



HAL
open science

Initial Plan for Dissemination, Collaboration and Standardization Activities

David O'Callaghan, Kathryn Cassidy

► **To cite this version:**

David O'Callaghan, Kathryn Cassidy. Initial Plan for Dissemination, Collaboration and Standardization Activities. 2010. hal-00687170

HAL Id: hal-00687170

<https://hal.science/hal-00687170>

Submitted on 12 Apr 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Enhancing Grid Infrastructures with
Virtualization and Cloud Technologies

Initial Plan for Dissemination, Collaboration and Standardization Activities

Deliverable D3.1 (V1.0)
18 October 2010

Abstract

This document outlines the dissemination, collaboration and standardization activities of the StratusLab project, giving details of progress to date, and plans and mechanisms for future activities.



StratusLab is co-funded by the
European Community's Seventh
Framework Programme (Capacities)
Grant Agreement INFSO-RI-261552.



The information contained in this document represents the views of the copyright holders as of the date such views are published.

THE INFORMATION CONTAINED IN THIS DOCUMENT IS PROVIDED BY THE COPYRIGHT HOLDERS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE MEMBERS OF THE STRATUSLAB COLLABORATION, INCLUDING THE COPYRIGHT HOLDERS, OR THE EUROPEAN COMMISSION BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THE INFORMATION CONTAINED IN THIS DOCUMENT, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Copyright © 2010, Members of the StratusLab collaboration: Centre National de la Recherche Scientifique, Universidad Complutense de Madrid, Greek Research and Technology Network S.A., SixSq Sàrl, Telefónica Investigación y Desarrollo SA, and The Provost Fellows and Scholars of the College of the Holy and Undivided Trinity of Queen Elizabeth Near Dublin.

This work is licensed under a Creative Commons Attribution 3.0 Unported License
<http://creativecommons.org/licenses/by/3.0/>



Contributors

Name	Partner	Sections
O'Callaghan, David	TCD	Editor
Cassidy, Kathryn	TCD	Contributor

Document History

Version	Date	Comment
0.1	16 July 2010	Initial TOC skeleton
0.2	1 October 2010	First draft
0.3	8 October 2010	Incorporated feedback
1.0	18 Oct. 2010	Final revision for release

Contents

List of Tables	6
1 Executive Summary	7
2 Introduction	9
2.1 Purpose	9
2.2 Organization	9
3 Dissemination	10
3.1 Aims and Objectives	10
3.2 Background Information	10
3.2.1 Target audiences	10
3.2.2 Partner Strengths	11
3.3 Methods	13
3.3.1 Raising Awareness in Related Communities	13
3.3.2 In Print and Online	13
3.3.3 Publishing, Demonstrating and Disseminating Results	14
3.3.4 Upcoming Events	15
3.3.5 Software Releases	15
3.4 Roadmap	16
3.5 Evaluation	16
3.5.1 Website and on-line presence	16
3.5.2 Presentations and Talks on StratusLab	18
3.5.3 Press	18
3.5.4 User Community Surveys	19

4	Collaboration & Standardization	21
4.1	Aims and Objectives	21
4.2	Collaborations	21
4.2.1	Distributed Computing Infrastructure Projects.	22
4.2.2	Related Projects	23
4.2.3	Standards and Policy Bodies	25
5	Conclusions	28
	References	29

List of Tables

3.1	Roadmap of Dissemination Activities	17
3.2	Dissemination Metrics and Targets	18

1 Executive Summary

This document describes the dissemination, collaboration and standardization activities of the StratusLab project used to help achieve the project's goals:

- Dissemination aims to raise project awareness, inform users and the public about the project results, and interact with user communities.
- Collaboration will ensure that StratusLab benefits from the work of other EU projects, and similarly that those projects can use the outputs of StratusLab.
- Standardization is necessary in order to ensure that the StratusLab project adheres to best practice and also drives the standardization process in a direction consistent with the project's vision.

Each partner in the consortium has particular contacts in other projects and these will be leveraged both for dissemination purposes and to develop collaborations.

The project will use a wide variety of dissemination methods. Awareness of the StratusLab project and products will be promoted by the partners in the related groups, projects and fora in which they are involved. The media will be used heavily to inform the general public, and particular user groups about the project and its results. The website forms the core of this dissemination effort, all public documents are available there, along with information about the project, its current activities and status, and a range of dissemination materials to include documentation and training materials.

