



EGI-InSPIRE

INTEGRATION OF CLOUDS AND VIRTUALISATION INTO THE EUROPEAN PRODUCTION INFRASTRUCTURE

EU DELIVERABLE: D2.6

Document identifier:	EGI-D2.6-258-v.9.doc
Date:	01/03/2011
Activity:	NA2
Lead Partner:	EGI.eu
Document Status:	FINAL
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/258

Abstract

Virtualisation and cloud computing have demonstrated how new technologies can enable dynamic execution environments or on-demand elastic service deployment with new, clear cost measurements and business models. Due to the financial constraints being felt throughout Europe, and that ICT policies and services tailored to the current e-infrastructure user communities do not always meet the needs of new communities, EGI needs to evolve to provide a more flexible, efficient e-infrastructure in order to attract new users from all disciplines. Therefore, this report is designed to build the foundation for integrating virtualisation and cloud technologies into EGI to better address the evolving user needs. It analyses the technology benefits and issues, economic aspects of delivering such resources, with a short- and long-term view to identifying why, where and how these technologies have a place within the EGI.

COPYRIGHT NOTICE

Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration. EGI-InSPIRE (“European Grid Initiative: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe”) is a project co-funded by the European Commission as an Integrated Infrastructure Initiative within the 7th Framework Programme. EGI-InSPIRE began in May 2010 and will run for 4 years. This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 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, and USA. The work must be attributed by attaching the following reference to the copied elements: “Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration”. 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.

I. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Sergio Andreozzi	EGI.eu/NA2	01/03/2011
Reviewed by	Cal Loomis Eric Yen Tryfon Chiotis	NA2	24/01/2011
Approved by	AMB & PMB		1/01/2011

II. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
ToC	07/01/2011	Table of Contents (ToC)	Sergio Andreozzi, Sy Holsinger, Steven Newhouse/EGI.eu
1	13/01/2011	First draft of initial content	Sy Holsinger, EGI.eu
2	14/01/2011	Reorganisation of content	Sergio Andreozzi, Sy Holsinger/EGI.eu
3	21/01/2011	Second draft of content for internal review	Sy Holsinger/EGI.eu
4	25/01/2011	Revised version from first internal review	Sergio Andreozzi, Sy Holsinger, Steven Newhouse/EGI.eu
5	28/01/2011	Revised version from second internal review	Sy Holsinger/EGI.eu
6	02/02/2011	Final version for external review	Sergio Andreozzi, Sy Holsinger, Steven Newhouse/EGI.eu
7	11/02/2011	Revisions from external reviewer comments	Sy Holsinger/EGI.eu
8	15/02/2011	Final version for AMB approval	Sy Holsinger/EGI.eu
9	01/03/2011	Revised from PMB comments and approved	Steven Newhouse/EGI.eu

III. APPLICATION AREA

This document is a formal deliverable for the European Commission, applicable to all members of the EGI-InSPIRE project, beneficiaries and JRU members, as well as its collaborating projects.

IV. DOCUMENT AMENDMENT PROCEDURE

Amendments, comments and suggestions should be sent to the authors. The procedures documented in the EGI-InSPIRE “Document Management Procedure” will be followed: <https://wiki.egi.eu/wiki/Procedures>

V. TERMINOLOGY

A complete project glossary is provided at the following page: <http://www.egi.eu/about/glossary/>.



VI. PROJECT SUMMARY

To support science and innovation, a lasting operational model for e-Science is needed – both for coordinating the infrastructure and for delivering integrated services that cross national borders.

The EGI-InSPIRE project will support the transition from a project-based system to a sustainable pan-European e-Infrastructure, by supporting ‘grids’ of high-performance computing (HPC) and high-throughput computing (HTC) resources. EGI-InSPIRE will also be ideally placed to integrate new Distributed Computing Infrastructures (DCIs) such as clouds, supercomputing networks and desktop grids, to benefit user communities within the European Research Area.

EGI-InSPIRE will collect user requirements and provide support for the current and potential new user communities, for example within the ESFRI projects. Additional support will also be given to the current heavy users of the infrastructure, such as high energy physics, computational chemistry and life sciences, as they move their critical services and tools from a centralised support model to one driven by their own individual communities.

The objectives of the project are:

1. The continued operation and expansion of today’s production infrastructure by transitioning to a governance model and operational infrastructure that can be increasingly sustained outside of specific project funding.
2. The continued support of researchers within Europe and their international collaborators that are using the current production infrastructure.
3. The support for current heavy users of the infrastructure in earth science, astronomy and astrophysics, fusion, computational chemistry and materials science technology, life sciences and high energy physics as they move to sustainable support models for their own communities.
4. Interfaces that expand access to new user communities including new potential heavy users of the infrastructure from the ESFRI projects.
5. Mechanisms to integrate existing infrastructure providers in Europe and around the world into the production infrastructure, so as to provide transparent access to all authorised users.
6. Establish processes and procedures to allow the integration of new DCI technologies (e.g. clouds, volunteer desktop grids) and heterogeneous resources (e.g. HTC and HPC) into a seamless production infrastructure as they mature and demonstrate value to the EGI community.

The EGI community is a federation of independent national and community resource providers, whose resources support specific research communities and international collaborators both within Europe and worldwide. EGI.eu, coordinator of EGI-InSPIRE, brings together partner institutions established within the community to provide a set of essential human and technical services that enable secure integrated access to distributed resources on behalf of the community.

The production infrastructure supports Virtual Research Communities (VRCs) – structured international user communities – that are grouped into specific research domains. VRCs are formally represented within EGI at both a technical and strategic level.



VII. EXECUTIVE SUMMARY

As the economic crisis forces Europe to take a hard look at public expenditure, recurring themes have tended to arise around aspects such as streamlining staffing costs, evaluating green energy, and achieving economies of scale. For the IT industry, especially large-scale e-Infrastructures, emerging technological solutions in the commercial sector have become potentially attractive in the academic research arena, such as the consolidation of data centres and wide-scale adoption of virtualisation. The ability to provision resources 'on-demand' to meet the needs of particular research collaboration and the implementation of cloud computing and business models have shown the use of virtualisation to deliver 'Infrastructure as a Service', hosted environments to provide a 'Platform as a Service' and hosted applications to access 'Software as a Service'.

This report aims to evaluate these technologies, understand how they relate to EGI, and build a foundation for the integration of cloud and virtualisation into the European production infrastructure. More specifically, after a brief introduction of the overall landscape, the document puts in perspective the current structure and status of EGI and provides an overview of cloud computing technologies and operation models. A dedicated section also looks at how the e-Infrastructure community at large is tackling the issue of cloud computing through collaborations and publicly funded projects. The report finally offers the vision of EGI, what is driving the change, and a cost analysis and comparisons to current market offers. The report concludes with a short- to long-term strategic roadmap for evolving EGI towards virtualisation.



TABLE OF CONTENTS

1	INTRODUCTION	7
2	EGI: CURRENT STATUS	8
2.1	Infrastructure	8
2.2	Scale	9
2.3	Usage	10
3	CLOUDS: TECHNOLOGIES & OPERATING MODELS	11
3.1	Background	11
3.2	Deployment Models	11
3.3	Service Models	12
3.4	Features	13
3.5	Benefits and Issues for EGI	14
4	CURRENT COMMUNITY CLOUD ACTIVITIES	16
4.1	Dynamic Execution Environments	16
4.1.1	Worker Nodes on Demand Service	16
4.1.2	Batch and Server Virtualisation and Cloud Integration	17
4.2	Provisioning Grids in Clouds	18
4.2.1	Grid on Demand	18
4.2.2	RESERVOIR	19
4.2.3	StratusLab	20
4.3	Application Suitability	20
4.3.1	Venus-C	20
4.4	Summary	21
5	EVOLVING EGI	22
5.1	Drivers	22
5.1.1	Public and European	22
5.1.2	Organisational	24
5.1.3	Economic	26
5.2	Cost Analysis	27
5.2.1	Scenarios	27
5.2.2	Cost Estimate: EGI-InSPIRE	27
5.2.3	Cost Estimate: e-IRGSP2	28
5.2.4	Cost Estimate: Amazon Web Services (EGI-InSPIRE Proposal)	28
5.2.5	Summary	29
5.3	A Vision for Integrating Virtualisation in EGI	30
5.3.1	Architecture	31
5.3.2	A Virtualised Ecosystem	32
5.3.3	Collaborations	33
6	TOWARDS A TECHNOLOGY ROADMAP	35
6.1	User Oriented Objectives	35
6.2	Required Technology Capabilities	35
6.3	Risks	36
6.4	Follow-up	36
7	CONCLUSION	37
8	REFERENCES	38



TABLE OF TABLES

<i>Table 1: Cloud Computing Benefits and Issues</i>	15
<i>Table 2: Estimated Total Cost of the EGI production infrastructure during EGI-InSPIRE</i>	27
<i>Table 3: Cost of moving EGI to Amazon</i>	29
<i>Table 4: Additional Costs & Savings to EGI/AWS Cost Comparison</i>	30
<i>Table 5: Virtualisation Effect on User Groups</i>	33
<i>Table 6: Identified needs from the user communities</i>	35
<i>Table 7: Capabilities required to augment UMD to support virtualisation and cloud-like services</i> ..	35
<i>Table 8: Barriers and Solutions to the EGI Vision</i>	36

TABLE OF FIGURES

<i>Figure 1: EGI Federated Resource Layers</i>	8
<i>Figure 2: EGI Virtuous Cycle</i>	9
<i>Figure 3: EGI – Cores & Sites</i>	10
<i>Figure 4: Competitiveness Cycle</i>	22
<i>Figure 5: A Virtualised Ecosystem</i>	31



1 INTRODUCTION

The European Grid Infrastructure (EGI) provides access to a federated distributed computing infrastructure to European researchers for a variety of scientific domains who are facing the challenge of how to process the deluge of large-scale data being generated by their communities.

The infrastructure has evolved from predecessor capacity building projects such as the European Data Grid (EDG) and Enabling Grids for e-Science (EGEE) that enabled transnational access to computing, storage and networking resources. However, the current set of services tailored to the requirements of the initial scientific communities does not always meet the needs of new communities (e.g., lack of flexibility of user environments and use of different technologies). EGI therefore needs to evolve its service offering in order to become a more flexible infrastructure for attracting new users on a wider scale.

As grid was consolidating, the commoditisation of virtualisation and cloud computing started to emerge, demonstrating how new technologies can enable dynamic execution environments or on-demand elastic service deployment. New business models supporting clear cost measurements and better quality of service isolation have been demonstrated in many commercial environments.

