# Towards the European Open Science Cloud

## Executive summary

The objective of this paper is to propose the establishment of the European Open Science Cloud that will enable *digital science* by introducing *IT as a Service* to the public research sector in Europe. The rationale calls for a hybrid model that brings together public research organisations and e-Infrastructures with commercial suppliers to build a common platform offering a range of services to Europe's research communities. The exploitation platform will make use of, and cooperate with, existing European e-infrastructures by jointly offering integrated services to the end-user. This hybrid public-commercial cloud model represents a significant change from the status-quo and will bring benefits for the stakeholders: end-users, research organisations, service providers (public and commercial) and funding agencies. A time-line for a pilot and its subsequent expansion together with a funding model engaging all stakeholder groups is described.

## Contents

Introduction ..... 2
Why should Europe develop its own cloud for scientific data ..... 3
What would it take from a technical point of view to set up such a cloud ..... 6
What would be the range of services that could be offered ..... 7
What would be the time frame for development ..... 9
What would be the costs ..... 10
Additional Information ..... 12

Prepared by Bob Jones (IT department) on behalf of CERN
5 March 2015

# Introduction

Over the past decade Europe has developed world-leading expertise in building and operating very large scale Research Infrastructures (RIs) supporting unprecedented scales of international collaboration in science, both within and across disciplines. As these RIs ramp-up their operation, we enter a period of explosive data growth. The foundations for handling the "Data Tsunami" or "Big Data" have been laid in the last decade as we have moved from simple commodity computing ("Farms"), to commodity distributed computing ("Grid") and then commodity computing services ("Cloud"). These have prepared the ground for handling the large amounts of data being produced today. The era of "Data Intensive Science" has begun.

Cloud computing has the potential to reduce IT expenditure while at the same time improving the scope for flexible, high-quality new services. However, the adoption of cloud computing services by public research organisations is inhibited by many barriers related to procurement, trustworthiness, technical standards and legal terms of reference, risk of lock-in etc. The overall challenge is to overcome these barriers in order to boost productivity by stimulating the preparedness for wide adoption of competitive, secure, reliable and integrated cloud computing services.

In order to meet the growing demand for cloud services, public research organisations have the choice between purchasing and operating in-house IT equipment which requires capital investment and operations expertise or procuring external cloud services from providers on a pay-per-usage model. Procuring external cloud services implies infrastructure is no longer 'institutionalised' and the cost of cloud services can be found on the operations budget, calculated on a pay-per-usage basis, rather than the capital budget.

CERN, EMBL and ESA have, in the context of the Helix Nebula initiative[1], been performing deployment tests with a number of commercial providers. Simulation jobs from the ATLAS experiment were used as one test application and showed the use of commercial *Infrastructure as a Service* (IaaS) is technically feasible for simulation workloads. In December 2014 a price enquiry by CERN requesting 2000 concurrent Virtual Machines (VMs) for a period of 6 weeks confirmed interesting price points which meant the difference when compared to the cost of in-house provisioning was greatly reduced.

Looking at the underlying trends in the cloud services market, it is our assessment that the question is no longer *if* commercial cloud services will be become financially attractive but rather a question of *when*.

This paper answers key questions about the European Open Science Cloud by expanding on extracts from documents published by the EIROforum IT Working Group[2].

[1] http://www.helix-nebula.eu/

[2] e-Infrastructure for the 21st Century http://zenodo.org/record/7592 and http://zenodo.org/record/13148





## Why should Europe develop its own cloud for scientific data

Looking forward over the coming 10-15 years, there are exciting challenges ahead to capture, manage, and process the vast amounts of data likely to be generated, not only by the established fundamental research domains but in a growing range of scientific disciplines, large and small. This new frontier in computation will be defined by the needs of data-driven science, simulation, modelling and statistical analysis in areas from climate change to life sciences, art and linguistics. All will see incredible growth and accelerated breakthroughs due to unprecedented access to data and the computational ability to process it.

Europe must preserve its intellectual capital and provide the opportunity for it to be nurtured, developed and to grow. Major advances in technology that have taken the world by storm, from Linux to the World Wide Web, have often been conceived in Europe but ultimately exploited elsewhere. This is a loss to Europe, in terms of skills, employment and business. Jeremy Rifkin, in his recent book entitled *The zero marginal cost society*[3], highlights the need to link public funded research with the commercial sector for the benefit of society:

> *"The exponential curve in computing that helped bring the marginal cost of producing and sending information to near zero was primarily driven by global companies. On the other hand, recall that the Internet was invented by government scientists and university academics, and the World Wide Web was the creation of a computer scientist interested in promoting the Commons. GPS, touchscreen displays, and voice activated personal assistants (e.g., Siri) – the key technologies that make the celebrated iPhone "smart" – were the result of government funded research. Linux, Wikipedia, and MOOCs are inspirations that come largely from the social economy while Facebook and Twitter are commercial ventures whose success depends on building social Commons in the hopes of reaping financial gain."*

Helix Nebula has demonstrated the potential of a hybrid model bringing together service providers, research organisations, data providers and publicly funded e-infrastructures to support and transform publicly funded research into data driven knowledge that is valuable to the wider research community and downstream industries. The stakeholders have federated their efforts and resources permitting a sub-set of the service providers to develop a first product called HNX[4] that is being marketed in a range of business sectors. Progress towards the objectives[5] of Helix Nebula was summarized in the recent report, *Strategic Plan for a Scientific Cloud Computing Infrastructure for Europe: Three years on*[6] which highlights that:

> *"The Helix Nebula consortium's enthusiasm and motivation have facilitated the shift from a 'cloud active' to 'a cloud productive' phase to seize new opportunities and challenges. The partnership continues its commitment to offer an open, secure and trusted Cloud Computing Infrastructure for European science, businesses and society, and to become a leader in a highly competitive global market."*

A key value of Helix Nebula is that of a forum for the *supply-side* and the *demand-side* where issues of common interest (such as procurement models, contractual frameworks, service platforms etc.) can be addressed. Demonstrable progress has been made on all the stated goals of the original strategic plan but the lack of a production environment offering a comprehensive set of services is hampering innovation.