The results of the project will be published in relevant journals and conference proceedings. Demonstrations of the technology will be organised at system-administrator-focused conferences as well as trade shows and commercial events.

Each software release affords an opportunity for a renewed effort in raising awareness of the project, and a range of dissemination activities will be planned to coincide with releases. The details of these activities will be informed by the feature set of each release.

Plans for dissemination over the next year are in place, with activities including website updates, press releases, magazine articles, production of documentation and the running of training workshops. Much work has already been completed or is underway. The project has received good press coverage and the partners have given several introductory talks on the project at events across Europe. The

announce mailing list is in place and many potential users have signed up to receive project updates.

The plans for collaboration are advancing well with a concerted effort on the part of the DCI projects to define interfaces between the projects and work together. StratusLab has agreed mechanisms for collaborating with all of the DCI projects and with several other related projects. For each external project, one partner has been designated as the lead partner and they will drive the collaboration. StratusLab will collaborate with many projects at a technical level, with projects using the StratusLab software and providing feedback. StratusLab will also use the outputs of other projects where appropriate.

Standardization plans are still at an early stage. Potential targets for standardization have been identified along with a number of possible fora for this work. As the project progresses the partners will engage with appropriate standardization bodies to ensure that their work follows best practice and makes use of existing standards and that any new standardization in areas relevant to the project is driven and informed by the StratusLab vision.

2 Introduction

2.1 Purpose

The purpose of this plan for dissemination, collaboration and standardization activities is to provide guidance for StratusLab project members by setting out the aims and approach for these activities.

The target audience for this document are the project members responsible for dissemination, collaboration and standardization activities individually and collectively. In addition, it may be useful to other StratusLab project members, and collaborators in external projects to understand our approach.

2.2 Organization

The StratusLab consortium will carry out a set of activities focused on the dissemination of project results and the promotion of their commercial exploitation. A specific activity, WP3, has been provisioned for the purpose of establishing plans for dissemination and exploitation of project results during its lifetime and beyond, and for coordinating the execution of those plans. Further, WP3 will be contributing to standards, and concertation and cooperation activities with other EU-funded projects sharing similar research interests.

A separate activity, WP2, deals with interactions with users and related communities. These interactions go beyond merely dissemination activities as WP2 is tasked with requirements gathering, training, user evaluations, etc. There is, however, a clear dissemination aspect to the work of WP2, and to that end both work packages must coordinate closely. Some activities of WP2 are described in this document, and these are clearly identified as WP2 work in the text.

The objectives of WP3 are:

- Disseminate results of the project to resource providers, end-users, and the general public.
- Identify project contributions to standards bodies and standardization efforts.
- Coordinate interactions with related projects, developing Memoranda of Understanding between projects where appropriate.

Please note that this document does not cover the project's *Exploitation and Sustainability Plan*. This will be published at the end of the project's first year.

3 Dissemination

The StratusLab project is currently engaged in or is planning a wide range of activities devoted to dissemination for both specialised constituencies and the general public. This section describes these activities, along with a summary of work completed to date and a roadmap for future dissemination activities.

3.1 Aims and Objectives

Overall, the dissemination of the StratusLab project results have a high priority with the purpose of ensuring a high visibility of the project and its results across Europe and internationally.

The dissemination effort in StratusLab has several aims

- To raise project awareness
- To inform users and the public about the project results
- To interact with user communities in order to educate them about the project's benefits (in conjunction with WP2)

The StratusLab consortium aims to make the project results known to multiple target audiences including, but not limited to, industry partners, clients, journalists, researchers, students and other technical professionals. The most appropriate medium for each target audience will be used such as articles in trade journals, conference presentations, commercial brochures, demonstrations, web presence etc.

3.2 Background Information

3.2.1 Target audiences

The target audiences for StratusLab have been split into two broad categories, end-users and systems administrators. The end-users encompass domain scientists who will use StratusLab in their work, while the system administrators are the resource providers who will install the StratusLab toolkit on their sites, in particular those coming from a background in grid computing (from National Grid Initiatives and the European Grid Infrastructure).

