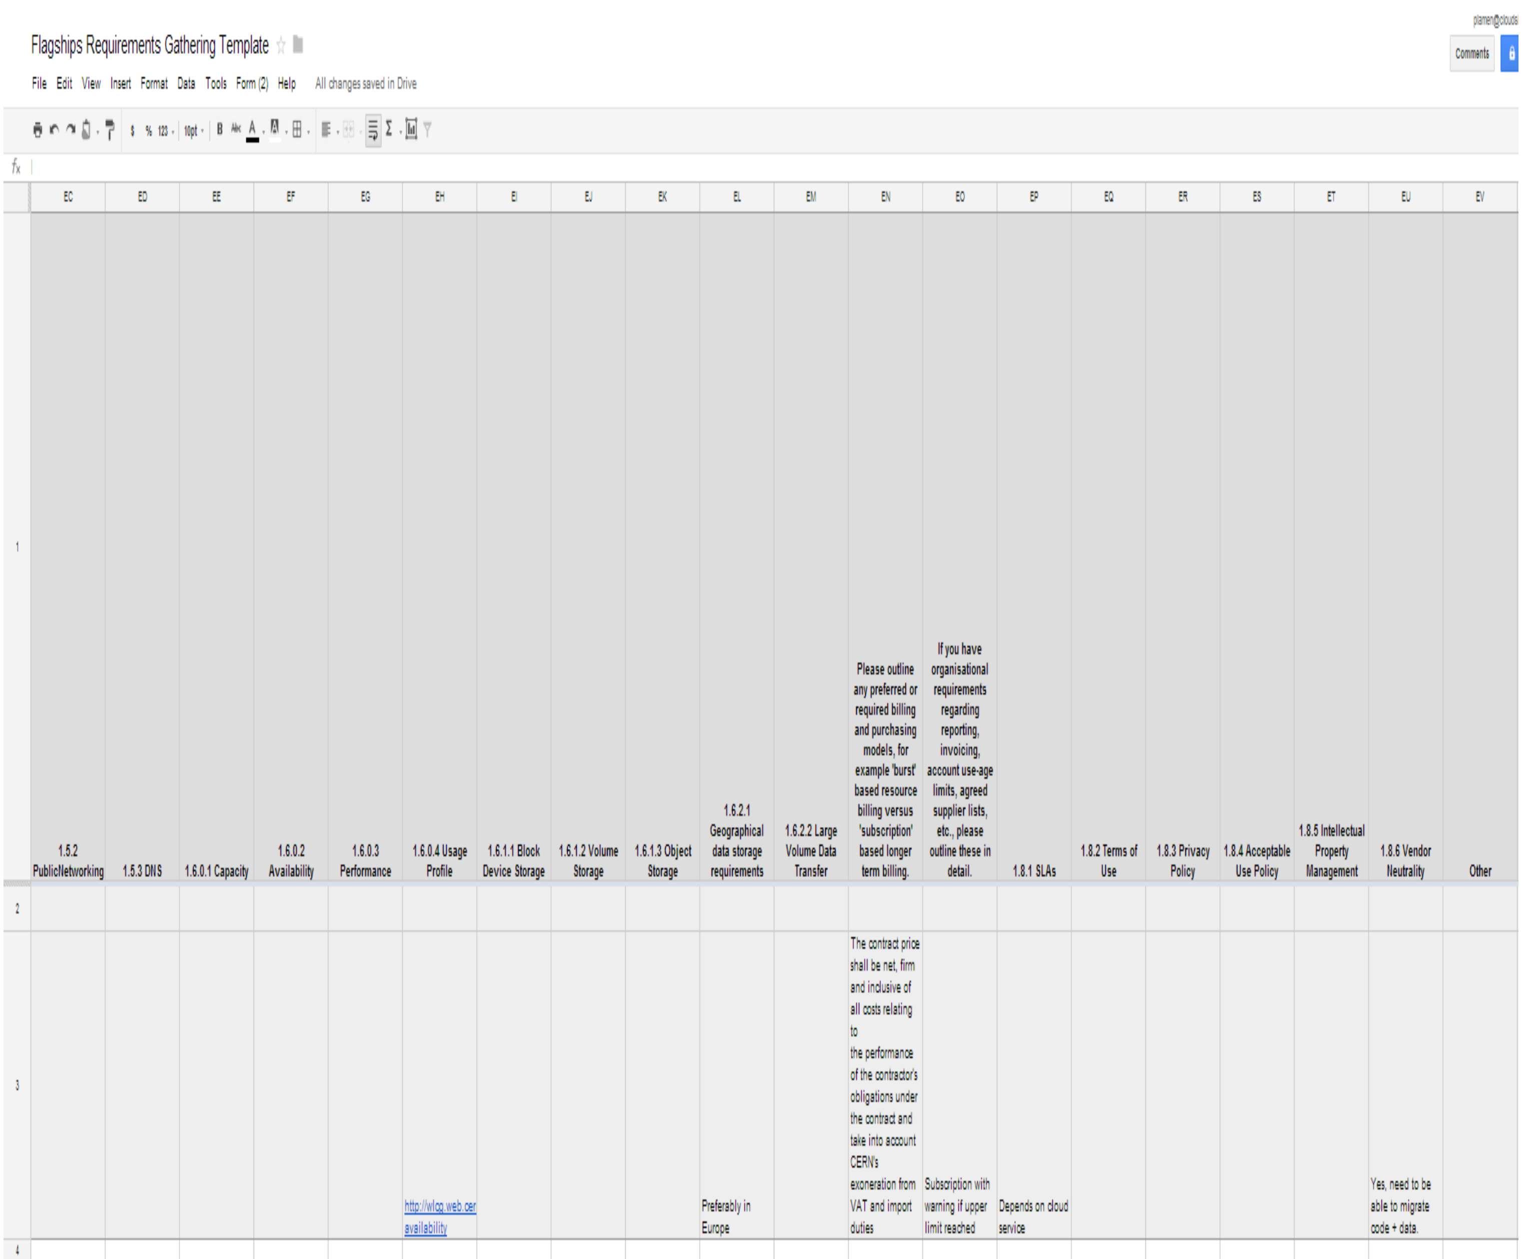

	EC	ED	EE	EF	EG	EH	EI	EJ	EK	EL	EM	EN	EO	EP	EQ	ER	ES	ET	EU	EV
1	1.5.2 PublicNetworking	1.5.3 DNS	1.6.0.1 Capacity	1.6.0.2 Availability	1.6.0.3 Performance	1.6.0.4 Usage Profile	1.6.1.1 Block Device Storage	1.6.1.2 Volume Storage	1.6.1.3 Object Storage	1.6.2.1 Geographical data storage requirements	1.6.2.2 Large Volume Data Transfer	Please outline any preferred or required billing and purchasing models, for example 'burst' based resource billing versus 'subscription' based longer term billing.	If you have organisational requirements regarding reporting, invoicing, account use-age limits, agreed supplier lists, etc., please outline these in detail.	1.8.1 SLAs	1.8.2 Terms of Use	1.8.3 Privacy Policy	1.8.4 Acceptable Use Policy	1.8.5 Intellectual Property Management	1.8.6 Vendor Neutrality	Other
2																				
3						http://wlcg.web.cer availability				Preferably in Europe		The contract price shall be net, firm and inclusive of all costs relating to the performance of the contractor's obligations under the contract and take into account CERN's exoneration from VAT and import duties	Subscription with warning if upper limit reached	Depends on cloud service					Yes, need to be able to migrate code + data.	
4																				

Figure 4: Sample screen captures of the aggregated data - 3

Our chosen tool also provides wider analysis of data entered allowing the interesting identification of trends which will provide further feedback in time, which in turn can be used to make the process more user-friendly and to adjust requirements when necessary.

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

Proposer Long-term Objectives:

Our long term objective is to make de novo assembly and genome annotation widely available to the scientific community maximize the scientific and social benefits we believe this service can bring. We also propose to expand into comparative a
genomes given that the confidentiality and legal issues of dealing with potential clinical data can be satisfied. We believe the provision of this service to human genome resequencing has great potential for benefitting the medical community. The cloud c
power and scalability to allow us to make this service ...

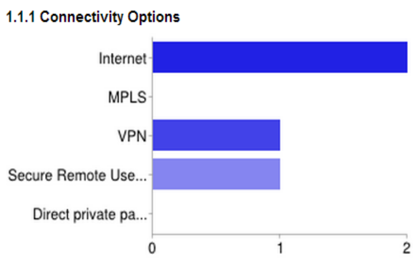

Internet	2
MPLS	0
VPN	1
Secure Remote User Access	1
Direct private patch to local switchж	0

People may select more than one checkbox, so percentages may add up to more than 100%.

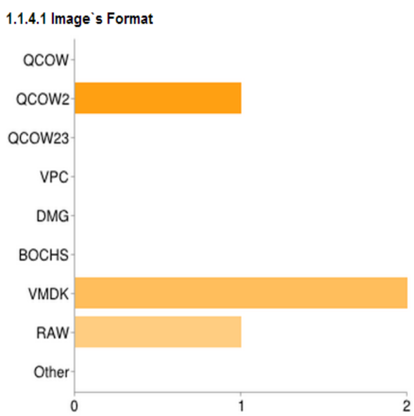

QCOW	0
QCOW2	1
QCOW23	0
VPC	0
DMG	0
BOCHS	0
VMDK	2
RAW	1
Other	0

People may select more than one checkbox, so percentages may add up to more than 100%.