---

[3] http://www.thezeromarginalcostsociety.com/

[4] http://hnx.helix-nebula.eu/

[5] Strategic Plan for a Scientific Cloud Computing infrastructure for Europe, https://cds.cern.ch/record/1374172/files/CERN-OPEN-2011-036.pdf

[6] https://cds.cern.ch/record/1706226/files/HelixNebula-D9_2.pdf





In parallel the European Cloud Partnership (ECP[7]) is working to bring together industry and the public sector to establish a Digital Single Market for cloud computing in Europe. It is assembling best practises for the contractual aspects of delivering cloud services including services level agreements, certification requirements, a code of conduct, and terms and conditions that respect European legislation. The ECP publication "*Establishing a Trusted Cloud Europe*" highlights the important role the public research sector can play in the development of Europe's cloud market.

We must take the step to make unprecedented scales of IT resources available to the next generation of emerging scientists, researchers and entrepreneurs, nurturing them from education to start-up activities and then sustainable businesses or research communities. As well as serving the direct needs of computing, the opportunity to innovate and explore new technologies is essential. The correct environment for innovation will allow many ideas to be tested, and explored, which will lead to the truly unexpected and ground-breaking discoveries and inventions that will shape this century. As such, the European Open Science Cloud is fully in line with the European Cloud Computing Strategy[8], adopted by the EC in September 2012, which aims at accelerating Cloud up-take across Europe. In this strategy, the EC calls upon Member States to embrace the potential of cloud computing.

The ESFRI RIs now entering production are securing long-term commitments from national funding agencies. As part of their computing models, such RIs are focussing on individual centres that provide tailored services. Such centres provide well-defined services for specific user communities' and frequently benefit from multiple funding sources and hence have a defined approach for sustainability. An established production example of this model is the WLCG *tier-1* centres[9] providing grid services to the LHC experiments. In the life science domain ELIXIR has identified existing *nodes*[10] in member states that are supported by national funding to offer bioinformatics services and participate in an international data consortium[11]. In the humanities domain CLARIN has a set of certified centres[12] that provide infrastructure services[13] and in the earth sciences domain LifeWatch has an IT Research and Innovation Centre in Amsterdam.

Put together these centres could form a network of interoperating facilities that would support workflows addressing the full research life-cycle. The underlying high-performance network, GEANT, can bring together these facilities and scientific instruments, which are big-data factories, with additional data providers hosting important datasets around the world. PRACE super-computers connected to the network could have a dual role: like scientific instruments they can be seen as big-data factories but they can also offer specialised capability-style data analysis services for datasets to be transferred over the network from the facilities or external data providers. Such data could be found, exchanged and replicated via services developed by EUDAT. Further capacity and services from the rapidly evolving commercial cloud services market can complement the underlying public infrastructure. Freely available compute resources may be available via the volunteer computing infrastructures managed via the International Desktop Grid Federation[14]. All the elements of the European Open Science Cloud must share a federated identity management system enabling single-sign-on access for users to any resource for which they have authorisation. A common digital object identifier service

[7] http://ec.europa.eu/digital-agenda/en/european-cloud-partnership
[8] http://ec.europa.eu/digital-agenda/en/european-cloud-computing-strategy
[9] http://home.web.cern.ch/about/computing/grid-system-tiers
[10] http://www.elixir-europe.org/about/elixir-nodes
[11] http://datafairport.org/
[12] http://www.clarin.eu/node/3822
[13] http://www.clarin.eu/sites/default/files/CE-2012-0037-centre-types-v07_0.pdf
[14] http://desktopgridfederation.org/

would promote the citation and re-use of any publication, dataset or software suite in the European Open Science Cloud.

The situation described above highlights the lack of a single focal point to bring together Europe's developments, federating capacity, policy and procurement activities into a consistent whole that can remove fragmentation and ensure the results of each activity are fully exploited. A closer alignment of these industrial and public (regional, national, and European) strategies would increase the rate of innovation.

The European Open Science Cloud represents, as stated in The Data Harvest Report[15], "*the future of science - a virtual science library spanning the world, where the vital catalyst is the ability to share data, in huge volumes, over vast distances, across disciplines and institutions. And then analyse, re-interpret, re-use and re-think it. And the future also belongs to small businesses creating new data products*."

[15] https://rd-alliance.org/data-harvest-report-sharing-data-knowledge-jobs-and-growth.html





## What would it take from a technical point of view to set up such a cloud

In considering the way forward, it is important that we foresee a European Open Science Cloud that supports a broad range of Europe's research communities including the "*long tail of science*" conducted by researchers that do not have access to significant in-house computing resources and skills. Consequently this should not be thought of as a one-size-fits-all solution. Rather an extended but coherent set of services and tools, organised into layers as described in more detail below, which must be available to meet the specific needs of each community and integrated via interfaces defined by open standards. This will allow the research communities to select the services or tools that they require, and only those, without additional complexity. One of the lessons from the grid experience is that unnecessary complexity should be avoided in the infrastructure layers. This common platform should also be able to act as the incubator for new businesses and scientific activities. It is essential that European industry engages with the scientific community in building and providing such services as part of a hybrid cloud model, but it is also important that the user community have a strong voice in the governance. In short, the European Open Science Cloud would become the incarnation of what the e-Infrastructure Reflection Group (e-IRG) refers to as the *single e-Infrastructure Commons* in its 2012 Roadmap paper[16].

Implementations for the majority of the infrastructure services of the European Open Science Cloud already exist at varying levels of maturity. The key challenges are integrating frequently changing technologies, managing the complexity and identifying the optimal organisational and financial models. Researchers must be convinced that they will not lose control of their precious data. The public research organisations can provide such guarantees. They can rapidly expand the available capacity by making use of commercial service providers offering commodity compute and data services as part of the hybrid cloud model. By keeping a "master copy" of the research data, the public research organisations can insulate the researchers from changes in service provider and technology.

The funding agencies that support the work of the researchers must also be convinced that the European Open Science Cloud offers financial advantages compared to building their own dedicated IT systems. The hybrid cloud model is perceived as both a threat and an opportunity by service providers. Commercial service providers see the market potential of the public research sector but fear their investments will be undermined by publicly funded service providers offering similar services at no cost to the user. Publicly funded service providers fear that commercial services will replace their existing services operated on their in-house infrastructure. Consequently, we foresee that the overall management and control of the European Open Science Cloud would be the task of the public research organisations so that a relationship of trust can be established with the research communities while ensuring non-discrimination, transparency and equal treatment for all commercial suppliers.