In June, two surveys were conducted by WP2 which provide input for dissemination activities targeted at these two types of audience. The surveys revealed

information about the backgrounds of potential users, for example, the majority identified themselves as working in research or educational institutes, most came from the Computer Science/Engineering, Bioinformatics and Physics/High-Energy Physics domains, and there were a large number of Computer Science Researchers among the respondents.

These results not only help the project to better identify its core user base, but they highlight areas where the project may need to invest additional effort in order to broaden the audiences reached by dissemination activities. For example, it was recognised that it is necessary to make a stronger effort in disseminating the project aims and results to commercial enterprises as they were under-represented among the respondents. This might be done by targeting more industry publications and events as well as academic ones.

More detail on the results of these surveys can be found in deliverable D2.1 *Review of the Use of Cloud and Virtualization Technologies in Grid Infrastructures* [25].

Finally dissemination to the general public is also important, to raise awareness of the project and for educational purposes. In targeting the general public it is necessary to focus on the economic and societal benefits of the project, its impact on science and research, and how this can be of benefit to the broader community. This is quite a different approach to that taken for dissemination to technical or scientific audiences.

3.2.2 Partner Strengths

While TCD will coordinate dissemination activities, all partners will be involved in the dissemination effort. Partners will leverage their own contacts and the other projects in which they are involved, nationally and internationally, to broaden the network of dissemination avenues for the project. This section gives an overview of the types of activity on which each partner intends to focus, along with details of any specific useful contacts which they may have.

3.2.2.1 Universidad Complutense de Madrid (UCM)

UCM will be involved in dissemination of results with technical and scientific publications at conferences, participation in standardization bodies, and diffusion in user communities of the technologies developed by the research group. Other dissemination instruments are the initiatives co-coordinated by the research group in Spain, such as the Middleware Activity in the Spanish Initiative in e-Science [20] and the Working Group on SOI and Grids of INES - Spanish Technology Platform on Software and Services [21].

3.2.2.2 Centre National de la Recherche Scientifique (CNRS)

CNRS will use the mechanisms available through the Institut des Grilles (IDG) [11] and France-Grilles [8] (French NGI partner of the European Grid Infrastructure) to inform both users and system administrators of the StratusLab results. These organizations include users and system administrators involved in grid technology

in France. At grid workshops and conferences intended for French audiences; LAL and IBCP will present StratusLab results at these events to ensure the maximum visibility of the StratusLab in France. StratusLab has already been presented at the GRISBI Bioinformatics School [10] in Roscoff, France (27-30 September).

3.2.2.3 Greek Research and Technology Network (GRNET)

GRNET will promote the results of the project in big infrastructure initiatives like the European Grid Initiative [5]. In particular, as GRNET is a large-scale service provider, it will disseminate results that demonstrate that the StratusLab approach has benefits of value for adoption in large production settings in Europe and elsewhere. To this end, GRNET will capitalise upon the contacts and the relationships it has build over the last years with network and infrastructure providers around the world, and upon its key position in many e-Infrastructure projects.

3.2.2.4 SixSq Sàrl (SixSq)

SixSq Sàrl will integrate the StratusLab open source cloud distribution into its portfolio of technologies and demonstrate the benefits of the distribution to existing and prospective customers. SixSq employees, as well as existing customers and collaborators will be trained on the distribution so that they can propose its use where appropriate. In particular, the benefits of the distribution will be discussed with the European Space Agency (ESA) [7], an existing SixSq customer. SixSq will attempt to seek support in introducing the technologies to the different ESA laboratories and missions. Finally, SixSq has an existing network of fellow SMEs, investors and business analysts mainly in Switzerland, Germany in the UK. By demonstrating to these contacts the potential of the StratusLab distribution to help small and large customers deploy their own cloud infrastructures, it should be possible to create enthusiasts within this network who will themselves go on to disseminate StratusLab to their own customers and contacts.

3.2.2.5 Telefónica Investigación y Desarrollo (TID)

Telefónica I+D is a company 100% owned by the Telefónica TELCO group. Its main mission in the group is to develop R&D for business units of the rest of companies of the group. Results of the StratusLab project will be disseminated in the Telefonica group's specific Innovation Forums and Telefonica I+Dgcs Communications journal. In addition, there will be individual presentations targeted to those business units of the Telco related to Service provisioning (for example Information Systems Departments at Telefonica España as internal service provider or Telefónica Soluciones as an ASP for the external market). The project results will, thus, be disseminated across Europe and Latin America where Telefonica's presence is stronger.