On top of this general trend in ICT provision, the economic and social crisis has exposed some structural weaknesses in the economy, leading policy makers to redefine EU strategic priorities and vision to effectively tackle other long-term challenges. Furthermore, financial constraints of most of the European states have caused many funding issues for the NGIs, EIROs and the EGI community as a whole, thus promoting a reassessment of its ICT provision and alignment to other public sectors and policies [R21].

It is under these contexts that EGI is defining how to better address the evolving user needs by exploiting these emerging technologies. EGI already started developing a vision for the future of the infrastructure [R20] and, throughout this report, builds the foundation for the integration of clouds and virtualisation into the EGI, provides a detailed analysis of the technology benefits and issues, economical aspects of delivering the new services and sets out the context for defining both technology and implementation roadmaps.



2 EGI: CURRENT STATUS

EGI’s principle mission is to create and maintain a pan-European Grid Infrastructure enabling sharing of digital resources for computing, storage, and data, facilitating research across diverse scientific communities. In order to guarantee the long-term availability of a generic e-Infrastructure for all European research communities and their international collaborators, EGI.eu works in collaboration with its participants (e.g., NGIs, EIROs).

2.1 Infrastructure

From the infrastructure viewpoint, the smallest resource administration domain in EGI is called a “resource centre”. It can be either localised or geographically distributed and provides local resources and the functional capabilities necessary to make those resources accessible to authorised users. Resource centres federate together into a “resource infrastructure provider” that is a legal organisation responsible of establishing, managing, and operating directly or indirectly the operational services to an agreed level of quality needed by the resource centres themselves and the user community. Each resource infrastructure provider holds the responsibility of integrating them in EGI through the coordination of EGI.eu to enable uniform resource access and sharing for the benefit of their consuming end-users. In Europe, Resource Infrastructure Providers are NGIs and EIROs [R1].

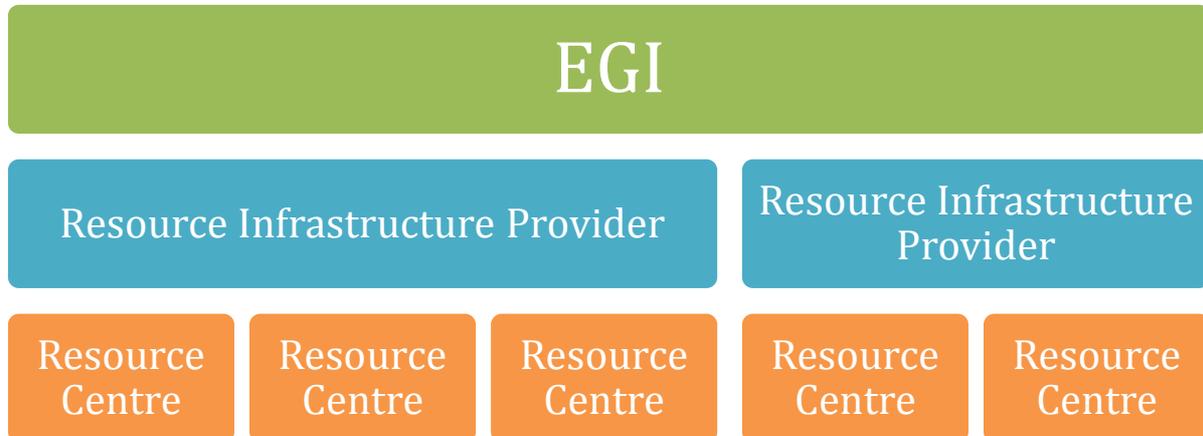


Figure 1: EGI Federated Resource Layers

EGI services are provided locally by Operations Centres and globally by EGI.eu in collaboration with some partners in the community. Local and global operations services are mutually dependent and can be complemented by additional services customised for local Virtual Organisations (VOs) and local Resource Centres. EGI integrates a wide range of distributed operational tools through a tiered architecture that is generally applicable to all EGI operations services [R1].

EGI.eu has the responsibility of coordinating the core needs of the EGI Community (end users and operations), managing the delivery of software, meeting requirements, and deployment of its technical services being provided from its partners (e.g., NGIs, EIROs). EGI.eu creates a virtuous circle that could easily transform into a vicious circle if specific measures are not put in place and feedback is either not obtained or ignored (See **Figure 2**).

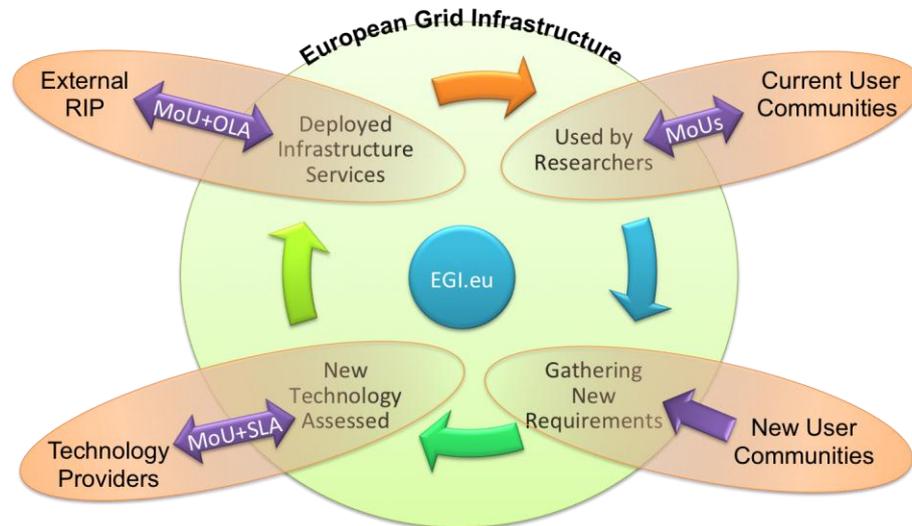


Figure 2: EGI Virtuous Cycle

This virtuous circle is expanded through processes and mechanisms during each step of the cycle that is executed through external work such as:

- 1) Establishing Memoranda of Understanding (MoUs) between EGI.eu and current user communities;
- 2) Gathering and prioritising new requirements that support current users and allow EGI to support new user communities;
- 3) Communicating requirements to External Technology Providers and coordinating the deployment of these new technologies when they become available;
- 4) Establishing Service Level Agreements (SLAs) between resource providers for the operation of these deployed technologies.

Through SLAs, Technology Providers agree to deliver software components to EGI that, in total, implement the functionality of one or more capabilities defined in the UMD Roadmap [R3]. However, the providers are given the free choice as to which capability to implement, when they will deliver it, but for it to be used by EGI, it must be integrated into EGI's support structure [R4]. Overall, the External Technology Providers provide the innovation needed by EGI to satisfy its users that cannot be found in the commercial or mainstream open source community and is therefore highly endorsed by EGI for use within the production infrastructure.

2.2 Scale

It is the European scale of EGI that provides both its value and its operational challenges. As can be seen below, EGI comprises over 250,000 cores with a federated model spread over 50 countries with around 60 sites with more than 1000 cores and over 130 sites with less than 100 cores, more than 330 sites in total. EGI is the largest multi-disciplinary e-Infrastructure in the world.

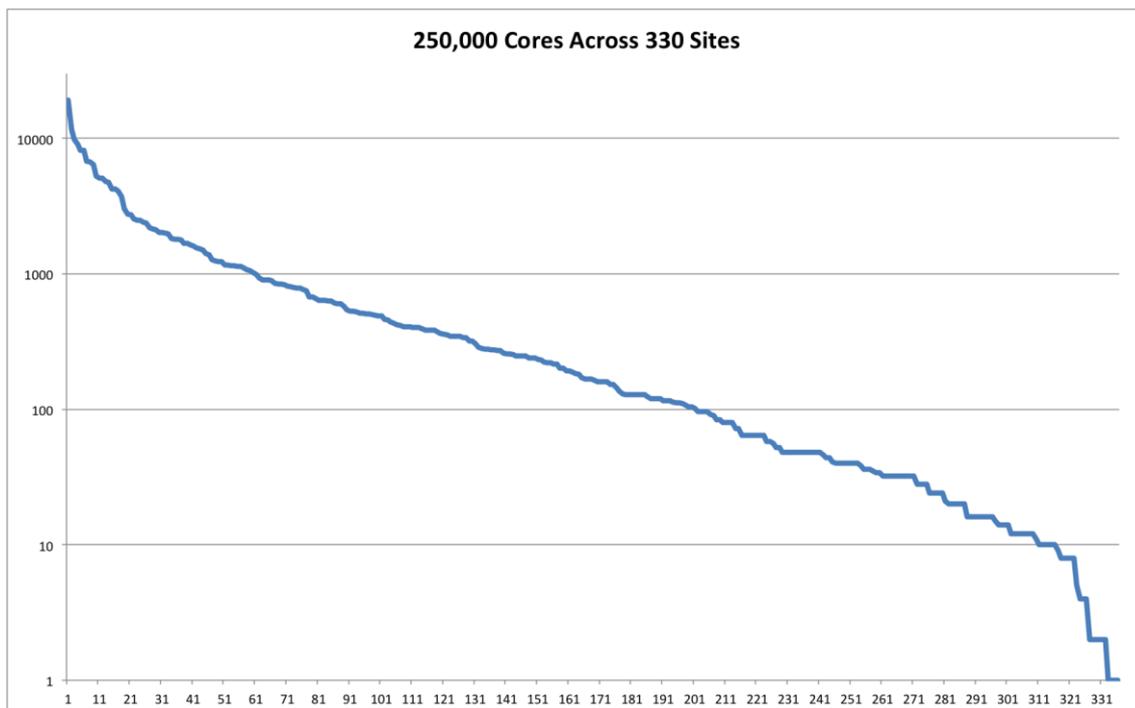


Figure 3: EGI – Cores & Sites

2.3 Usage

Many of the current grid users engage using command line interfaces. However, a significant number prefer web portals or domain specific environments that interact directly with grid resources. Applications with more advanced and user-friendly interfaces have been developed over the years to manage capabilities such as job management and workflow coordination. Nevertheless, there are an untold number of potential users who are effectively alienated by the complexity and inflexibilities of even these interfaces.