Through the work of initiatives, such as Helix Nebula, it has become clear that it is essential to separate the roles of end-user (the researcher making use of the services) and customer (the organisation sponsoring the consumption of the services by the end-user) and ensure that services, both commercial and publicly funded, are offered as free at the point of use. Helix Nebula has also shown that the public research sector can generally predict its medium-term needs for cloud services. Such foresight means the sector can procure large quantities of commodity commercial cloud services on longer-term contracts at significantly better prices than available for on-demand purchases. Joint procurement in the public research sector would permit small scale users, typical of the long-tail of science, to benefit from these economies of scale and a rich set of cloud services.

[16] http://e-irg.eu/documents/10920/12353/e-irg_roadmap_2012-final.pdf





## What would be the range of services that could be offered

We propose creating a common platform that builds on the experience of the last decade and is flexible enough to adapt to technological and service innovations. Such a platform must provide the underlying layers of common services and yet be adaptable to the very different and evolving needs of the research communities. The proposal has 3 distinct layers of services:

1. European and international networks (GEANT); services for identity management and federation across all European research and education institutions and integrated with other regions of the world (eduGAIN);
2. A group of facilities to provide cloud and data services of general and widespread usage. Linking data services with cloud computing capacity to offer data analysis platforms will present users with a comprehensive environment supporting the full lifecycle of science workflows.
3. Software services and tools to provide value-added capabilities to the research communities, in a managed repository:
    a. The tools to provide those research communities that have access to large sets of resources the ability to federate and integrate those resources and to operate them for their community, potentially sharing with other communities;
    b. Tools to help build applications: e.g. tools to manage data, storage, workflows, visualisation and analysis libraries, etc.
    c. Tools and services to allow researchers to integrate everyday activities with the European Open Science Cloud: collaborative tools and services; office automation, negotiated licensing agreements etc.;
    d. Tools to help research communities engage the general public as citizen scientists.

This portfolio of services, using as a starting point those listed by e-IRG and the High Level Expert Group on Scientific Data[17], will be made available under a set of terms and conditions that are compliant with European jurisdiction and legislation with service definitions implementing recognised policies for trust, security and privacy notably for data protection. Training and consultancy services can also be made available to users of the European Open Science Cloud via competence centres.

As an example, the following is the portfolio of services already provided by CERN to its users:

- A virtual multi-tenant compute environment to provision and manage networks of virtual machines on-demand[18].
- A 'dropbox' style service for secure file sharing over the internet[19] based on the ownCloud[20] product.
- A point-to-point reliable, automated file transfer service for bulk data transfers[21].
- The zenodo OpenAIRE compliant open access repository[22] for publications and supporting data and software allowing users to create and control their own digital libraries. Persistent digital object identifiers will be assigned to all publically available uploads so as to make them citable and permit the creation of a digital data continuum spanning from experimental data through to publications including links to commercial publishers.

[17] http://cordis.europa.eu/fp7/ict/e-infrastructure/docs/hlg-sdi-report.pdf
[18] http://clouddocs.web.cern.ch/clouddocs/
[19] http://cernbox.web.cern.ch/
[20] http://owncloud.com/
[21] http://information-technology.web.cern.ch/services/file-transfer
[22] http://www.zenodo.org/features

- A long-term archiving service[23].
- Integrated Digital Conferencing[24] tools allowing users to manage their conferences, workshops and meetings. The digital conferencing tools already include commercial services such as Vidyo[25].
- Training services: Online training material[26] will be made available and re-enforced with advanced training via brief residential programmes and secondments, such as CERN's openlab summer student programme[27], which has recently been expanded with the engagement of EIROforum organisations and more companies.

The services are accessible in a single sign-on (SSO) manner supported by a federated identity management system and initial integration with eduGAIN has been performed[28]. A Service Portal, based on the service-now[29] commercial product, provides a common support interface for the users. These services are implemented in a cloud environment using the OpenStack[30] open source software suite. OpenStack has been selected as the framework for the cloud services because it has emerged as a clear leader in the vibrant cloud management market. CERN has extensively tested the software and has adopted it for the management of its production resources at its data centres in Geneva and Budapest. OpenStack is distributed under the Apache 2 license and supported by a growing consortium of more than 450 private companies and public organisations with a well-defined and transparent governance structure representing a global community of more than 17,000 people across 87 countries. More than 1400 developers are contributing to the software. All of these attributes means OpenStack does not require public funds to be maintained and currently offers one of the best environments for innovation.

A European Open Science Cloud offering such a comprehensive range of layered services would become a critical ICT infrastructure for the European Research Area and would need to be protected by identifying vulnerabilities and ensuring an operational security plan is in place to minimize the detrimental effects of disruptions.

[23] http://castor.web.cern.ch/
[24] http://indico-software.org/
[25] http://information-technology.web.cern.ch/services/fe/vidyo
[26] http://indico.cern.ch/categoryDisplay.py?categId=88
[27] http://openlab.web.cern.ch/news/cern-openlab-summer-student-programme-invites-2013-applications
[28] Deliverable D9.4 (DS5.5.1) Towards Horizon 2020 - The Enabling Users Experience, GN3PLUS14-557-118, http://www.geant.net/Resources/Deliverables/Documents/
[29] http://service-now.com/
[30] http://www.openstack.org/

## What would be the time frame for development

A pilot for the European Open Science Cloud could be initiated in 2015 and delivered in 2018. The pilot would bring together the stakeholders:

- Research Infrastructures *(ESFRI and the future ATTRACT collaboration[31]);*
- Research organisations (*EIROforum members etc.*)
- European e-Infrastructures *(GEANT, EGI, PRACE, EUDAT, OpenAIRE);*
- Commercial cloud service providers *(Helix Nebula, etc.);*
- End-users including the *long-tail of science.*

These stakeholders would work together to deliver the pilot corresponding to a 10% scale implementation of the European Open Science Cloud. The pilot activities would include:

- Produce a technical architecture for the hybrid cloud that can build on the existing public and commercial developments;
- Agree a security model that is compatible with EU data protection legislation;
- Assemble and deploy a 10% scale prototype;
- Verify the business model to ensure it can be sustained beyond the pilot phase;
- Establish an inclusive governance structure where all the stakeholders are represented and avoids a monopoly of any supplier or research group;
- Develop a roadmap for full-scale implementation.