Telefonica I+D is researching the development of business models related to the integration of solutions for public administrations and corporate clients with the collaboration of enterprises, local SMEs, research centres and academia. These models are been applied in NESSI [16] and INES [21] platforms, and the MOR-

FEO [14] Open Source Communities.

3.2.2.6 Trinity College Dublin (TCD)

TCD will coordinate the WP3 dissemination activities of the StratusLab project. It will also disseminate its results in scientific publications, conference presentations, and in the user communities of the StratusLab technologies (and ensure there is a forum for small sites). It will manage the project website and wiki. It will foster the “StratusLab weblogs” with articles by project members. It will also actively disseminate the project results within the Irish e-infrastructure community (e-INIS) [4] in training courses and e-Learning courseware.

3.3 Methods

The StratusLab project will make use of a range of methods to meet the dissemination aims and objectives. The general methods to be employed are outlined below.

The project will incorporate to the largest extent possible the communication best practices collected by the European Commission and will acknowledge the funding from the European Commission through the FP7 programme via the appropriate use of the FP7 logo, EU flag, EC/e-Infrastructures, and the like in all dissemination activities and materials. Scientific publications of the project will explicitly acknowledge the European Commission’s financial support through the StratusLab project.

3.3.1 Raising Awareness in Related Communities

StratusLab partners will raise awareness of the project in meetings, forums and distribution lists of related European projects in which they participate, in the Concertation Meetings of European Projects and technical groups, in the work groups and workshops of the Open Grid Forum [19], and in the Open Source Communities where the consortium participates, such as Globus Alliance and Morfeo.

3.3.2 In Print and Online

The project is attempting to create a strong on-line presence. The project webpage is available at <http://www.stratuslab.eu/> and provides general information about the project, along with news and status updates on the project, an archive of public documents and deliverables, links to presentations and other dissemination material, plans and a *Current Activities* section which outlines the development and research tasks in progress at any given time.

This site is in active use by the project members as well as forming the core of the project’s external interface, and it is regularly updated. Continued updates and improvements to the website content are planned.

The site integrates twitter and RSS feeds to provide latest news about software releases, press releases, attendance at events, and other newsworthy items. A project weblog is also planned which will form a hub for online discussion of the

project by users and other interested parties.

WP3 has begun to analyze the search engine results for the website and an effort to improve the site's page ranking is ongoing. It is hoped that this will result in an increase in website traffic.

There has been an attempt to integrate tools to encourage visitors to bookmark and republish pages of interest to them, for example, the website uses a plug-in which allows visitors to instantly create a bookmark to the site in a range of external sites and portals, including facebook, twitter, and others.

The project partners maintain several weblogs such as the DSA Research blog [3], and the Morfeo Cloud Technologies blog [15], and also publish posts in GridGurus [9] and OGF Thought Leadership Series [17] and these have been and will continue to be used as dissemination tools for the project.

The project will also be promoted through articles in magazines and other general interest periodicals. The project partners will issue press releases to announce project events and news to local, national, and international press.

3.3.3 Publishing, Demonstrating and Disseminating Results

The results of the project will be described and made available to Technology Platform groups in which the consortium partners participate, such as INES and NESSI. StratusLab results will be shared with these groups while at the same time insights and feedback can be solicited on the value of StratusLab to the service computing community.

In addition results will be publicised at scientific meetings and conferences at national and international levels. The project partners participate on the Program Committees of several relevant annual events on parallel/distributed processing, Grid computing, Cloud computing and virtualization and can push to have StratusLab featured prominently, for the purpose of general project awareness, and also to describe specific results which are relevant to the scientific communities involved. Seminars and demonstrations will be organised about the new technology incorporating the research results and its application for research centres and businesses.

Publications in the most relevant national and international journals and conferences will aid in both raising project awareness and dissemination of results. All partners have undertaken to identify possible publication targets relevant to the specific area of the project in which they are involved. It is likely that such publications will become more important later in the project when there are more concrete results to report. A full list of relevant Calls for Papers is maintained on the project wiki [27].

The project intends to exploit networking through project participants and collaborating projects to reach system administrators and other target users. A number of conferences and bodies will offer dissemination opportunities for different target users. With respect to more in-depth knowledge transfer, in particular for system administrators, WP2 will conduct seminars, demonstrations and tutorials at events. In addition documentation will be contributed to, or linked from, relevant system administration websites and wikis.