The current middleware stacks adopted in EGI provide similar capabilities and in general require skilled system administrators to manage the process of maintaining and upgrading the infrastructure for particular communities. This typically involves managing a cluster of machines and maintaining the up-to-date operating systems and middleware applications on these machines. The most commonly deployed solution, gLite, imposes a significant constraint in following exactly a particular version of Scientific Linux. The key benefit of this is that the administrator, and hence the end user, has a highly stable system on which to work. The downside includes an inability to run applications depending on different flavours of UNIX (e.g., Debian or BSD).

Another limitation in the current middleware deployment model is that it is managed by the local system administrators, making the infrastructure inherently static and bound to specific software stacks and computing models. It stops user communities being able to deploy technologies or updates to their software environments at a timescale that suits them. An example of where this is particularly restrictive is in training where an academic might want to enable access for 100 users for 2 days to a particular software environment on some sites, after which the accounts, the environment and files created can be removed.

3 CLOUDS: TECHNOLOGIES & OPERATING MODELS

3.1 Background

According to the National Institute of Standards and Technology (NIST) [R36], cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction [R25].

Cloud computing is fundamentally shifting the economics of IT. For consumers, it facilitates elastic consumption, self-service and pay-as-you-go pricing with computing capabilities ranging from data storage and processing to software available instantly and on-demand. It allows large data centres to standardise and pool IT resources and automate many of the maintenance tasks previously done manually to produce significant economies of scale.

Cloud computing represents a new way of delivering computing capability (i.e. a business model or philosophy) rather than a new technology, while virtualisation techniques have helped make corporate servers more efficient by allowing multiple applications to run on multiple operating systems on the same machine. The new underpinning economic model has gained increasing popularity and global investment. According to IDC's analysis, the worldwide forecast for cloud services in 2009 was estimated to be in the order of \$17.4B. The estimation for 2013 amounts to \$44.2B, with the European market ranging from €971M in 2008 to €6B in 2013 [R5].

Virtualisation, which underpins cloud computing, has led to widespread changes in commercial data centres. Virtualisation can be typically defined when a Virtual Machine (VM) is created as a "representation" of a physical machine using software that has its own set of virtual hardware hosting a single or multiple operating system(s) in which applications can be loaded. Using virtualisation, each VM is created with consistent virtual hardware regardless of the underlying physical hardware that the host server is running. A VM can be further customised by adding or removing additional virtual hardware as needed by editing its configuration.

The rest of this section introduces the different deployment models of virtualisation technology to deliver cloud computing capabilities to different user communities, defines the common service models built on top of the virtualised resources, and the features commonly offered by such cloud environments. The section concludes by summarising the open issues around cloud technologies as they relate to EGI.

3.2 Deployment Models

There are different models upon which a cloud infrastructure can be deployed. In this section, those that have emerged to date are presented [R26].



Private Cloud - also referred to as internal cloud or on-premise cloud, intentionally limits access to its resources to service consumers that belong to the same organisation that owns the cloud. The infrastructure is managed and operated for one organisation only, primarily to maintain a consistent level of control over security, privacy, and governance.

Public Cloud - also referred to as external cloud or multi-tenant cloud, this model essentially represents a cloud environment that is openly accessible. It generally provides an IT infrastructure in a third-party physical data centre that can be utilised to deliver services without having to be concerned with the underlying technical complexities.

Community Cloud - refers to special-purpose cloud computing environments where resources are pooled together and managed by a number of related organisations participating in a common domain or vertical market. It may be managed by the organisations or a third party and may exist on premise or off premise.

Hybrid Cloud - is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardised or proprietary technology that enables data and application portability.

Federated Cloud - a composition of a number of private clouds working in collaboration to deliver an integrated cloud resource to specific user communities that is based on the aggregation of deployed clouds, that remain unique entities but are bound together by standardised or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

3.3 Service Models

Understanding the core architectures of cloud computing is vital in pursuing the right cloud computing solution. Each organisation chooses a cloud service (along with a deployment model, described above) based on their specific business, operational, and technical requirements. Below is a short overview the three principle service models and their market value [R27].

Cloud Infrastructure as a Service (IaaS) - The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

IaaS offers basic computing services, from compute nodes to data storage, which customers can combine to build highly adaptable computer systems. The market leaders are GoGrid, Rackspace and Amazon Web Services (the computing arm of the online retailer). Forrester Research [R38], a



consultancy firm, predicts that revenues generated by computing infrastructure as a service will grow to nearly \$56B by 2020.

Cloud Platform as a Service (PaaS) - The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

PaaS is an operating system living in the cloud. Such services allow developers to write applications for the web and mobile devices. Offered by Google, Salesforce.com, and Microsoft Azure. This market is also fairly easy to measure, since there are only a few providers and their offerings have not really taken off yet. Forrester puts revenues at a mere \$311M.

Cloud Software as a Service (SaaS) - The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

SaaS includes web-based applications such as Gmail (Google's e-mail service) and Salesforce.com, which helps firms keep track of their customers. This layer is by far the easiest to gauge. Many SaaS firms have been around for some time and only offer such services. Forrester estimates that these services generated sales of \$11.7B in 2010.

3.4 Features

Cloud computing is generally characterised by a number of features described below [R25, R28]:

On-demand self-service - A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically through programmatic interfaces or through web management portals without requiring human interaction with each service's provider through virtualisation and automation technologies.

Publicly accessible - Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).

Multi-tenancy - The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources



but may be able to specify location at a higher level of abstraction (e.g., country, state). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.

Rapid provisioning, scalability and elasticity - Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear unlimited and can be purchased in any quantity at any time.

Accounting - Cloud systems automatically control and optimise resources use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g. storage, processing, bandwidth, active user accounts, etc.). Resource usage can be accounted for, monitored, controlled, and reported providing transparency for both the provider and consumer of the utilised service. Payments are then associated with actual usage.

3.5 Benefits and Issues for EGI

Commercially available public clouds have been designed to satisfy general computing requirements such as e-commerce and transactional communications that are typically less sensitive to bandwidth and latency. As clouds become more mature, however, it is anticipated that clouds of different “flavours” will be deployed to meet the requirements of different user communities such as those that are currently dependent on EGI (e.g., research computing). Therefore, while all of the potential benefits and issues of general cloud computing are relevant to the research computing community, their needs will not always be met by commercial cloud providers. The notion of science clouds will force an emphasis on specific benefits and issues for these user communities that are not provided or available commercially.

The question now for EGI is to understand how the adoption of virtualisation technology within its current infrastructure composed of federated resource providers should deliver a cloud computing environment for its users, and how it implements such an environment. There is certainly interest in the current and potentially new user communities in exploiting cloud computing resources.

Anecdotal experience over the last few years within various communities will be supplemented by the VENUS-C project in the next two years as it attempts to further understand the suitability of certain applications for cloud computing (See Section 4.3). Within EGI Resource Infrastructure Providers, they have also been exploring the deployment of dynamic execution environments (Section 4.1) and the provisioning of grid sites in cloud infrastructures (Section 4.2).

The following table summarises the overall benefits and current issues of cloud computing from an EGI perspective [R29; R32; R35].

Benefits	Issues
Clear business models	Accounting (<i>significant additional technical work for cloud metering services</i>)
Commoditisation of compute capability	Application redesign may be needed to exploit full potential of the new cloud services
Data centre / resource consolidation potential	Maintaining compliance with public regulations and internal IT policies relating to data harder (<i>*issue remains in certain user communities regardless of grid and cloud</i>)
Ease of application deployment	Data Access, Portability and Interoperability between clouds
Efficient energy usage (Green IT)	Software Licensing
Identity and Federation management	Performance Management: Abstraction vs. Control (<i>Virtualisation layer and beyond – e.g., network and storage</i>)
Improved reliability	Security (<i>Loss of ownership, control, availability, guarantees and 100% user responsibility</i>)
Improved server utilisation	Service Level Agreements
Managing surge requirements with on-demand resources.	Trust (<i>system admins' reluctance to allow users to run their own VMs on the infrastructure</i>)
Virtual ownership of resources (change ratio from CapEx to OpEx)	

Table 1: Cloud Computing Benefits and Issues

4 CURRENT COMMUNITY CLOUD ACTIVITIES

Within the EGI Community, the convergence of Grid and Cloud Computing (i.e. the convergence of federated distributed resources with greater on-demand elasticity of user defined environments) is seen as providing many potential benefits. Currently, there is no architecture or roadmap for such convergence, though efforts are being made within the community to understand the critical issues. This section, which builds on the concepts described previously, highlighting a non-exhaustive list, as there are a number of initiatives and efforts through the community, with summaries and the features being examined.

Focusing on the EGI context, these are a number of technical issues being explored:

- **Provisioning** - Cloud-technology can make it possible to create virtual Grid sites on any resources. An entire site can be virtualised, running all basic Grid services in the Cloud, thereby improving service availability and giving providers more flexibility in how they deliver these services to their user communities.
- **Dynamic Execution Environments** - Many data analysis tasks result in an application processing a data file through a batch processing queue. Different user communities need different applications and different environments, which if installed directly onto a physical machine are hard to integrate together, or result in tying particular applications to a particular hardware if deployed separately. Instead ‘virtual machines’ can be run on demand to meet the specific needs of certain applications or jobs. One piece of hardware can then run several operating systems simultaneously (within a ‘hypervisor’) giving the user communities and resource providers much more flexibility.
- **Scale Out: A ‘hybrid’ solution** - Combining private virtualised infrastructure used to host a whole grid site or a dynamic execution environment with public cloud resources to expand the computing capacity at a site.
- **Application Suitability:** Potentially not all applications are “cloudable” and it is important for the user community and the resource providers to understand which applications can be adapted to this environment.

4.1 Dynamic Execution Environments

4.1.1 Worker Nodes on Demand Service

The Worker Nodes on Demand Service (WNoDeS) [R37] is an INFN-developed architecture, which makes it possible to dynamically allocate virtual resources out of a common resource pool. It aims to expand and exploit existing infrastructures (e.g. EGI) through sharing and virtualisation. There are no resources specifically allocated for virtualisation, but each worker node in the pool can run a regular grid job, or a virtual machine.

WNoDeS is software, also developed by INFN, which builds around a tight integration with a LRMS (batch system), a virtualisation infrastructure (KVM) and the WNoDeS framework itself. It permits a full integration with existing computing resource scheduling, monitoring, accounting, and security



workflow. It provides on-demand virtual resources that are worker nodes, but also VLANs, to dynamically isolate virtual machines according to users' requests.