Expanding this pilot to provide more capacity, functionality and serve more RIs and communities would take an additional 3 years. The expansion from the pilot to the complete European Open Science Cloud would include engaging more RIs/communities and encouraging more European cloud service providers to participate. The Research Data Alliance[32] would offer a convenient means for engaging more research communities and expanding beyond Europe. The total capacity of the hybrid cloud would need to be significantly increased and functionality expanded by offering higher-level added-value services (Platform as a Service, Software as a Service, etc.) To ensure the European Open Science Cloud can act as a vehicle for innovation across the research and business sectors it will be necessary to expand coverage to the broader public sector (including eGovernment) and the private sector (i.e. industries that use research data including energy, pharmaceutical, insurance etc.) Interaction with other regions of the world (Africa, Asia, Latin America, USA etc.) would also need to be addressed to satisfy the collaborative nature of global research communities.

The Helix Nebula initiative has explored and documented[33] many aspects of a suitable governance structure and identified several possible approaches. Having an organisation that could oversee the procurement process, certify and enrol service providers as well as handle the contractual arrangements between suppliers and customers with centralised billing would simplify the operation and expansion of the European Open Science Cloud.

[31] http://www.attract-eu.org/

[32] https://rd-alliance.org/

[33] D8.1 A study of governance models for public-private cloud partnerships, http://www.helix-nebula.eu/sites/default/files/D8_1%20A%20study%20of%20governance%20models%20for%20public-private%20cloud%20partnerships.pdf

## What would be the costs

The European Open Science Cloud builds on the investments already made in Europe for the publicly funded e-infrastructures and commercial cloud services. Through Horizon 2020, the EC and national funding agencies have recently confirmed their commitments to GEANT, AARC (Authentication and Authorisation for Research and Collaboration), EGI, OpenAIRE, EUDAT and PRACE. In particular, EGI's commitment to an open science commons forms the basis of the recently approved EGI-Engage project. In order to ensure full synergies, DG CNECT foresees that e-infrastructure projects will be grouped into clusters of related projects. This new phase of funding for the cluster of e-infrastructure projects offers the EC a window of opportunity and a means to focus on establishing the European Open Science Cloud.

In addition to a focussed cluster of e-infrastructure projects, what is missing is the engagement of the commercial cloud service providers, the RIs and communities as well as a vehicle for bringing all these stakeholders together. There must be an incentive to encourage the research community to adopt cloud services, gain trust in service providers and encourage commercial cloud providers to consider the public research sector as a potential market.

Public procurement of cloud services is promoted as a means of encouraging innovation and reducing the time to market for new products and services. Forming a network of public research organisations that can procure cloud services will attract the interest of service suppliers as well as national and regional funding agencies. The majority of this procurement funding will be directed to service providers and the approach has the advantage of permitting the procuring organisations to choose which services and providers receive these funds and thus represent a change to the established funding model for public sector IT services.

CERN has been working with the EC, in the context of a Horizon 2020 support action[34], to explore co-funding for pre-commercial procurement of innovative cloud services. The EC has published details of suitable funding calls with a deadline for submission in April 2015. Such funding calls offer an opportunity for public research organisations to gain financial support from the EC for the joint procurement of state-of-the-art cloud services. The main steps involved include:

- A group of research organisations commits to participate in a joint procurement activity. The target total procurement sum needed is in excess of 5M€;
- CERN would act as the coordinator and submit a Pre-commercial procurement (PCP) proposal to the funding call by April 2015;
- Assuming the proposal is successful, the project would start in January 2016 to develop a common set of specifications (technical and contractual) in conjunction with the publicly funded e-Infrastructures and industry suitable for integration into the European Open Science Cloud pilot and address the needs of data intensive research;
- An intermediary with experience in commercial utility markets would act as lead procurer to aggregate the demands of the research organisations and attain the best market prices from multiple suppliers;
- The EC co-funding would reimburse the procurements made by the procurers (maximum 70%) and cover the costs for project management, communication and service integration activities.

The timeline and phases of the PCP project to establish the European Open Science Cloud pilot are shown in the figure below.

[34] http://www.picse.eu/

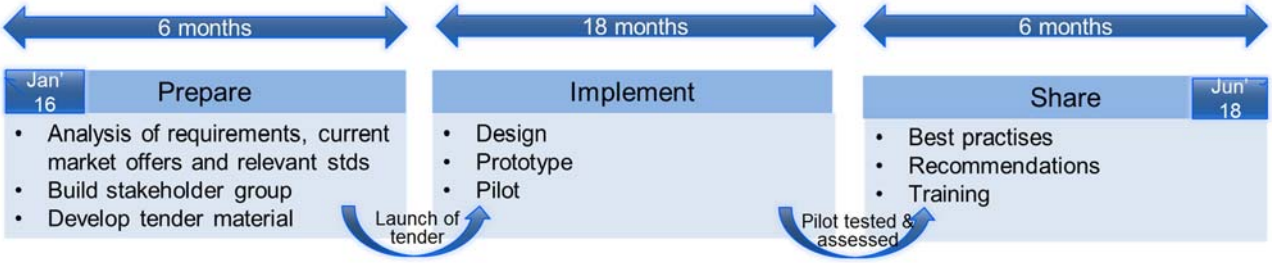


In order to expand from the pilot to the full-scale European Open Science Cloud, the PCP funding model used in the pilot phase could be repeated for the higher-level services. Additional capacity for lower-level services that have been successfully tested during the pilot phase could be supported by Public Procurement of Innovative services (PPI) actions. It is estimated that a total of 50M€ of EC co-funding for the combined PCP/PPI activities would be required during the period 2017-2020. The expansion of the pilot to the full-scale European Open Science Cloud would also require continued investment in the public e-infrastructures during the same period to increase their capacity (e.g. networking and storage - notably for long-term data preservation where many issues remain to be addressed). Computing resources at the research organisations will need to be increased so that they can continue to guarantee the safe keeping of a copy of the rapidly growing body of research data. These actions would need to be supplemented by investment in strategic and application software in order to build and share expertise in ensuring that applications are capable of exploiting evolving computing architectures.

