



EGI recommendations for the Horizon 2020 Work Programme 2016-2017

EC DG CONNECT E2 public consultation on Cloud Computing and Software

Michel Drescher (Michel.Drescher@egi.eu)

With contributions from:

EGI management &
EGI Federated Cloud Task Force

17 October 2014

Version 1

<https://documents.egi.eu/document/2320>

COPYRIGHT NOTICE

Copyright © 2012 EGI.eu. This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA.

The work must be attributed by attaching the following reference to the copied elements: "Copyright © EGI.eu (www.egi.eu). Using this document in a way and/or for purposes not foreseen in the license, requires the prior written permission of the copyright holders. The information contained in this document represents the views of the copyright holders as of the date such views are published.



TABLE OF CONTENTS

Executive summary	3
Recommendation 1: Improving the Cloud service reference model	4
Description	4
Rationale.....	4
Required changes	4
Impact.....	4
Beneficiaries	4
Recommendation 2: Cross-Cloud identity environments	5
Description	5
Rationale.....	5
Required changes	5
Impact.....	5
Beneficiaries	5
Recommendation 3: European data portability policies	6
Description	6
Rationale.....	6
Required changes	6
Impact.....	7
Beneficiaries	7
Recommendation 4: Expand the portfolio of virtualised computing hardware	7
Description	7
Rationale.....	7
Required changes	7
Impact.....	8
Beneficiaries	8
Recommendation 5: Intensified support for Open standards interoperability	8
Description	8
Rationale.....	8



Executive summary

Over the course of three years, EGI has taken on the task of designing and deploying a federation of institutional private Clouds in Europe while applying its decade-long experience and expertise in operating a federated scientific Grid infrastructure for the benefit of interested research communities to make use of this new model of resource sharing. While the majority of EGI Federated Cloud members are academic institutions by nature, it is open to commercial and public administration members as well. With this approach EGI has access to a wide variety of requirements on both the demand side through its engaging user communities and the supply side coming from its Cloud service provider members.

Key to the success of EGIs federated Cloud task force (and the corresponding infrastructure deployed in Europe) is the fundamental concept of implying and (to a certain extent) enforcing consistency in the delivery of services through and by the members of the federation. Consistency may manifest in a variety of ways, including:

- Exposed service access interfaces governed through publicly defined standards (such as OCCI, CDML, and others),
- Access control and identity management procedures (through a shared AAI)
- Common resource description and consumption models

Even though EGI has come a long way in the process of assembling and operating a federated Cloud infrastructure, there are still a number of gaps requiring attention on the way implementing the Digital Agenda for Europe. Consequently, the recommendations formulated in this contribution all share the aim of increasing consistency, and thus comparability of Cloud services in a visional single European market:

The first recommendation tackles the semantic confusion and shortfall of appropriate terminology for Cloud service offerings, followed by the second recommendation to harmonise user identity managed in multi-Cloud (including federations) environments - in a way this proposes user identity portability across Cloud service providers. Just as important, data portability improvements are described in the third recommendation to bring the classic economic free flow of goods and services into the digital market. Stemming from EGI's scientific and academic background, recommendation four suggests expanding the portfolio of hardware virtualisation support for processing components such as GPGPUs, general-purpose ASICs and others. Last but not least, EGI suggests to implement an impartial activity able to assess (and perhaps certify) standards compliance and interoperability informing Cloud demand side and support Standards support requirements in procurement processes.

Recommendation 1: Improving the Cloud service reference model

Description

Fill the term “resource” with meaning and understanding. Reflect on reference models and reassess the validity of core parts, such as the typical deployment models IaaS, PaaS and SaaS.

Rationale

The term “resource” is an abstract term for used for a wide variety of things that are generally considered as consumables that a customer of the Cloud service may utilise. Its use ranges from very detailed hardware-oriented components such as RAM, or CPUs (CPU cores) to something as difficult to capture as some classes of software provided as a service. The notion of “everything is a resource” is overly simplistic and rather hides the different types of resources and related intricacies.

This corresponds to the currently popular use of the terms “XaaS”, “*-aaS” or “Anything as a service”, which all reflect both the current confusion and contending definitions of Cloud Computing and the need to expand on the experienced constrictions of the terms IaaS, PaaS and SaaS.

Required changes

The proposed ontology should be developed from existing material. It needs to provide viable classifications of resources and develop definitions of associated capacities, capabilities, and consumption models. It might be delivered in the form of a Glossary, Compendium, formal ontology or any other way that is easily digested and shared.

The technical activity should be accompanied by appropriate dissemination and training activities to promote and facilitate uptake and usage across all stakeholders.

Impact

Many discussions still spend time on agreeing on a common understanding on terms, context and semantics before the actual topic at hand can be reasonably discussed. Even worse, if this does not happen, participants’ differences in understanding the topic are often overlooked and not addressed, which may cause a false understanding of agreement where in fact cross-purpose discussion may cause even further confusion on the topic.