WNoDeS has been in production mode at the INFN WLCG Tier-1 Centre since November 2009. It currently exposes both a grid entry point, allowing distributed submissions to be run on user-specified VMs, and a local jobs submission interface. It also has preproduction solutions for an integrated cloud computing web interface. Generally, there are from 1200 to 2000 VMs running at the INFN Tier-1, with plans to extend the virtualisation framework to all 8k Tier-1 cores. WNoDeS has also been installed at an Italian Tier-2 site.

The big virtual cluster composed by real and virtual nodes strains the components involved in the framework, like the LRMS and the network file system, addressing a number of scaling issues, where solutions sometimes involve the technology providers directly. WNoDeS provides different entry points to the virtualisation infrastructure:

- gLite grid interface: WNoDeS enables the possibility for grid users to select at job submission time and the virtual image that will be used for the instantiation of the virtual worker node that will execute the job by reusing the current grid interfaces.
- Open Cloud Computing Interface (OCCI): Defined by the Open Grid Forum this interface is still being implemented in parallel with a web application that provides a more user-friendly experience. The OCCI layer supports the same authentication and authorisation technologies used by the grid infrastructure.
- Local job submission: Local batch jobs can be run on both virtual and real execution hosts. WNoDeS offers the same virtualisation framework to the local users who usually do not use grid interfaces, providing direct access to the batch system.

4.1.2 Batch and Server Virtualisation and Cloud Integration

In 2009, CERN [R39], the European particle physics organisation that runs the Large Hadron Collider, started to develop an Infrastructure as a Service (IaaS) setup. Since then, significant progress has been made in the implementation of the new system. In spring 2010, about 500 recent batch worker nodes were added temporarily to the system, which allowed large-scale tests of the new infrastructure. The batch computing farm, which makes a critical part of the CERN data centre, can now use this IaaS model to provision a large number of virtual batch worker nodes. By making use of the new equipment, both the virtual machine provisioning systems and the batch application itself have been tested extensively at large scale. This has demonstrated that the system can sustain 15,000 or more concurrent virtual batch worker nodes.

CERN has also embraced server virtualisation and cloud computing technology to improve CPU utilisation and the delivery of computing resources to scientists around the world. CERN, which uses Red Hat's version of the Xen hypervisor as well as Microsoft's Hyper-V, has recently installed private cloud software from Platform Computing to automate the process of managing the virtual infrastructure.



Platform has provided its Platform LSF software as a “private cloud” tool that aggregates servers, storage, networking tools and hypervisors to create a shared pool of physical and virtual resources. An announcement from Platform Computing credits the software with helping CERN build “the world’s largest cloud computing environment for scientific collaboration” [R13].

So far, CERN is running a few hundred VMs on the Intel-based x86 servers that make up its batch environment, which serves the scientific community. CERN could potentially have 60,000 or more VMs running batch jobs in the future, however, they want to aggressively move batch jobs to VMs over the next year hoping to improve system utilisation by about 15% or 20%, but that depends partly upon user acceptance.

4.2 Provisioning Grids in Clouds

4.2.1 Grid on Demand

The Grid on Demand Project [R22] was carried out within the System- and Network-Engineering department at the University of Amsterdam [R40] focusing on one question: Can grid computing be offered as a cloud service? Cloud compute services are seemingly provided without limits and promises an almost infinite number of resources that can be added and removed dynamically. The grid has traditionally offered a dedicated computing platform for compute intensive scientific (e-Science) applications. The answer to the question posed is carried out by combining the properties of both grid and cloud for the purpose of supporting current or newly developed e-Science applications with a sudden demand for compute power. The goal of this work was to use the elasticity and scalability of cloud computing (IaaS) while providing the abstraction of a grid interface on top of the virtualisation of cloud.

The implementation was realised by extending an existing Amazon Machine Image, containing the Ubuntu Lucid Linux operating system, with Torque Resource Manager and Globus Toolkit. To test the performance of Grid on Demand, a comparison needed to be made with a real cluster having a grid interface. An actual e-Science workload generated a representative load on both a real cluster and Grid on Demand. When the test was performed using an existing grid application utilising the grid interface (Globus Toolkit) it showed that Grid on Demand could be used in existing environments without modification, though only the minimum grid services were applied (Resource and Connectivity Layer) delivering the most generic grid resource that was not configured towards specific usage.

An interesting result from the study was the workload execution time. On the local cluster, it was almost equal to the job execution time when the number of jobs in the workload was less than the number of CPU’s available. The workload execution time in the local cluster was 4 times longer than the job execution time when the number of jobs is more than 3 times the number of available CPU’s. Grid on Demand scales with the number of jobs and the total workload execution time of roughly 100 jobs is close to the workload of 30. The difference between the 100 and the 30 job workload is still significant and is due to the sudden increase of the pending time around job number 70



suggesting that the cloud provided a more consistent resource offering a lower job execution time as the number of jobs increases.

Though further investigation needs to be carried out to test variations, the conclusions of the research project are that cloud resources can be leveraged to augment the grid. It demonstrated how e-Science applications could use Grid on Demand when there is a high demand for resources for a short period of time, as an elastically scalable solution.

4.2.2 RESERVOIR

RESERVOIR [R12], “Resources and Services Virtualisation without Barriers” was a three-year project, partially funded under the European Commission’s Seventh Framework Programme, which ended in January 2011. The RESERVOIR consortium, led by IBM, with thirteen leading industrial, research and academic partners from across Europe used requirements derived from use cases brought by industrial partners in the project, which cover e-Government, utility computing, business computing, and telco applications. Its main objective was to seamlessly enable deployment and management of complex IT services across distributed administrative domains and geographies.

The emerging model of cloud computing is characterised by elastic and location-independent resource pooling typically hosted in large data centres, which may have tens or even hundreds of thousands of physical machines. The RESERVOIR approach, however, contended that no single compute cloud could be large enough to meet rapidly scaling demands on its infrastructure without having to expensively overprovision its physical infrastructure. RESERVOIR’s research had focused on solving this problem by enabling the migration of virtualised resources across federated clouds, while guaranteeing security, and meeting Quality of Service (QoS) requirements.

RESERVOIR demonstrated the ability to create an infrastructure that allows for live migration of virtual machines, moving to physical hosts, which may not share common storage, or may reside on different subnets or even different clouds.

In addition to its research goals, another aim of the project was to create technologies that could be exploited by the European community to build an infrastructure for a cost-competitive, service-based online economy by merging virtualisation and business management technologies. These results are available in the form of the RESERVOIR Framework, which is downloadable from the RESERVOIR website [R12]. This framework groups all the open source software, and the detailed specifications of the proprietary code that are necessary to help the user build a RESERVOIR cloud. RESERVOIR supplies the architecture for a service-oriented infrastructure, built on open standards and new technologies. The architecture is composed of three main layers, with functionality such as security and a “messaging bus” cutting across all layers: Virtual Execution Environment (VEE), VEE Management layer (VEEM), and a Service Management layer.

Professional integration, certification, and technical support that many enterprise IT shops require for internal adoption is now available through a commercial organisation, C12G Labs. The new



company contributes to the OpenNebula project and allows it not to be tied exclusively to public financing (research grants, etc.), contributing to its long-term sustainability.

4.2.3 StratusLab

StratusLab [R41] is developing and deploying cloud technologies with the aim of simplifying and optimising the use and operation of distributed computing infrastructures such as the EGI. The target users run from systems administrators and technicians, to community service administrators and researchers. StratusLab expects administrators to install a StratusLab cloud on their physical infrastructure, and then to install grid services in this cloud. Once the cloud layer is in place, the opportunities arise to grant cloud access to community service administrators, software engineers and researchers to deploy VM-based appliances and services to meet their specific needs.

The project integrates, distributes and maintains a sustainable open-source StratusLab cloud distribution to bring cloud to existing and new grid sites. The StratusLab distribution is based on existing cutting-edge open source software, such as OpenNebula and Claudia, with additional features, innovative services and cloud management technologies developed in the project. The developers and integrators incrementally deliver a production grade distribution that is being demonstrated through the operation of production-level grid sites during the project.

StratusLab is a two-phase project, with two major software releases scheduled in May 2011 and 2012. In addition, the project operates a six-week continuous release cycle, to deliver incremental improvements and additional features on a regular basis. Development is based on an agile process that allows the developers to react to changes, requirements and opportunities identified through interactions with users and other projects including EGI.

In the first phase, the project focuses on cloud computing for resource provisioning in grid sites. This entails development of the StratusLab cloud platform and creation of virtual appliances for the scientific application domains in the project.

The StratusLab infrastructure will also serve as an important platform for assessing the economical impact of cloud technologies in the provision of grid services both in terms of human resources (e.g. for administration and system maintenance) and environmental costs (power consumption, carbon footprint, etc.). Alongside the reference infrastructure, StratusLab has been hosting a public appliance repository for virtual machine images in advance of its first release. In the roadmap for release 1.0, the development of an *appliance marketplace* will be designed to meet the requirements of StratusLab users and the HEPiX Virtualisation Working Group, with an eye on EGI plans in this area.

4.3 Application Suitability

4.3.1 Venus-C

VENUS-C (Virtual multidisciplinary EnviroNments Using Cloud Infrastructures) [R14], a project funded under FP7, brings together industrial partners and scientific user communities. Its aim is to develop and deploy a cloud computing service for research and industry communities in Europe by offering



an industrial-quality, service-oriented platform based on virtualisation technologies facilitating a range of research fields through easy deployment of end-user services. Current user communities involved are: bioinformatics, systems biology, drug discovery, civil engineering, civil protection and emergencies and data for science.

The VENUS-C solution is an open and generic Application Programming Interface (API) at platform level for scientific applications, striving towards interoperable services. The VENUS-C platform will be based on both commercial and open source solutions supported by the Engineering data centre, Microsoft Azure and its European data centres, along with two European High Performance Computing centres, The Royal Institute of Technology (KTH, Sweden) and the Barcelona Supercomputing Center (BSC, Spain).

Azure offers a multi-layer solution, including computing and storage power, a development environment and immediate services, together with a wide range of services that can be consumed from either on-premise environments or the Internet. From an open source perspective, the Eucalyptus and OpenNebula solutions are being evaluated, while the Emotive middleware for clouds is offered by the Barcelona Supercomputing Centre, thus demonstrating interoperability and ultimately portability to VENUS-C users.

The main output from VENUS-C will be a series of user scenarios showing how the cloud computing model can benefit different scientific communities. VENUS-C will expand the supported communities by means of an open call for up to twenty short experiments to exploit the VENUS-C cloud platform through the cloud resources provided within the project. The first call will be open until 11 April 2011 aiming to extend the current user scenario portfolio and enable a new generation of research applications to validate the infrastructure for advancing scientific discovery.