As stated earlier, an objective of the pilot PCP project would be to verify the business model to ensure it can be sustained beyond the pilot phase. The *Cloud for science and public authorities*[35] report commissioned in 2013 by DG CNECT includes an analysis of funding models for pan-European cloud e-infrastructures and the following policy recommendations:

> *"The EC should promote and sustain the spontaneous movement towards the integration and federation of clouds at the EU level, avoiding the risk to develop top-down infrastructures totally dependent on public funding and unable to adapt to the multidimensional characteristics of demand. This will allow supporting the provision of cloud services for science and research across Europe, filling the gap between the actual offering of computing resources and the emerging, pent-up demand.*
>
> *The EC should however make sure that the provision of cloud services for science and research covers the whole of the EU27, avoiding any potential "digital divide" between the large/small countries and between Big Science/Small and Medium Science (SMS) projects, insuring equal access to clouds for all researchers."*

The research organisations in the procurement group fund computing capacity for a well-defined set of researchers. The business model supporting the European Open Science Cloud will need to include mechanisms so that the costs of serving a larger and more diverse set of user groups, including SMS projects, can be shared among the stakeholders. For example, the exploitation of scientific data by downstream industries could rely on services from commercial service providers charged to the companies consuming the data. Policy changes are also likely to be needed to permit many researchers in the public sector to buy computing capacity under the operating expense budget. Consequently it is expected that a revision of the

[35] Cloud for science and public authorities, Ref. Ares(2013)2703248 - 18/07/2013, ISBN:978-92-79-31142-0, DOI:10.2759/25446; http://cordis.europa.eu/fp7/ict/e-infrastructure/docs/final-report-clouds-study.pdf

procurement processes of research organisations will be necessary and such changes can build on the results of the on-going Procurement Innovation for Cloud Services in Europe (PICSE[36]) project.

## Additional Information

This appendix provides some additional information concerning the implementation of the PCP pilot project.

**What is the role of the research organisations?**

Research organisations have multiple roles within the PCP pilot project. Research organisations will be members of the buyers group participating in the joint procurement. These research organisations will also have a role in deploying and integrating the procured services into the European Open Science Cloud as well as engaging their user communities. The research organisations need to extend the IT services they make available to their user communities and have the organisational capacity, links to e-infrastructures and relevant skill-sets to take on the role of *champions of change* to establish the European Open Science Cloud. In some cases a buyer may be a national funding agency that decides to delegate a number of its responsibilities in the project to a research organisation. These responsibilities may include contributing to the preparation of the tendering material and handling the deployment of the procured services.

**Can the PCP project be used to procure hardware, cloud services or both?**

From discussions with the EC, we understand that the PCP project can be used to procure **innovative** cloud services and hardware that are intended to become part of joint European IT systems. The European Open Science Cloud paper highlights that a key role for the research organisations is to guarantee the safe storage of a copy of the data. Procuring **prototype** hardware that could be deployed at the data centres operated by the buyers and focussed on storing the data, i.e. tertiary storage, would be consistent with this role. The innovation aspects could be addressed by procuring next-generation solutions that can be collectively referred to as a 'cold storage' technology[37]. The technology choices would be part of the PCP pilot project tender process. The procurement of production-ready, operational hardware is **not** allowed in a PCP project.

In terms of cloud services, the focus would be on commercial IaaS since these are not yet in production usage in the public research sector. The IaaS services should be capable of supporting not only simulation (i.e. CPU-intensive) but also analysis (i.e. I/O intensive) workloads.

The relative weighting between prototype hardware and cloud services procurement needs to reflect the vision of the overall hybrid cloud. We estimate that allocating no more than 20% of the procurement funds for prototype storage hardware and at least 80% for cloud services would be appropriate in order to both stimulate the adoption of cloud resources whilst ensuring a safety net for the data.

Practically, the hardware and cloud services would appear as separate lots within the same joint tender. Suppliers have free choice about which lot(s) they wish to bid for.

[36] http://www.picse.eu/

[37] http://www.forbes.com/sites/tomcoughlin/2014/02/03/storage-observations-from-the-open-compute-project/

**What kind of expertise should the people have that are funded by the personnel money available through the project?**

The majority of EC funding in the PCP project will be used to partially reimburse the procurements made via the joint procurement process. There is limited EC funding available for *coordination activities* including the management of the project, engaging external experts that assist with the procurement process and to lead communication/outreach/dissemination tasks. EC funds cannot be used to perform R&D and technical activities in the public organisations that will be beneficiaries of the project. We propose that the EC funding for the coordination activities is assigned as follows:

- Project director, administrative assistant, procurement officer (CERN)
- PCP legal expert (TRENTO RISE)
- Cloud services legal expert (Cloud Legal Research Project)
- Cloud services broker (Strategic Blue)
- Communications (TRUST-IT)
- Engagement with public e-infrastructures and non-HEP research communities (EGI.eu)
- Organisation of events and publication of material.

The profiles of the experts are provide below:

TRENTO RISE

TRENTE RISE is a Not-for-profit Association founded under Italian law by FBK and UNITN. TRENTO RISE is a pole of excellence for research, innovation and top-level training in the field of ICT and promotes partnerships between the research and business sectors. TRENTO RISE has extensive experience with organising public sector procurements[38] and has participated as a legal expert for procurement processes in a number of FP7 PCP projects including Cloud4Europe and NYMPHA-MD and has delivered webinars about PCP[39].

Cloud Legal Research Project

The Cloud Legal Research Project (CLP) aims to address the various aspects of the legal and regulatory status of cloud computing that will be essential to its successful development. The CLP research findings demonstrate thought leadership in several complex and difficult areas of law and regulation that are of vital importance to governments and businesses globally. The multi-year project was started in October 2009 by members of the Centre for Commercial Law Studies (CCLS) at Queen Mary, University of London with financial support from the Microsoft Corporation. CCLS has extensive relevant expertise and experience and a leading reputation in all aspects of computer and communications law and regulation. Prof. Christopher Millard of CCLS has published an acclaimed book on Cloud Computing Law[40] and Prof. Kuan Hon has delivered webinars about cloud service procurement in the public sector[41].

[38] http://www.trentorise.eu/pre-commercial-procurements
[39] https://indico.cern.ch/event/372086/
[40] http://www.law.qmul.ac.uk/staff/millard.html
[41] http://indico.cern.ch/event/306750/

Strategic Blue

Founded in 2009 by former Morgan Stanley commodity traders, Strategic Blue was originally a consultancy dedicated to leading Morgan Stanley's research into the viability of trading cloud computing as a commodity. Strategic Blue became independent of Morgan Stanley in 2010, and is funded by a number of high net worth individuals from banking, trading, hedge funds, systems integration and oil industries. In 2013, Strategic Blue was selected to join the elite startup accelerator programme run by Techstars.