Beneficiaries

All stakeholders in the Cloud services market will benefit from developing and training a common language. Customers will benefit from clearer understanding of services at hand, service providers may find it easier to position themselves and their added value in an understandable manner.

Recommendation 2: Cross-Cloud identity environments

Description

Define the scope and surrounding policies and processes concerning customer identity in a multi-Cloud usage scenario.

Rationale

Uptake of Cloud services is steeply rising, and consumers are increasingly facing the situation of making use of more and more Cloud services simultaneously. However, this one-to-many scenario is complemented by the converse, i.e. the many-to-one, and to make it even more complex, many-to-many relationships with Cloud services.

Contemporary Cloud service business models often over-simplify the situation into a “Credit cards are the new identity” paradigm backed by simple username-based credentials merely used to apply and enforce quotas on resources offered as services. While this model fits well with the typically isolated business model still prevalent in the commercial Cloud market, it fails in properly addressing the collaborative nature of research and economy envisioned in the Digital Agenda for Europe.

While service federations and identity federations already exist in Europe (for example, EGI and TERENA, respectively, representing the academic sector in Europe) serving the one-to many and many-to-one problem reasonably well, there is a gap that yet needs addressing the many-to-many problem in collaborative research and business in Europe.

Required changes

Representative for many other scenarios, highly-collaborative academic projects typically entail many-to-many relationships between researchers and research infrastructure providers. Introducing the concept of a “virtual organisation” helps managing the technical complexity, but needs further development and expansion around the definition of a tenant (is it the customer, or is it the VO?) and associated issues concerning privacy, data protection, identity mapping/management, operational security, service management, responsibilities, and more.

Impact

A common, or a set of commonly accepted approaches to addressing this problem has the potential to boost the otherwise constricted collaboration at large scale across Europe. It may enable a lightweight process of defining, and implementing collaborative projects over a period of time, without having to worry about the typical “legalese” of many-to-many business relationships.

Beneficiaries

We expect that no less than the entire Cloud services consumer ecosystem will benefit from addressing this problem for as long as collaboration of otherwise independent entities is a key component of the activity, be it research, commercial joint ventures, or public administration.



The second immediate beneficiary will be Cloud service providers offering services to collaborative endeavours, in that the outcome of activities in this area will help building reassurance that customer duties will still be met when scaling out Cloud services from individual relationships to to collaborations.

Recommendation 3: European data portability policies

Description

Data portability issues are an important factor in Cloud service procurement. Clear procedures and policies describe terms and conditions on the technical and business level.

Rationale

Cloud service uptake is steeply rising even though unresolved issues with data portability in the Cloud still exist - perhaps mostly due to the commonly adopted strategy of deferring the issue resolution to a later point in time when it may (as opposed to proactive issue management).

The current situation leaves Cloud customers with very little to no influence (let alone control) over the terms and conditions of data portability between Cloud providers - since “data is the new gold” this inequality arguably enables Cloud providers to hold customers ransom over the data stored in their service.

Three scenarios should be considered in scope for data portability policies and T&Cs:

Data intake; a customer would like to engage with a Cloud service provider but needs bulk data intake/ingestion preceding the operational service operation.

Data extraction; a customer desires to extract data from a Cloud service provider (e.g. to implement its own disaster recovery process).

Data removal; similar to household removal services, a customer needs data being transported from one Cloud service provider to another (for example, to switch business from one to the other, or to improve data availability across Europe)

Required changes

While tackling this issue seem manageable in IaaS scenarios it is less clear or outright impossible to date when engaging with CSPs in PaaS or SaaS scenarios: Considering CRM, Service Desk, and similar SaaS existing today, how would a customer be able to achieve either data portability scenario today, or even in a future single digital economy in Europe? How might the exported data structures look like, how would it be made available?

What would be applicable besides technical aspects of data portability, harmonisation of policy aspects across Europe should be scoped and promoted, such as pricing structures, request entitlement, fulfilment, and last but not least data privacy.



Impact

The biggest benefit of a clear data portability framework applicable to a single digital European market would allow transporting the principle of “free flow of goods and services” from the classic European economy into the digital market, ideally simplifying the data portability problem to a question of mere logistics.

Beneficiaries

Clearly, Cloud service consumers would be the major beneficiaries of a data portability framework as it would provide an instrument to demand service providers implementing in their services, contributing to a reduction of vendor lock-in for data stored in any type of Cloud service. Secondary beneficiaries are clearly SME Cloud service providers for which such a framework would create a level playing field on which they can compete in service delivery quality instead of investing significant effort and resources into legal and business competition that only major players can afford.

Recommendation 4: Expand the portfolio of virtualised computing hardware

Description

Increase support for virtualisation of computing hardware beyond general-purpose CPUs, RAM and Storage.

Rationale

Current Cloud computing models focus on utilising general-purpose computing hardware to achieve efficiency and exploit economies of scale in competitive commercial (mainstream) Cloud computing use cases.