4.4 Summary

e-Infrastructures and the innovative technologies that power them and the demanding researchers that use them are a strong and ever-present mechanism enabling researchers, developers, and technology and resource providers to all work toward a common goal. It is essential EGI.eu uses every means available, on behalf of the community, to communicate and collaborate in these strategic areas (i.e. establishing MoUs, etc.) in order to answer the needs of the current users and to continue building new communities. Forming and maintaining these communities will be how the EGI and its stakeholders will survive, thrive, and evolve.

5 EVOLVING EGI

EGI needs to expand its resource infrastructure of compute and data resources to include new types of resources (e.g. desktop grids, virtualisation and high performance computing) in response to its current and new user communities. It has a process to collect and prioritise requirements from a multi-disciplinary user community to drive its development. Many of these requirements relate to user communities needing more flexible on-demand access to resources and a greater range of environments and services to those currently provided.

The adoption of virtualisation technologies within EGI could evolve the infrastructure from a collection of relatively isolated systems to a virtualised fabric of resources. These will enable it to meet current and new user requirements, and provide the opportunity for the optimisation and delivery of an infrastructure of platforms and of services by specialist providers – either academic or commercial – those best able to provide them most efficiently. Such an ecosystem of providers can achieve extreme economies through specialisation and scale in particular aspects required by the community as a whole. Providers that specialise by only offering a limited portfolio of services are able to differentiate themselves by minimising their internal diversity and their management cost to provide a commodity to other users [R15].

The remainder of this section examines the drivers guiding the evolution of the infrastructure towards virtualisation due to its public funding, organisational structure and economics. The costs in providing the current EGI resources cloud providers are analysed to explore the economic issues in adopting virtualisation. The section concludes with EGI’s vision for integrating virtualisation into EGI.

5.1 Drivers

5.1.1 Public and European

European e-Infrastructure is publicly funded either by national or European level funds. The Digital Agenda for Europe, which more generally identifies a number of issues for ICT provision in the public sector at a European level, has three key areas for EGI: Borderless Services, Standards and Interoperability, and Innovation.



Figure 4: Competitiveness Cycle



5.1.1.1 Borderless Services

The Internet is borderless, but online markets (the resource infrastructure providers within EGI), both globally and in the EU, are still separated by multiple barriers that inhibit collaboration.

Cloud computing, while removing many of the barriers to accessing resources, offers a separate set of issues. Removing “borders” or crossing national or continental territories means that the physical location of your data becomes a concern. If a dispute arises, what will be the place of jurisdiction? Other issues, such as responsibility of data, liability coverage for breach of privacy such as the data centre getting hacked, intellectual property rights, third party access, etc. follow on from this concern.

Also, many of the problems are humanistic. Individual countries are concerned with safeguarding national sovereignty in order to conserve knowledge and technological competence as well as protecting data privacy and sensitive industrial information - fear of losing jobs, as developing locally based IT infrastructure will avoid workers having to relocate elsewhere and avoiding the under-utilisation of existing local data centres and rendering them obsolete [R16].

EGI has historically overcome some of these issues through the spirit of collaboration and minimalistic policies governing usage, accounting and authentication.

5.1.1.2 Standards and Interoperability

To achieve the portability, interoperability, and economies of scale that clouds offer, it is clear that common design principles must be widely adopted in both the user community and marketplace. To this end, a private-to-public cloud deployment trajectory will be very common, if not dominant. The current market of a few resource and technology providers raises concerns about technology lock-ins that have persisted previously in the technology community.

This trajectory can be used to define a progression of needed common practices and standards, which in turn, can be used to define deployment, development and fundamental research agendas. The cloud standards landscape and the standards process should be driven by major stakeholders (e.g., large user groups, vendors, and governments) to achieve scientific and national objectives. It is therefore necessary that stakeholders actively engage in driving this process to a successful conclusion.

There are different clouds from companies such as Microsoft, Amazon, IBM, and Google, but with an evident lack of interoperability between them. Interoperability has not been a huge focus around the cloud computing space, other than general statement of “support” from the larger cloud computing providers without specific detailed plans. There has been a slight push through Standard Development Organisations (SDOs) with increasing participation by industry such as in the Distributed Management Task Force (DMTF) Open Cloud Standards Incubator and Open Virtualisation Format [R42]. Other dedicated groups comprise the Open Grid Forum (OGF) Open Cloud Computing Interface Working Group [R43], though overall, though not much as resulted tangibly or enough to shake up the market.



Data interoperability is a little more difficult, which deals with a few key concepts, such as semantic interoperability (i.e. the way that data is defined and stored on one cloud versus another). Also, another consideration is that of transformation and translation, so the data appears native when it arrives at the target cloud, or clouds, from the source cloud (or clouds) among other data issues previously mentioned – data governance and data security [R17].

5.1.1.3 Innovation

In today's economy, it is clearer than ever that Information and Communication Technology (ICT) is the most important driver of innovation and competitiveness. In addition to ICT, other key enabling technologies are revolutionising the products and services on offer as well as the way business is conducted in Europe and this revolution will continue in the future. The European Commission is trying to make sure innovation is thoroughly understood and approached comprehensively, thereby contributing to greater competitiveness, sustainability and job creation through formulating, influencing and, where appropriate, implementing policies and programmes to increase Europe's innovativeness [R18].

In regards to ICT, specifically e-Infrastructures, EGI.eu, on behalf of the EGI community, has positioned itself at the forefront for innovation support through the deployment of technological innovation in a dynamic environment such as distributed computing and continues to push the boundaries as technology evolves to meet the needs of its user communities. Being able to expand the number and size of the supported user communities will require new technologies such as the emerging stabilisation of cloud computing and virtualisation and others such as desktop grids.

EGI.eu has also enabled software innovation in order to provide a reliable persistent technology platform with tools and services built on middleware extending past gLite, into UNICORE, ARC and Globus. It also supports research innovation by providing a stable infrastructure for data driven research as well as opening up new opportunities for international research such as European Strategy Forum on Research Infrastructures (ESFRI) [R44].

5.1.2 Organisational

In some niche sectors over the last decade, like engineering, banking and life sciences, federations of distributed computing resources (a.k.a. grids) between organisational units within an enterprise (i.e. departments within a company or research groups within a collaboration) are already well established. Increasing needs for data storage and computing resource in many different areas represent great opportunities for federation allowing local resources to be used by remote users when demand is low.

From a technology and business point of view, grids have enabled organisations to use their resources more efficiently, by supporting large scale processing on demand, through consistent access to shared resources and data no matter what the user's location, thereby empowering distributed user communities. However, the technology approaches taken within the grid community



in the last decade have not lowered the barriers to adoption sufficiently, despite the maturing of the underlying approaches, to grow the user communities.

The emergence of commodity virtualisation, which underpins all cloud computing activity, provides the route by which the large-scale, on-demand delivery of storage and computing, that was the original vision of the Grid, might now be delivered.

5.1.2.1 Efficiency

Today, business user communities are now widely adopting the “as a Service” approach. The actual objective behind this is minimising costs and therefore outsourcing everything that is not in the core business of the company (e.g. computing power, storage, services, applications, etc.) or does not add value.

A similar opportunity now exists within EGI. The utilisation of infrastructure, platforms, and software “as a Service” provided from outside the EGI Community as part of the regular production infrastructure either paid for on demand or as a regular contract, would be enabled with a move to virtualisation. For some usage patterns, such a model will be indistinguishable to the end-user from using any other resource within EGI, and may be delivered with greater efficiency. The cloud business model allows end-users or resource infrastructure providers who are not willing to own and manage their own resources but need to use or deliver data and computational power. Such providers may come from other NGIs, through commercial providers, or other organisations within the community.

5.1.2.2 Consolidation

The pressures (e.g. staffing costs, green energy, economies of scale, etc.) that produced the consolidation of data centres and wide-scale adoption of virtualisation in the commercial sector are beginning to be felt in the academic and research sector. People have invested in local infrastructures (hardware and personnel) because they offer guaranteed capacity and instant access. Consolidation may occur unless the Infrastructure can provide both. However, many campuses are encouraging the move of departmental or group level computing resources into central locations where they can be managed and supported by dedicated staff.

These consolidation pressures may continue beyond the campus to a regional, national or European level if the public sector wishes to deliver economic efficiencies comparable to commercial providers for similar resources. Clouds indeed are a potential solution, but must be deployed with this in mind.

5.1.2.3 Security

While security and reliability are often cited as potential hurdles to public cloud adoption, the increased need for them leads to the level of investment required to achieve operational security and reliability. Deploying virtualisation within private or public cloud computing environments calls for a strategy to ensure a secure move into this complex and dynamic model. Legacy security solutions for a physical data centre tend to impede adoption of these technologies because they are not virtualisation or cloud-aware, among others. Large commercial cloud providers often bring deep expertise to bear on this problem than typical corporate IT departments, thus actually making cloud systems more secure and reliable.

However, reliability and security will likely continue to improve as public clouds are still in a relatively early stage of development. This has already been shown in areas such as public cloud email, which are generally more reliable than most on-premise implementations within the bounds provided by the terms and conditions of use, which may provide no guarantees towards data locality, ownership and availability. Security issues around data and applications normally occur when systems are out of date. Within PaaS, the automatic patching and updating of cloud systems greatly minimises this. Currently, there are no fundamental reasons why public clouds would be less secure, but are actually likely to become more secure due to strict security policies commercial providers must enforce as well as the level of expertise they bring [R5].

5.1.3 Economic

5.1.3.1 Cost of Power

Electricity costs are rapidly rising to become one of the largest elements of total cost of ownership, currently representing 15%-20%. Power Usage Effectiveness, a measure of how efficiently a computer data centre uses its power [R6], tends to be significantly lower in large facilities than in smaller ones. While the operators of small data centres must pay the prevailing local rate for electricity, large providers can pay less than one-fourth of the national average rate by locating its data centres in locations with inexpensive electricity supply and through bulk purchase agreements further reducing energy costs.

5.1.3.2 Infrastructure Labour Costs

While cloud computing significantly lowers labour costs at any scale by automating many repetitive management tasks, larger facilities are able to lower them further than smaller ones. While a single system administrator can service approximately 140 servers in a traditional enterprise, in a cloud data centre the same administrator can service thousands of servers. This allows IT employees to focus on higher value-add activities like building new capabilities and working through the long queue of user requests with which every IT department contends.