Dr James Mitchell, CEO and co-founder of Strategic Blue is a Cambridge physics graduate with a doctorate in DNA Nanotechnology from Oxford University. He was one of the most successful deal originators / structurers on the commodities trading desk at Morgan Stanley, pioneering numerous landmark deals in the electricity markets. Today, he is better known for his work architecting a traded wholesale market in cloud computing, and is regarded as the world's foremost expert in how the cloud computing market will develop. James was originally hired into Morgan Stanley by Strategic Blue's other co-founder, John Woodley, who is credited with being the first person to be hired by a Wall Street bank from a power utility, and as a result is unrivalled in the breadth of his experience in how utility markets develop.

Strategic Blue offers three broad types of services:

- Cloud Options – this core business involves helping clients who find they have a difference between the terms on which they would ideally like to purchase cloud computing, and the terms on which the cloud provider(s) are willing to sell at their lowest prices. Strategic Blue will step into the billing chain between cloud customer and cloud provider(s) and inject the required combination of a) billing services, b) finance and c) risk appetite in order to achieve the optimal financial result for all concerned.
- Cloud Options Pricing Insights – access to proprietary analysis of published historical and current public cloud pricing from vendors such as Amazon Web Services.
- Consulting – access to the team's insight on cloud procurement and pricing best practice, based Strategic Blue's unique background in commodities trading.

TRUST-IT

Trust-IT is a UK SME specialised in analysing and communicating Information and Communication Technologies across Europe and globally. Trust-IT supports in carrying out feasibility and market studies, design reports, sustainability activities and business models.

For the past 13 years Trust-IT has been working very closely with small and large enterprises, government and research to drive forward new ideas, pinpoint market opportunities, create multi-stakeholder partnerships and international networks. In driving home the many benefits of IT, Trust-IT has taken the communication of innovative ICT results to a completely new level of digital innovation while aggregating user communities in various ICT, e-Infrastructure and data infrastructure domains.

It coordinates the recently funded CloudWATCHhub.eu project and is the founder of the Cloudscapeseries.eu workshops. It has also coordinated distributed computing support actions such as SIENA and OGF-Europe, co-authoring the SIENA Roadmap on Distributed Computing Infrastructure for e-Science and Beyond, a 'reference document for standardization efforts' (Neelie Kroes, Vice President of the European Commission). It has gained extensive knowledge on cloud by working closely with international experts, service providers, technology developers and user communities from many different fields. TRUST-IT also has played an active role in pioneering cloud initiatives spanning RESERVOIR, VENUS-C, VISIONCloud, Helix Nebula and the industry-sponsored Cloud4Science initiative and more recently the PICSE project.

EGI.eu

The Stichting European Grid Initiative (also referred to as "EGI.eu") is a not-for-profit foundation established under Dutch law to coordinate and manage the European Grid Infrastructure (EGI) federation on behalf of its members: National Grid Initiatives (NGIs) and European International Research Organisations (EIROs). EGI.eu offers a variety of services to the wider EGI community such overseeing infrastructure operations, coordinating user community support, working with technology providers, representing EGI in collaborative projects, steering strategy and policy development, organising flagship events and publicising the community's news and achievements. EGI.eu and its members provide EGI, a pan-European infrastructure of publicly funded computing, storage and data resources to support excellent research and innovation in Europe. Building on over a decade of investment by national governments and the European Commission, EGI supports more than 22,000 researchers across many scientific fields with a wide range of technical services such as high-throughput data analysis, federated cloud, federated operations and community building.

**Must all buyers commit to the same level of procurement?**

It is not necessary for all buyers to commit to the same level of procurement. A minimum procurement commitment from each buyer will be necessary to justify their participation. A combined procurement commitment across all the buyers in excess of 5 million euros is required to make the PCP pilot project viable.

**How will the procured resources be controlled?**

The decision about the usage of the cloud services and prototype storage hardware procured remains with the buyers.

Each buyer will be allocated procured cloud services and prototype storage hardware corresponding to their relative fraction of the total procurement. As all the foreseen buyers are members of the WLCG collaboration, the primary end-user group will be members of the high-energy physics community. The PCP pilot project offers an opportunity for the buyers to promote how they can serve multiple research communities via the European Open Science Cloud and consequently it is proposed that, as a group, the buyers agree to make a fraction of the cloud services accessible to a range of organised scientific communities as well as the long-tail of science and citizen scientists. This will increase the impact of the PCP pilot and help reduce the digital divide across European member states and beyond. Specifically the EGI-Engage funds 8 competence centres serving these research infrastructures and communities:

- biobanking and medical research (Biobanking and Bimolecular Research Infrastructure - BBMRI),
- art and humanities (Digital Research Infrastructure for the Arts and Humanities – DARIAH),
- life science research (ELIXIR),
- earth sciences (the next generation incoherent scatter radar – EISCAT_3D, the European Plate Observing System – EPOS, and climate modelling),
- biodiversity and ecosystem research (LifeWatch),
- structural biology and brain imaging research (MoBrain supporting WeNMR and Integrating Structural Biology – INSTRUCT

**What is the relationship to other projects and existing e-infrastructures (notably WLCG and the tier-1 centres, INDIGO-DataCloud and EGI)?**

The proposed model is that of a hybrid cloud which will involve the existing European e-infrastructures:

WLCG and the tier-1 centres

The concept of performing joint procurement has been discussed with the WLCG overview board in 2014 and it is expected that all the buyers will be European members of the WLCG collaboration[42]. The PCP pilot project will form a consortium separate from the WLCG collaboration. The PCP pilot project consortium will sign a Horizon 2020 grant agreement with the EC and a consortium agreement among themselves. The consortium agreement will have an annex that determines how the buyers will work together as a group.

Those WLCG tier-1/2 centres operated by buyers in the PCP consortium would be suitable sites to host the procured prototype storage hardware.

The foreseen buyers are members of the WLCG[43] collaboration and would like to have their procurements count towards their pledged resources[44]. It is proposed that majority of the procured cloud services and prototype storage hardware are allocated to WLCG and taken into account in the buyers' pledges with the remaining fraction made available to other research disciplines, the long-tail of science and citizen scientists.