The project will also use the press and general-interest publications as a venue to describe the project results. As project results are generated or new features are implemented in the software, articles will be written which describe the benefits of these to target communities. Publication of these in relevant domain-specific or popular-science publications will not only educate potential users about the project results, but will further raise awareness of the project in these target communities. For example, an article is currently planned for the online magazine International Science Grid This Week [12] describing the results of the user and administrator surveys.

3.3.4 Upcoming Events

StratusLab intends to have a presence at many events. This may take the form of stands and booths, talks, presentations, panel discussions, posters and paper submissions.

Some of the events at which a presence is planned are listed below. This is not intended to be an exhaustive list, but merely a sample of upcoming events in the short term to which the project hopes to submit papers, present talks, or have some presence at. A full list of relevant events is maintained in the project wiki [27].

- Open Grid Forum 25-28 October, 2010 [18]
- EMI all-hands meeting 22-24 November, 2010 [6]
- EGI User Forum 11 – 15 April 2011 [5]
- CloudCamp [1]

It is also intended to demonstrate the StratusLab distribution and seek feedback at trade shows such as VMWorld [23], the Virtualization Conference and Expo [22] and similar commercial events.

3.3.5 Software Releases

Public releases of the StratusLab distribution provide an opportunity not only to re-new contact with existing users, but to further publicise the StratusLab project via release announcements, press-releases, emails and other means. Each scheduled release will include an internal dissemination plan, with details of the proposed activities to publicise the availability of the new version. The feature set in a particular release will inform this plan, for example, the venues chosen to publicise the release may differ depending on whether the new functionality is aimed particularly at users or at administrators.

The first release of the StratusLab distribution is expected mid-October 2010.

A significant dissemination effort is planned to coincide with this release. The release will be announced on the StratusLab website, and notification will be sent to user communities via the StratusLab Announcements mailing list, and also potentially via other mailing lists and fora in which the project members are involved.

A Press Release will be written and circulated in order to further disseminate within both the technical press, and to the general public.

Partners who maintain Weblogs will be asked to post entries which mention the StratusLab project and link to the website and the software.

The release notification will also be posted on the StratusLab twitter and RSS feeds.

An article is currently being planned for iSGTW which describes the results of the user surveys which were carried out by WP2, and this article will mention the software release and provide a link to the website for more information.

3.4 Roadmap

Table 3.1 presents a roadmap of dissemination activities for the first year of the project. Only those in the relatively near-term can be considered concrete plans, but later events are listed to give an idea of the types of events being planned. The details of those in the longer-term are likely to change as the project progresses, and many additional activities will be added.

3.5 Evaluation

Table 3.2 lists the project's dissemination metrics that will be used to understand the impact of the dissemination activities. This will be supplemented by active analysis of each dissemination activity to understand how to improve the overall effort. The remainder of this section describes the dissemination activities completed or currently underway.

3.5.1 Website and on-line presence

The StratusLab website went live in May 2010 and forms a core part of the project's dissemination strategy as the main interface with the general public and the first contact point for many potential users.

The website is regularly updated to ensure that the content is of interest to users and the current status of the project is always readily available.

To date, the project Twitter and RSS feeds have been used to announce the project kickoff, participation in major events, publication of project documents and more informal updates on development. The use of these tools will be continually reviewed and expanded.

The number of visits to the website is still low, but steadily increasing. WP3 is investigating the use of Google Analytics and other tools to improve the search engine ranking and number of visits.

In order to spread awareness of the site website the address ("stratuslab.eu") should be included prominently in other project dissemination activities.