Currently, many institutes and universities are using PhD students as network administrators, but implementing a virtual layer, would not necessarily remove a position, but could streamline human resources to more dedicated or productive roles. This potential new role would open up not only the need for, but also the opportunity for developing new skills, as many things are different in a virtual environment. Some of these areas are highlighted in Table 5.

5.1.3.3 Buying Power

Operators of large data centres can get discounts on hardware purchases of up to 30% over smaller buyers. This is enabled by standardising on a limited number of hardware and software architectures. For the majority of the mainframe era, more than 10 different architectures coexisted. Large-scale buying power is difficult in this heterogeneous environment. With cloud, infrastructure homogeneity enables scale economies. Leveraging EGI's buying power and influence through the collection of requirements from and for the EGI community could provide significant cost savings.



5.2 Cost Analysis

5.2.1 Scenarios

Adoption of cloud computing platforms and services by the general scientific community is still in its infancy as the performance and monetary cost-benefits for scientific applications are not perfectly clear. For EGI, virtualisation offers the opportunity to deploy different software environments on demand, customised to the needs of the individual user communities. These virtualised resources could be provided from within the existing network of resource infrastructure providers or through commercial providers. This section provides some financial estimates of the cost of using cloud providers alongside the existing infrastructure, assuming that the functional requirements from the user communities within EGI, or a subset of them, could be provided without any reduction in capability.

5.2.2 Cost Estimate: EGI-InSPIRE

The total cost of delivering an integrated pan-European production infrastructure is estimated by the EGI-InSPIRE consortium as €335M over four years. The activity described in the project's description of work [R10] relates to less than 25% of the total effort provided across Europe by its partners. The European Commission's (EC) contribution of €25M is therefore less than 10% of the overall investment being made by the NGIs and EIROs. EGI can only exist by coupling the considerable existing national investments being made by the partners, with the investment from the EC, in order to provide the European-level coordination and governance necessary to accelerate the integration of these independent national activities. The staff effort within EGI and costs around providing the production infrastructure are broken down in Table 2 below.

Cost	Description	Notes	Value
Cost A	EGI Global Tasks within EGI-InSPIRE	Average staff effort 44 FTE	€17.2M
	NGI International Tasks within EGI-InSPIRE	Average staff effort 113 FTE	€45.1M
	General Tasks within EGI-InSPIRE	Average staff effort 17 FTE	€6.7M
Cost B	Additional effort within the NGIs for the International Tasks	Estimated staff effort 224 FTE. Calculated from the effort recorded in the EGI_DS Functions that is not directly supported within the EGI-InSPIRE project.	€89.1M
Cost C	Additional effort to support the internal NGI/EIRO activities	Estimated staff effort to operate the NGI/EIRO infrastructure – 300 FTE	€111.2M
Cost D	Hardware costs	€10.7M a year for replacing the current compute clusters every 3 years.	€32.1
Cost E	Running costs	Annual electricity costs of €9.3M are based on 93,000 CPUs totalling 170,000 cores	€37.2M
TOTAL			€338.6M

Table 2: Estimated Total Cost of the EGI production infrastructure during EGI-InSPIRE



5.2.3 Cost Estimate: e-IRGSP2

The closest cost comparison of grid and cloud available to date based on actual figures has been through the e-Infrastructure Reflection Group Support Programme (e-IRGSP2) [R45]. This financial exercise examined comparable elements of European e-Infrastructure to Amazon EC2 in order to better understand the different costs relating to the two service models [R24]. The final research paper is being drafted as part of an e-IRGSP2 deliverable 4.3b to the EC, “Final Legal Issues Report” [R33].

The cost estimation referred to 2009 and corresponded to the yearly cost of EGI including both capital expenses (i.e. depreciation of CPUs, storage and auxiliary equipment, etc.) and operating expenses (i.e. personnel costs, software, electricity costs and premises cost, etc.). The yearly EGI cost had been estimated within the range of €55M-€118M. Overall, calculations were based on information gathered through a detailed questionnaire completed by seven NGIs that participated in the study, complementary commercial and industry data and references found in literature.

The CPU core hour cost without storage depreciation was compared to Amazon EC2 offers. Storage depreciation was excluded for the comparison as Amazon sells storage services separately. Except for the Amazon “standard small instances” and “micro instances”, EGI cost per CPU hour seemed to be less costly. However, small and micro instance configurations (performance and memory) seem to be less advanced than an average grid computing node, therefore they are not directly comparable. In the calculations, the CPU core hour cost without storage depreciation ranged from €0.0569/CPU core hour (90% utilisation) to €0.1356/CPU core hour (60% utilisation). The Amazon equivalent ranged from €0.0899/CPU core hour (90% utilisation) to €0.1033/CPU core hour (60% utilisation).

The output of the financial exercise cannot be suitable for answering the question about what would be the cost if cloud computing was to replace either fully or partially the grid, but serves as a reference point for moving forward. A different, more comprehensive analysis is needed in order to properly address this issue.

5.2.4 Cost Estimate: Amazon Web Services (EGI-InSPIRE Proposal)

Table 3 below offers a general concept cost analysis of running the some 93,000 nodes and 170,000 cores of EGI and 100 PB of storage onto Amazon using the online tool “AWS Simple Monthly Calculator” [R23] equivalent to the infrastructure described and costed in the EGI-InSPIRE proposal. Both high-end and low-end figures are provided for standard instance types using EC2 at 70% usage and transferring data into Simple Storage Service (S3).

Type	Amt.	Description	Cost	Assessment
EC2 Small Instance	170,000	<ul style="list-style-type: none"> • 1.7 GB memory • 1 EC2 Compute Unit (1 virtual core w/ 1 EC2 Compute Unit) • 160 GB instance storage • 32-bit platform • I/O Performance: Moderate 	14.2M/yr.	Assuming each core is one compute unit, which is normally not the case within Grid. EC2 small instance offers less computing power than the current EGI ¹
EC2 Large Instance	85,000	<ul style="list-style-type: none"> • 7.5 GB memory • 4 EC2 Compute Units (2 virtual cores w/ 2 EC2 Compute Units ea.) • 850 GB instance storage • 64-bit platform • I/O Performance: High 	25.3M/yr.	Grid jobs normally run multiple cores, but EC2 large instance offers more computing power than the current EGI ²
S3 - Storage	100 PB	Interface to store and retrieve data	71M/yr.	For simply hosting data
S3 - Data Transfer In	100 PB	One time move of data	10.4M/yr.	Based on moving 100 PB divided into 12 parts (months) and does not consider moving any data out which has a considerable additional cost
TOTAL		Cores+Storage+Data Transfer	95M-107M/yr.	Converting these numbers into 4 years (length of EGI) costs range between €383-€427M compared to €69.3M of only EGI-InSPIRE's est. hardware and running costs and almost 25% more of total costs

Table 3: Cost of moving EGI to Amazon

A variety of factors have not been taken into consideration, such as the cost of networking within EGI, nor can the cost of systems staff (implicitly included in the Amazon model) can be accurately attributed in the EGI model.

5.2.5 Summary

Considering the baseline provided by the EGI-InSPIRE project (Section 5.2.2) two direct comparisons are provided:

- E-IRGSP2 (Section 5.2.3): Considering only direct compute costs, the cloud is more expensive than grids. However, the cost of networking and overheads (e.g. energy & staff) are ignored.
- Amazon (Section 5.2.4): Examining the pure computational elements of EGI with Amazon shows that the cost becomes broadly comparable – €89.2M (avg. Amazon EC2 costs) vs. €69.3M (EGI) – especially if only a modest proportion (e.g. 10%) of the staff C costs (€111.2M) are attributed to system administrators.

¹ EGI average core: 2200 SI2000 -> EC2 small instance: 1700 SI2000 77%

² EGI average core: 2200 SI2000 -> EC2 large instance: 3400 SI2000 150%



A very simple comparison of networking costs has €41.6M for moving all the stored data in/out of the Amazon each year while the 4 year EC contribution to GEANT3 (which also has considerable national co-funding) is €93M, which supports other user communities beyond EGI.

While the technical suitability of cloud computing may vary from application to application, it is also clear that the economic benefits will also be dependent on the application and the data it stores, and transfers to its cloud based resources. Any serious financial discussion requires better and more transparent costs of the various infrastructure components as they are used by a particular application or community. Even if it is cheaper to host an application in-house, the application may have peaks in demand that can be handled by external resource bursts into a public cloud if no federated resources are available.

Items not included in Calculation	
Additional Costs	Potential Savings
Overhead costs	Negotiation with cloud provider for discount/reduction based on volume
Grid operations and support will still exist, even if streamlined or reduced, therefore costs need to be added on top of cloud costs	Consolidation of sites (reduce personnel / overhead costs, energy, hardware costs)
Network costs (Data transferred out)	Linking GEANT to commercial cloud providers for research community data transfers
Administrative and contracting fees with cloud providers.	

Table 4: Additional Costs & Savings to EGI/AWS Cost Comparison

Cost comparisons between academic and commercial offerings are always challenging as presented, but also as cloud providers' have the ability to pass savings directly and efficiently on to users as operating costs are reduced.

Moreover, the most important conclusion moving forward is that even if the costs are equal, or have a small margin either way, there are structural, organisational, and political barriers to completing the outsourcing of these infrastructures and federation of existing resources still remains the best short-term solution for increased efficiency and reduced costs.

5.3 A Vision for Integrating Virtualisation in EGI

The commoditisation of hardware, software, and networking over the last decade has fuelled the establishment and expansion of the infrastructure that was inherited by EGI. The next few years will see the impact of commoditised virtualisation within the provision of resources for the European Research Area providing improved efficiencies both in terms of human resources and in their energy footprint. The changes that have been already been seen in commercial data centres for transactional workloads and have led to the cloud computing business model will inevitably impact the way data centres in the research community are provisioned.

Therefore, the impact of the whole scale adoption of virtualisation in the data centre will benefit that organisation. To the production infrastructure as a whole it offers the opportunity of deploying different software environments, customised to the needs of individual user communities, with minimal overhead to the individual resource centre and doing so on demand from the end-user.

As a consequence of this change, resource centres in EGI would be able to support the different service environments required by the increasingly diverse application communities using the production infrastructure. Introducing a virtualisation layer would move the software deployment decisions away from the sites and back into the virtual organisations using the infrastructure.