INDIGO-DataCloud

The recently approved Horizon 2020 INDIGO project aims develop and deliver software components allowing execution of applications on Cloud and Grid based infrastructures, as well as on HPC clusters. Both public and private sector cloud service providers are already offering IaaS (Infrastructure as a Service) cloud resources, while INDIGO will focus on the PaaS (Platform as a Service) and SaaS (Software as a Service) levels. Consequently the PCP pilot project (focusing on procuring IaaS level services) and INDIGO (PaaS & SaaS) will be complementary and it will be possible for INDIGO services to be deployed across the commercial IaaS services procured via the PCP pilot project.

EUDAT

The recently funded EUDAT2020 project will further develop the prototype data services (B2SHARE, B2SAFE, B2STAGE, B2FIND, B2DROP) operated by a number of publicly funded centres (referred to as data nodes by EUDAT). These data nodes can be integrated into the European Open Science Cloud as data repositories. In addition, the EUDAT activities to develop data quality assessments for data repositories could provide useful input for the certification aspects of the tendering material to be developed by the PCP pilot project.

GEANT

GÉANT is the pan-European research and education network that interconnects Europe's National Research and Education Networks (NRENs). GÉANT's mission is to deliver world-class network services with the highest levels of operational excellence to research and education communities within Europe and beyond, helping to develop talent and providing opportunities for communication across the divides of resources and distance,

[42] http://wlcg.web.cern.ch/collaboration
[43] http://wlcg.web.cern.ch/
[44] http://wlcg-rebus.cern.ch/apps/pledges/resources/

and so promoting the free, unimpeded movement of scientific data and knowledge. In the context of the Helix Nebula initiative, a number of commercial service providers across Europe have already been connected to the GEANT network so that scientific flagship applications could be executed using commercial cloud services and publicly funded services operated by EGI and in-house resources provided by CERN, ECMWF, EMBL, ESA and PIC.

AARC

The Authentication and Authorisation for Research and Collaboration (AARC) project aims to develop an integrated cross-discipline AAI framework, built on production and existing federated access services (National Identity Federations and eduGAIN), to serve researchers, students and educators. By using the AAI framework users can access different services with the same credentials [Single Sign-On - SSO] and service providers can offer resources in a more controlled and consolidated way. AARC will also pilot SSO access for commercial cloud services for research communities and consider both technical/architectural solutions and legal and policy aspects. The proof of concepts will involve services from the main e-infrastructures in Europe (e.g. EGI, EUDAT, PRACE, eduGAIN), as well as libraries and ESFRI cluster projects (ELIXIR, DARIAH and so on). Service providers will include both academia and industry, and both web-based and non-web based services.





EGI

The recently approved EGI-Engage project includes a task on cross-border procurement of e-Infrastructure services within NA2[45]. The PCP pilot project will offer a realistic implementation of the objectives of this task. Through this task, the PCP pilot project will also be able to engage with more Research Infrastructures as potential procurement groups, notably BBMRI, EPOS and DARIAH.

The EGI Fed Cloud could provide the means by which the procured cloud services are made available to the users. This would require some limited technical work to allow the APEL module supporting EGI Fed Cloud accounting to be used for WLCG accounting but would encourage EGI resource providers to move from batch (grid) to cloud which means that they are more aligned with general purpose provisioning that has lower operating costs. Using EGI to provide end-user access to the procured services would also offer the advantage of providing a ready-made solution for collecting the impact metrics (i.e. number of users, distribution by country & discipline, etc.)

The PCP pilot project tender is open to commercial cloud service providers and public providers that are allowed, by their national laws, to respond to tenders. EGI has already completed surveys of grid sites and identified a small number (approximately 10 sites) that are able and willing to respond to calls for tender. These sites could potentially form a consortium that submits a bid in response to the PCP tender. However, EGI sites that are owned or funded by a buyer in the PCP pilot project are excluded from being part of such a consortium.

OpenAIRE2020

OpenAIRE2020 will expand its existing scholarly communication infrastructure to workflows and processes, from publications to data, software, and other research outputs, and the links between them, and strengthen the relationship of European Open Access (OA) infrastructures with other regions of the world, in particular Latin America and the U.S.A. OpenAIRE has created operating links to related data infrastructures, registries and initiatives (EBI, arXiv, CERN, DataCite, re3data.org, OpenDOAR, ROMEO/FACT, ORCID, EGI, EUDAT), and has established the foundations for an integrated Open Scholarly Communication Infrastructure, offering technical and support services to a wide range of stakeholders. The network of National Open Access Desks (NOADs) of OpenAIRE provides an OA communication & outreach platform with active presence in 33 European countries.

PRACE

The publicly owned and operated PRACE installations offer capability style HPC resources to researchers selected via calls for proposals and peer reviewed according to scientific excellence criteria. As a result of the calls, selected researchers are allocated CPU hours on scarce supercomputer installations. In comparison the PCP pilot project will procure capacity style cloud services hosted on commodity cluster installations operated by commercial providers.

PRACE does not offer long-term storage facilities and is currently looking to establish partnerships with other e-infrastructures that can provide such services. The European Open Science Cloud could provide a channel

[45] NA2.2 Strategy, Business Development and Exploitation https://wiki.egi.eu/wiki/EGI-Engage:WP2

by which PRACE is able to offer long-term storage for the results of capability HPC simulations and permit their post-analysis on commodity cloud services.

**How would the very broad science user community be involved in steering the project (biomedical research, climate, social sciences…)?**

The EC requires PCP projects to measure the impact of the procurement and we propose that this can be achieved via metrics such as the relative percentage of the total capacity procured by the buyers via the PCP pilot project compared to other means; the numbers of users making use of these resources and their distribution both geographically and by scientific discipline.

During the preparation of the tender specifications for the cloud services, input will be solicited from the broad user communities that are supported by the group of buyers and the demand-side organisations that are members of the Helix Nebula initiative.

The decision about the usage of the procured cloud services remains with the buyers. For example, in the case of WLCG, the LHC experiments determine who has access to the pledged resources. As outlined in response to the questions above, we propose that a fraction of the cloud services procured by the buyer-group is made available to other disciplines, the long-tail of science and citizen scientists with access to these resources be made available via the EGI Fed Cloud and the recently announced eGrant portal[46].