Figure 5: Sample of Graphical representation of the data gathered

3. Framework in Context of the Final User and Service Requirements (D3.3)

The newly created repository will be used to collate details of the requirements and capabilities of all new prospective supply and demand side institutions applying for membership to the Helix Nebula consortium.

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

Prospective new members will be asked to complete the relevant form so all pertinent information can be collected upfront, allowing a professional and informed review of each prospective member.

The back-end repository is automatically appended with the details submitted by any new member on the completed online forms. This allows a simple organised central store of relevant member capabilities and requirements.

It is anticipated that further refinements of these requirement frameworks will be necessary as work undertaken from other work packages in the Helix Nebula consortium matures. The tools deployed are flexible and have been designed to easily incorporate additional requirements as they come to fruition.

In particular, areas surrounding service architectures and business & delivery models remain work in progress and are subject to intensive discussion and review. It is anticipated that the requirements framework will be significantly expanded to include more detailed information in relation to these critical areas as the project progresses. These additional areas will be incorporated directly into the existing requirements gathering framework. Capturing these fine grain requirements will be critical in guiding the supply side in relation to these less technical deployment areas and in subsequently guiding new flagship deployments to concur with supply side cloud offerings.

The central repository created from these frameworks will provide a critical role as an input for D3.3 in allowing the organisation and subsequent analysis of new member applications.

We intend to further improve the online tool to gradually allow submission of data over time by avoiding the session state. This will allow collaboration amongst colleagues within demand side flagship institutions and the gathering of data over an extended time frame rather than the current single consolidated submission. Additionally, we intend to introduce advanced features for both demand and supply side parties for the online framework including supporting version control and knowledge aggregation as aspects of the online tool and back-end management databases and spreadsheets.

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