5.3.1 Architecture

Providing secured, authorised, and accounted mechanisms across Europe for starting virtual machines on remote sites is in many ways no different from the currently agreed-upon procedures for starting jobs on remote sites, providing appropriate policy and certification models. Virtual organisation managers, or operations staff acting on their behalf, would prepare, deploy, and monitor the software required by that application domain. Site administrators would still retain control over which virtual organisations access the resources and the quantity of local resources (compute and storage) allocated to each. Communities could choose to deploy the middleware that they currently use (gLite, UNICORE, ARC or dCache) into the virtual machines run on the infrastructure, or choose to deploy software services coming from within their own communities.

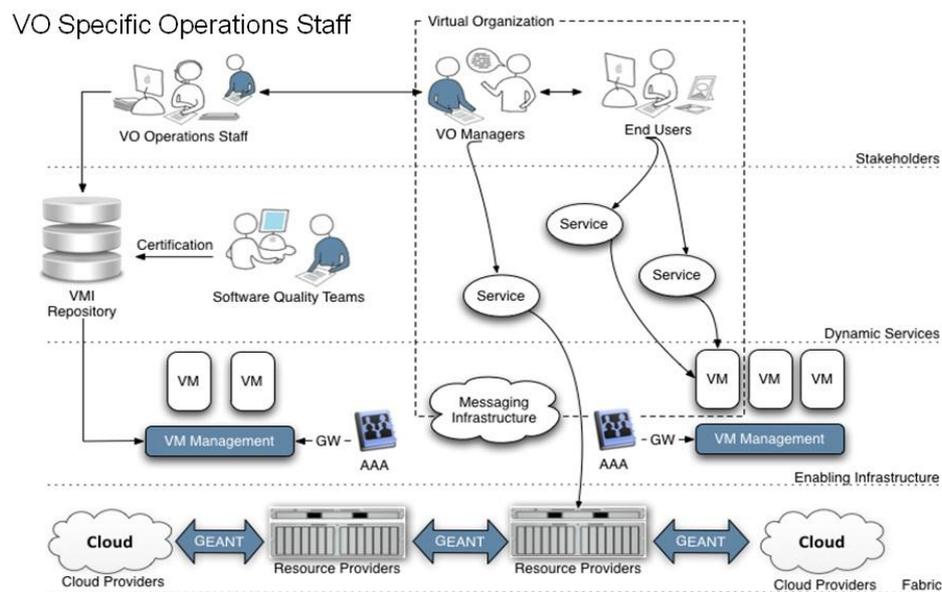


Figure 5: A Virtualised Ecosystem

Through the virtualisation layer deployed on each site, different virtual organisations would be able to deploy the software needed by their community at an update cycle appropriate to their own work. The workload produced within the European Research Area is primarily based around data. The high performance research networks around Europe enable the rapid movement of data between sites — many of which have the ability to store many petabytes of data. This capability and



associated cost is distinct from those of many commercial cloud providers, mainly because the business model within the research networks makes this usage free to the end-user [R19].

However, for an activity that is more computing- rather than data- focused, virtualised architecture provides a bridge to commercial cloud providers as additional VMs that can be deployed into a commercial cloud. Also, their services can be integrated into the broader infrastructure available for a VO. Access to such facilities could be facilitated on a site-by-site basis or across the production infrastructure as a whole.

This architecture is discussed in greater detail in the DCI Collaborative Roadmap [R20]. The report provides a vision for developing a pan-European production infrastructure built from federated distributed resources.

5.3.2 A Virtualised Ecosystem

The move within EGI to use external technology providers has decoupled the provider of services to end-users (i.e. EGI) from being the developer of these end-user services (i.e. the technology providers). Virtualisation offers the opportunity of also decoupling the deployment and operation of end-user services (i.e. the services end-users interact with) from the site to the user community and enabling resource providers outside of the EGI Community to provide services (e.g. commercial cloud providers). The development of an EGI ecosystem around the adoption of virtualisation provides more flexibility in the provision of these individual elements, and therefore the sustainability of the ecosystem as a whole. The major difference between the current operating model and one with virtualisation/cloud shifts the responsibility from resource providers to community operations staff.

While the use of virtualisation by resource providers can be transparent to end-users, unless it is exposed to them as a capability they (or their community) can use, it provides no direct benefits. Commercial cloud providers are able to offer that capability and user communities have been able to assess the effectiveness of this technology for their applications and the resulting cost. However, like any technology, it will not work for all applications, and the current service providers will not provide all the desired functionality. In particular, aspects of collaboration, result sharing in virtual organisations, and many of the more complex data management aspects are not covered. This provides an opportunity for service providers in EGI to offer a cloud oriented service, with the collaboration and data management aspects its users require, where no commercial service exists, and to use commercial service providers where it is technologically and economically beneficial to do so.

Such a mixed (research and commercial) service model would allow resources to be provided 'on-demand' to meet the needs of particular research collaboration, but to balance the cost of its overall delivery. It would allow the pay-per-use business model used in the commercial world to infrastructures (IaaS), hosted environments (PaaS) and hosted applications (SaaS) to be integrated seamlessly alongside the academic resource providers offering a virtualised compute resource – but currently without the direct integration with the GEANT network [R20].

The user community within EGI consists of a combination and sometimes an intertwining of five different “users”: End users; Application Developers; Operations Staff; and Resource and Technology Providers. These groups are evaluated and outlined in Table 5 below and how integrating virtualisation would affect these groups.

Group	Current Experience	Virtualisation Effect	
		Positive	Negative
End User Community	Enjoy direct access to resources; Issues around complexity and inflexibilities of environments	Complete mobility; Self-service; Increased reliability and flexibility	Potential performance loss due to virtualisation overheads
Application Developers	Restricted to specific operating systems and middleware	Greater portability and ease of provisioning through appliances	Potential redesign of some applications
Community Operations Staff	Substantial effort dedicated to detecting problems, coordinating the diagnosis, and monitoring the problems through to resolution	Simplifies operations through the ability to quickly move virtual workspaces between physical server resources	Need to orchestrate VM provisioning across resource centres
Resource Providers	Staff effort restricts the environments and user communities that they can support	No longer directly involved in deploying community specific software	Loss of direct control in deploying user environments; Need to manage VM allocation to communities
Technology Providers	Need to provide new innovative tools and services, while supporting existing legacy software	Opportunity to develop and provide required management tools as virtualisation usage is increased	Need to ensure software can run and be configured in virtualised environments

Table 5: Virtualisation Effect on User Groups

5.3.3 Collaborations

An ecosystem can be viewed as a number of independent activities that interact and are dependent on each other. Within EGI, the collaborative ecosystem consists of EGI.eu as the coordinating body, Virtual Research Communities representing the end-user communities, resource infrastructure providers that coordinate resource centres at a national or domain level, technology providers, to a name a few. The adoption of virtualisation within the infrastructure makes it easier in the future for resource infrastructure providers to be commercial cloud providers and the clear need for groups previously embedded within the end-user communities to act as intermediaries between the resource providers and the end-users. Currently, these dependencies are captured within MoUs (focusing on infrastructure, technology providers, virtual research communities, projects, etc.) or OLSA/SLAs (with resource providers or technology providers). These MoUs can describe many activities, but typically include common plans around dissemination, representation to ensure the exchange of requirements, and the development of joint roadmaps.



The DCI Collaborative Roadmap describes the individual interactions between the six European DCI projects and shows how the provision of e-Infrastructures in Europe could evolve over the next 3 years and the contributions that each project may make towards this future by working with each other [R20]. Overall, EGI-InSPIRE provides a route for the deployment across Europe of new technological innovations into production once they have shown sufficient robustness and value to the EGI community. European Middleware Initiative (EMI) and Initiative for Globus in Europe (IGE) provide a source of innovation in the short-term, and are expected to expand over time with the inclusion of technology and procedures developed within the StratusLab project. VENUS-C will eventually provide best practices and potential success stories to the EGI community on the applicability of cloud computing for scientific computing, while EDGI will provide desktop and cloud resources to various European research communities.

Specific SLAs will be defined to govern the expected operational interactions on the provision of third line support and security incident handling. This will take place initially for two projects, EMI and IGE, which will deliver software components and support for the EGI community.

The shared vision of the DCI collaboration provides an added value response to the evolving European strategic landscape. The DCI community can influence EU policy decisions only if it acts jointly. Stronger external representation will need to go hand in hand with strong internal coordination.

Outside of direct technology oriented collaborations, EGI sees the European Strategy Forum on Research Infrastructures (ESFRI) as a large user community able to take full advantage of EGI. The recent ESFRI Roadmap [R34] has defined the preparatory phase funding for most projects with a big push to come in FP8 and beyond. To date, 44 projects cover domains such as: Social Sciences and Humanities; Environmental Sciences; Energy; Biological and Medical Sciences; Materials and Analytical Facilities; Physical Sciences and Engineering; and e-Infrastructures. Importantly, these projects will involve data intensive science requiring national commitments in a European context with global collaboration and shared access for the long-term (10-20+ years).

6 TOWARDS A TECHNOLOGY ROADMAP

The mission of EGI is to guarantee the long-term availability and access to a generic e-Infrastructure for all European research communities and their international collaborators. While several scientific communities are collaborating based on the current infrastructure, others have not yet moved their software environment into EGI for lack of technical compatibility or for limitations in the supported computing mode. The adoption of mature open-source technologies such as virtual machine management environments and hypervisors will enable cloud computing services to be hosted by EGI's resource providers. This opens the way for the evolution of EGI towards a more generic and flexible infrastructure able to better meet the needs of more user communities. After the extensive analysis provided in the previous sections, the user requirements, technical capabilities, and critical risks that need to be addressed are summarised below. Further work should concentrate on technology selection and implementation plans.

6.1 User Oriented Objectives

Need	Description
Scale out to new communities	A number of user communities do not yet collaborate through EGI because either their applications cannot run in the available OS or the edge services are not suitable for their computing model
Rapid provisioning of new capabilities	User communities should be able to deploy their own services in EGI through some authorised power-users from within their own community
Service quality assurance	User communities should be able to isolate resource use between different communities for QoS predictability, not affected by execution of parallel activities

Table 6: Identified needs from the user communities

6.2 Required Technology Capabilities

In this section, the technology capabilities that are needed to be available to the EGI community, either through UMD or from other sources, to address the emerged user requirements are listed.