However, particularly academic use of Cloud services in research impose different requirements on virtualised infrastructure. For example, GPGPU platforms have gained significant momentum and popularity in the scientific community to solve specific classes of algorithmic problems. Such types of hardware typically do not provide hardware-level virtualisation (in contrast to general-purpose CPUs), which makes it next to impossible to implement efficient resource sharing of such alternative computing platforms.

Required changes

Looking beyond popular virtualised computing hardware such as RAM, CPU, and (within limits) NICs alternative (and to some extent more exotic) hardware components should be subject to virtualisation support both in hardware and corresponding drivers in the deployed hypervisor.

The range of applicable hardware may include GPGPUs, general purpose FPGAs and DSPs. However, with peripherals interfaces reaching data transmission capabilities in the range of GBytes per



seconds (e.g. USB 3 and Thunderbolt, to name all but two) any type of pluggable peripheral device may raise interest for sharing.

Impact

Cloud services business has been very much driven by the economic needs of the commercial and industrial communities, and implemented using commodity hardware serving general-purpose use cases, delivering research-driven high-performance and high-throughput computing. Cloud services can benefit from hardware-supported virtualisation of less general-purpose hardware components.

Leading edge research often faces the dilemma of compromising: Will commodity hardware be used in order to drive down the cost of ICT infrastructure even though other hardware platforms might be a better fit for the application domain at hand, or will the application domain needs have priority over the TCO of the necessary ICT resources? Effective virtualisation support in hardware for such components may help solving this dilemma.

Beneficiaries

It is certainly possible to increase application efficiency and performance through parallelising the application domain, and scale out resource consumption using Cloud computing paradigms. But this is only possible to a certain degree (subject to the application domain at hand), while it still would have to integrate with POTS hardware for which efficient virtualisation support is available.

Any research domain can benefit from extended hardware virtualisation support, spanning academic, commercial and industrial research in Europe and beyond.

Recommendation 5: Intensified support for Open standards interoperability

Description

Support the implementation and uptake of open standards beyond current involvement, focussing on openness of standards.

Rationale

The past and current work programmes included implementation, support for, and promotion of open standards in the scope and eligibility for funded projects. Through organising two DG CONNECT E2 concertation meetings, the CloudWatch project (in which EGI.eu represents EGI) elicited a number of position statements from FP7-funded projects related to their implementing certain open standards such as OCCl, CDMI, and others. However, it is currently unknown to CloudWATCH how well these implementations interoperate with other known implementations outside those projects, and no indication or references were given by the projects themselves. Public dissemination and general invitations to interoperability events did not yield the level of participation among those FP7 projects as we expected.



This may leave a gap of assurance to an outsider of the project of the level of maturity and compliance of software claiming to implement a bespoke standard. Such gaps inevitably will interfere not only with any third party interest in re-using the developed software, but also with any public funding agency in the return of investment in that aspect of their funding schemes.

Required changes

Future programmes should call for CSA supporting and aiding projects in producing the required evidence of standards implementation. Support should be given through the entire lifetime of the supported project (where possible) demonstrating strategic selection of open standards, implementation strategies, and progress towards achieving compliance with the selected standards.

Further, work programmes shall call for CSAs to charter the landscape of standards, their level of openness (through assessment of the SDO's documented processes and policies) resulting in a portfolio of standards, possibly including a badge of approval indicating compliance with the CSA's published assessment framework.

EGI recommends to liaise with international funding agencies and public administrations to explore the inception and operation of an impartial council chartered to produce reliable standards related information for Cloud service demand side to include standards support in their procurement process. For example, a council similar to the US Government's Cloud Standards Customer Council might be envisioned to implement in Europe.

Impact

Open standards, by design, lack the control of one single entity over the specification(s) comprising the standard. While being often considered a significant design flaw, this lack of control of one ensures that changes to the standard will not come unexpected and through well-documented channels. Openness ensures non-discriminatory inclusion (or at least proper processing) of any voice and contribution according to open processes and decision-making.

Putting emphasis on documented evidence on standards compliance allows creating a level playing field particularly for SMEs that do not have the effort and market dominance to maintain an industry (or pseudo) standard, so that they can compete in excellence of service delivery for a defined service baseline rather stick to fragmentation and limitation of the technical capabilities, while retaining the potential for innovation beyond that - on comparable terms provided by the standardised baseline.

Inherently, standards codify shared understanding of the target domain at hand. In other words, standards drive commoditisation of the services and resources covered, hence further driving down the cost of procurement and lowering the vendor lock-in (for as long as the procured service is entirely covered by standards).

Beneficiaries

Pushing for evidence of standards compliance allows Cloud service consumers, particularly the public administration and publicly funded institutes to start including demand for support of certain standards into their procurement and tender processes. It therefore it helps protecting the initial investment in projects for implementing or integrating with bespoke standards by circumventing the typical chicken-and-egg problem of demand versus supply. Using a clear surrogate demand (perhaps by stating "we require you to implement certain standards in exchange for financial remuneration") will set the market mechanics into motion by changing the current stalemate situation into a situation of market opportunity.