**How would success of the pilot project be judged?**

The pilot project has a series of clear objectives against which its success can be judged:

- The foreseen buyers are all members of the WLCG collaboration and the primary research community will be high energy physics. The performance/quality/reliability of the procured services will be benchmarked/compared to the existing WLCG production grid.
- Increase the adoption of commercial cloud services in the publicly funded research sector. This can be measured as the increase in the total spending on commercial cloud services across the publicly funded research sector.
- Improve the cost-effectiveness of publicly funded research sector IT systems through efficient joint procurement of cloud services. A study of the cost effectiveness of e-infrastructures compared to commercial offerings was performed by the eFISCAL project[47] and a total cost of ownership study[48] was performed by SAP on specific CERN in-house services within the context of Helix Nebula. This analysis can be updated by measuring the cost of procuring cloud services via the PCP pilot compared to global market leaders. Financial comparison will be more complicated than technical benchmarking because WLCG does not collect such information so we will compare the cost of procured services to those of EGI and also compare to procuring directly from the global IaaS leader (Amazon).
- Produce a practical implementation of the European open science commons.

[46] https://www.egi.eu/news-and-media/newsletters/Inspired_Issue_14/eGRANT.html

[47] http://www.efiscal.eu/files/deliverables/D2%203%20Executive%20Summary%20-%20Computing%20e-Infrastructure%20cost%20calculations%20and-business%20_models_vam1-final.pdf

[48] D7.3 Costing exercise comparing in-house vs. cloud based operation for the CERN flagship use-case and incorporation of qualitative cloud adoption criteria targeting prospect members http://www.helix-nebula.eu/publications/deliverables/d73-costing-exercise-comparing-in-house-vs-cloud-based-operation-the-cern

- Provide a means by which Research Infrastructures could collectively provision their computing needs.

The pilot project should be seen as the first phase in the establishment of the European Open Science Cloud.





**What is the "ultimate purpose" of the 100% model of which the 10% pilot is procured via this PCP project?**

The ultimate purpose is to operate a sustainable European Open Science Cloud based on *IT as a Service* serving Europe's Research Infrastructures, communities and related business sectors. It will reduce fragmentation in Europe and be a fully functional implementation of the open science commons. Its combined capacity will largely surpass that which is currently available via existing public e-infrastructures and the in-house facilities of research organisations. The underlying hybrid cloud model foresees a continuous need for IT services provisioned via publicly funded data centres to guarantee long-term data preservation and commercial service supplier independence. A significant difference compared to the current model is that funding agencies and research organisations will no longer provision services *exclusively* from their in-house resources. Services will be provisioned from commercial suppliers when they are not available in-house or can be provisioned externally on better terms (i.e. at shorter notice, lower cost or better performance etc.) As outlined in the introduction to the paper, we see the scale and range of services being provisioned from commercial suppliers to gradually increase over time as the cloud market matures. The PCP pilot project provides the funding agencies and research organisations with the opportunity to prepare themselves for this new model and be in a position to continue to offer the highest quality of IT services to their users from a range of sources.

**What is the special expertise of CERN that makes it best placed to lead this effort?**

CERN is a major procurement organisation purchasing approximately 400 million euros of goods and services annually. CERN's procurement needs provide opportunities for European companies to develop innovative products and services. CERN has extensive experience in procuring IT equipment and services in a multi-national environment in accordance with a published procurement process embodying the principles of transparency and impartiality[49].

CERN has a long-standing R&D relationship with the IT industry via the CERN openlab[50] which has been expanded to include several research organisations serving various scientific disciplines and has recently published a whitepaper on future IT challenges in scientific research[51].

To handle the increasing data management requirements resulting from the upgrades of the LHC accelerator and detectors, CERN needs to expand the capacity of the IT services it provides to the physics community. Following a public tender in 2012, CERN contracted a second data centre with Wigner in Budapest[52], effectively "Datacenter as a Service". This provided the incentive for the computer service operations to be reviewed and modified to ensure remote operations would be possible. The contract will be subject to a competitive renewal and CERN is investigating various options to move to the next stage, including procuring commercial cloud services rather than data centre space and the successful execution of the PCP pilot project will provide CERN will clear evidence on the suitability of this approach.

Interaction between commercial cloud services and publicly funded e-infrastructures is an important requirement for the research communities and has been investigated as part of the Helix Nebula initiative that was form by CERN, EMBL and ESA. Three initial flagship use cases from high energy physics, molecular biology and earth-observation have been used to validate the approach, enable a cost-benefit analysis to be

[49] http://issuu.com/fpcern/docs/cern_procurement_brochure?e=9058305/4400413
[50] http://openlab.web.cern.ch/
[51] http://press.web.cern.ch/press-releases/2014/05/cern-openlab-publishes-whitepaper-future-it-challenges-scientific-research
[52] http://press.web.cern.ch/press-releases/2012/05/cern-awards-major-contract-computer-infrastructure-hosting-wigner-research





undertaken and the next stage of the Science Cloud Strategic Plan[53] developed and approved. The rationale for the European Open Science Cloud has been inspired by the publication *Helix Nebula – The Science Cloud: A catalyst for change in Europe*[54]. The work of Helix Nebula on an architecture model has produced a number of scenarios exploring the means by which publicly funded infrastructures can interoperate with commercial cloud services and produced a report on their integration with a set of recommendations[55]. Helix Nebula has also produced a number of reports directly relevant to the Cross-border procurement of e-Infrastructure services: *D7.1 Analysis of the existing supply and demand side – Big Science - business processes for the procurement of IT infrastructure services*[56]

[53] Helix Nebula D9.2 Strategic Plan for a Scientific Cloud Computing Infrastructure for Europe: Three years on. http://www.helix-nebula.eu/publications/deliverables/d92-strategic-plan-scientific-cloud-computing-infrastructure-europe-three

[54] http://www.helix-nebula.eu/sites/default/files/HelixNebula-NOTE-2013-003.pdf

[55] Helix Nebula D6.2: Roadmap for the integration and interoperation of commercial cloud with e-Infrastructures http://www.helix-nebula.eu/publications/deliverables/d62-roadmap-the-integration-and-interoperation-of-commercial-cloud-e

[56] http://cds.cern.ch/record/1529847/files/HelixNebula-D7_1.pdf