Capability	Description
VM image repository	Ability to manage a repository of VM images integrated with the Grid authentication system, enabling special users to upload and describe VM images or normal users to retrieve
VM image discovery	Ability to discover the characteristics and location of images
VM mobility	Ability to move a VM from a node to another even on a different site; this capability may require a common format
VM management	Ability to configure, deploy, monitor, and decommission a virtual machine
Dynamic execution environment selection	Ability for a end-user to select a specific VM where to run its application
Dynamic execution environment deployment	Ability for a end-user to request the provision a customised virtual environment available at a resource centre to be instantiated for running its application
Accountability	Ability to account the VM usage per user and per group in order to provide the basis for defining quotas or support the billing - necessitates significant additional technical work to take advantage of metering services

Table 7: Capabilities required to augment UMD to support virtualisation and cloud-like services

6.3 Risks

In this section, the critical risks are identified that, if not met, may cause any transition to virtualisation across the European production infrastructure to fail.

Critical Factors	Mitigation
There are many players offering competitive solutions, the selection of the most appropriate platform may not be easy	A small group of experts should develop possible technology roadmaps that should be later reviewed, critiqued and validated by a larger group
The selected technologies may not provide all the features needed by the EGI community	EGI.eu should engage the selected technology providers via formal agreements (MoUs + SLAs) to make sure that they can be part of the requirements gathering process and communication via the Technology Coordination Board
Future sustainability and governance of the e-Infrastructure to communities that have not been actively involved in its development is not clear or assured.	Without improving the flexibility and responsiveness of the infrastructure providers to different communities and their service needs, will be very hard to expand user communities. Focused engagement with new communities will be needed to develop this capability and for them to adopt it
Providers will have less control over the application environment and may not be willing to delegate access to their resources	Policy relating to the use of virtualised environments within EGI need to be established and implemented in the deployed systems, e.g. identifying repositories of trusted VMs for use by production sites.

Table 8: Barriers and Solutions to the EGI Vision

6.4 Follow-up

In order to follow up with the technology roadmap definition, a larger consultation is needed. The decision in which technology to invest in is not straightforward and the need to coordinate the development of multiple technologies within the current operating infrastructure, given the impact and investment needed to make these changes, means building a consensus across all of EGI's stakeholders. Data is one of the largest challenges of e-Science, which is also not directly specified in terms of integrating cloud and virtualisation technology, but an issue for moving into the future. Based on the analysis and the context set out by this document, it is recommended to organise a dedicated workshop by late spring 2011 for which an outcome should be a technology roadmap for integrating virtualisation into EGI and updated every 12-18 months through established formal mechanisms as technology by its nature is ever-changing and evolving.

Preparatory work should include:

- Refining the list of UMD capabilities needed to augment EGI with virtualisation services
- Engaging candidate technology providers and better understand their detailed work plan, their availability to deliver EGI-specific needs and their own long-term sustainability options
- Consulting with the user communities to clarify specific needs in running their applications in virtualised environments or in deploying their own services.

Once the technology roadmap is defined, Memoranda of Understanding should also be established with the selected technology providers to formalise the implementation plan followed by Service Level Agreements by the end of 2011. EGI.eu has already signed MoUs with IGE and EMI to provide the software required by EGI's user community. Others are already underway or planned for throughout 2011 (i.e. StratusLab).



7 CONCLUSION

In these turbulent times of global economic uncertainty, financial constraints and fragile social cohesion, the EGI community simply cannot allow itself the luxury of staying passive and conducting “business as usual”. The EGI community must raise its own level of awareness in regard to the importance of this opportunity and react in a proactive and decisive way in order to fully embrace this moment for e-science and research. A shared determination and a common vision are needed to achieve this step that can qualitatively change the European scientific and research landscape.

Many campuses are encouraging the move of departmental or group level computing resources into central locations where they can be managed and supported by dedicated staff. This trend will inevitably continue over the next decade, forcing a greater integration between the client environment available at the researchers fingertips and the remote resources that they have access to ‘somewhere’ over the Internet. The ‘where’ of these resources will become increasingly less important to some communities, but of critical importance to those where their data is governed by legislation (e.g. medical, personal, financial, etc.). A researcher will have access to a pool of resources that are available to them through their roles within physical organisations (e.g. their employer), their funders (e.g. national resources), through their collaborations (e.g. international virtual organisations) or acquired commercially. Much more important will be the ‘how’ of configuring and exploiting these resources effectively for their own needs or those of their collaborators.

Cloud providers offering Infrastructure as a Service can be integrated seamlessly alongside the academic resource providers offering a virtualised compute resource – but currently without the direct integration with the GEANT network. Technology will always be evolving therefore it is essential that EGI.eu takes on its coordination responsibility in leading the evaluation of emerging technologies and facilitate the adoption of best practices where it makes sense and streamline efforts and resources where possible.

Overall, EGI provides human, technical and infrastructure services through the federation of national and domain specific resource providers to researchers in Europe and their international collaborators, a completely unique set of characteristics to anything else currently available. As presented throughout this report, the integration of cloud computing virtualisation technologies offers a wide range of technical, economical and organisational benefits and outlines a few sets of challenges that need to be addressed. Through the direct engagement with key experts over the next several months and the production of a defined roadmap, EGI.eu will ensure that any available opportunities will be thoroughly evaluated and where possible, implemented, ensuring the infrastructure continues to ever-evolve and improve for the current and new user communities it serves.

8 REFERENCES

R 1	MS407 Integrating Resources into the EGI Production Infrastructure – https://documents.egi.eu/document/111
R 2	D4.1 EGI Operations Architecture – https://documents.egi.eu/document/218
R 3	D5.1 UMD Roadmap - https://documents.egi.eu/document/100
R 4	EGI Service Level Agreement Template - https://documents.egi.eu/document/241
R 5	The Economic of the Cloud - http://www.microsoft.com/presspass/presskits/cloud/docs/The-Economics-of-the-Cloud.pdf
R 6	http://en.wikipedia.org/wiki/Power_usage_effectiveness
R 7	Borderless European Cloud Risks Fragmentation - http://www.businessweek.com/blogs/europeinsight/archives/2010/05/borderless_european_cloud_risks_fragmentation.html
R 8	Amazon EC2 Instance Purchasing Options - http://aws.amazon.com/ec2/purchasing-options
R 9	Cost-Benefit Analysis of Cloud Computing versus Desktop Grids - http://www-users.cselabs.umn.edu/classes/Fall-2010/csci8980-cloud/papers/cloud-seti-cost-anal_kondo_hcw09.pdf
R 10	EGI-InSPIRE Description of Work - https://documents.egi.eu/document/10
R 11	EGEE Comparative Study: Grids and Clouds - https://edms.cern.ch/file/925013/3/EGEE-Grid-Cloud.pdf
R 12	RESERVOIR (FP7 Project) – https://www.reservoir-fp7.eu
R 13	Server virtualisation, cloud software come to CERN - http://www.computerworld.in/articles/server-virtualisation-cloud-software-come-cern
R 14	VENUS-C (FP7 Project) - http://www.venus-c.eu
R 15	VMware CTO Blog - http://communities.vmware.com/community/cto/emea
R 16	Europe Insight Borderless European Cloud Risks Fragmentation - http://www.businessweek.com/blogs/europeinsight/archives/2010/05/borderless_european_cloud_risks_fragmentation.html
R 17	The Data Interoperability Challenge for Cloud Computing - http://www.infoworld.com/d/cloud-computing/data-interoperability-challenge-cloud-computing-259?source=footer
R 18	Enterprise and Industry Innovation - http://ec.europa.eu/enterprise/policies/innovation/index_en.htm
R 19	From EGEE to EGI - http://www.thedigitalscientist.org/feature/feature-eg-ee-egi-plain-talk-bob-jones-and-steven-newhouse
R 20	DCI Collaborative Roadmap - https://documents.egi.eu/document/207
R 21	EGI Role towards Europe 2020 - https://documents.egi.eu/document/317
R 22	Grid on Demand - http://staff.science.uva.nl/~deLaat/sne-2009-2010/p36/report.pdf
R 23	AWS Simple Monthly Calculator - http://calculator.s3.amazonaws.com/calc5.html
R 24	e-IRG meeting in Brussels, 16 December 2010 - http://www.e-irg.eu/e-irg-meeting-in-brussels-16-december-2010.html
R 25	NIST paper - http://www.nist.gov/itl/cloud/index.cfm

R 26	An Introduction to Service Oriented Computing - http://www.whatissoa.com/whatiscloud/p1.php
R 27	Tanks in the Cloud - http://www.economist.com/node/17797794
R 28	John Barr, The 451 Group - http://www.isc-cloud.com/2010/Program/Schedule/Cloudy-I-Can-See-Clearly-Now
R 29	A Perspective on Scientific Cloud Computing - http://dsl.cs.uchicago.edu/ScienceCloud2010/p11.pdf
R 30	http://en.wikipedia.org/wiki/Virtualisation
R 31	ENISA Cloud Computing Risk Assessment - http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment
R 32	Scientific Computing in the Cloud - https://agenda.cnaf.infn.it/getFile.py/access?contribId=33&sessionId=4&resId=0&materialId=slides&confId=364
R 33	e-IRGSP2 deliverable 4.3b - Final Legal Issues Report - http://www.e-irg.eu/publications/e-irgsp2-public-deliverables.html
R 34	ESFRI Roadmap Update 2010 - http://ec.europa.eu/research/infrastructures/index_en.cfm?pg=esfri-roadmap&section=update-2010
R 35	Review of the Use of Cloud and Virtualization Technologies in Grid Infrastructures - http://stratuslab.eu/lib/exe/fetch.php?media=documents:stratuslab-d2.1-v1.2.pdf
R 36	National Institute of Standards and Technology - http://www.nist.gov
R 37	Worker Nodes on Demands Service (WNoDeS) - http://web.infn.it/wnodes/index.php/wnodes
R 38	Forrester Research - http://www.forrester.com
R 39	European Organization for Nuclear Research (CERN) - http://public.web.cern.ch/public
R 40	University of Amsterdam The System and Network Engineering research group - http://www.science.uva.nl/research/sne
R 41	StratusLab Project - http://www.stratuslab.eu
R 42	Distributed Management Task Force (DMTF) - http://www.dmtf.org
R 43	Open Grid Forum (OGF) Open Cloud Computing Interface Working Group - http://occi-wg.org
R 44	European Strategy Forum on Research Infrastructures (ESFRI) - http://www.ec.europa.eu/research/esfri
R 45	e-Infrastructure Reflection Group Support Programme (e-IRGSP2) - http://www.e-irg.eu/about-e-irg/e-irgsp2.html