Table 3.1: Roadmap of Dissemination Activities

Activity	Date	Description
Website Updates	Ongoing	The website content will be continually updated to ensure that the site is informative and interesting to users
OGF30/Grid2010 [18]	25-28 October 2010	Keynote speech given at this conference by a project participant to include slides on StratusLab
CloudComp2010 [2]	26-28 October 2010	Keynote speech given at this conference by a project participant to include slides on StratusLab
ISC Cloud2010 [13]	28-29 October 2010	Keynote speech given at this conference by a project participant to include slides on StratusLab
First Release v0.1	Early Nov. 2010	To coincide with the first release of the StratusLab Distribution a wide range of dissemination activities are planned including press releases, mail shots, etc.
Documentation	Early Nov. 2010	Documentation will be produced to accompany the first release and this will be available on the website
iSGTW Article	November 2010	An article is planned in the iSGTW outlining the results of the user surveys and publicising the availability of the first version
EMI Meeting	22-24 November 2010	Presence at EMI all-hands meeting
Website Redesign	Q4 2010	A possible website redesign to make it more user-friendly and eye-catching has been proposed
Social Networking	Q4 2010	Further integration with Microblogging, Social Networking or other tools will be investigated
Search Engine Optimization	Q4 2010	WP3 has begun to analyze the search engine results for the website and an effort to improve the site's page ranking is ongoing.
CloudCamp	Q4 2010	Presence at CloudCamp meeting
Project Weblog	Q4 2010	A project weblog will be set up
EGI User Forum	11-15 April 2011	Presence at EGI User Forum
Training Event	Q4 2011	It has been suggested that a Cloud School could be run in late 2011, either in conjunction with another event or as a stand-alone event.

Table 3.2: Dissemination Metrics and Targets

Metric	Source	Target	
		Y1	Y2
Number of people on StratusLab announcement list	Mailer	25	75
Number of people on StratusLab discussion list	Mailer	50	100
Number of VOs server via sites with StratusLab distribution installed	Information System	10	30
Number of scientific disciplines served via sites with StratusLab distribution installed	Information System	3	7
Number of views of website	Web server	–	–

3.5.2 Presentations and Talks on StratusLab

StratusLab project members have been active in giving talks and presentations at grid and cloud events since the beginning of the project, and even before the project's official kickoff. Presentations have been given at the following events. A full list of StratusLab presentations is maintained in the project wiki [28].

- the GRISBI Bioinformatics School in Roscoff, France, 30 September 2010
- the EGI Technical Forum, 14-17 September 2010;
- as part of a course on Grid and e-Science at the Instituto De Física Corpuscular, Valencia, Spain, 6-9 July, 2010;
- the XtreamOS summer school, Ulm University Günzburg, Germany, 5-9 July, 2010;
- the EGEE User Forum/OGF25, 4 March 2009;
- HEPiX Spring 2009, 28 May 2009;
- EGEE 2009, 22 September 2009.

StratusLab had a booth at the EGI Technical Forum in September 2010. Fact-sheets and three posters covering the project in general; integration with Quattor; and benchmarking were designed and printed for the event.

3.5.3 Press

A project kickoff press release was created in June, in both general and localised forms. The localised versions were not only translated into the local language, but

also had a particular focus on the participation of local institutions. This press-release was widely picked up in both grid and cloud publications, as well as some more general scientific and educational publications. Coverage by major grid and cloud publications included Cloud Computing Journal, International Science Grid This Week, SmartGrid, HPCWire and others. Including smaller local technical and educational publications, the press-release was carried in over fifteen online publications.

The project has initiated contact with the widely read International Science Grid This Week online magazine. StratusLab members have participated in iSGTW Live Chat! on 14 July 2010 and mentioned StratusLab in the discussion. The project kickoff press release was also picked up by iSGTW (30 June 2010) and an announcement of the user and administrator surveys was printed on 7 July 2010. Further articles in iSGTW are planned.

Contact has been made with the e-ScienceTalk project which is responsible for publishing iSGTW, to discuss other the dissemination avenues which they might make available to the project, such as GridCafe and GridCast.

A full list of references to StratusLab in the press is maintained in the project wiki [29].

3.5.4 User Community Surveys

In addition to these dissemination activities, there has been a specific activity, led by WP2, to survey the StratusLab target user communities, and a part of this effort focuses on dissemination. Two online surveys were created, one for end-users and one for systems administrators. These surveys ran for four weeks in July.

There were three broad goals to support dissemination efforts of the project

- To inform members of those communities about existing software and services related to virtualization technologies.
- To increase awareness of the StratusLab project within the target communities.
- To expand the number of direct contacts within the target communities.

Results of the surveys confirm that progress has been made along these lines.

As these surveys were web-based, they could include links to other web sites. Questions that mentioned specific technologies systematically included links to relevant web sites. We could not track how many of those links were actually used. Nonetheless, we hope that interested respondents followed them to find out more about virtualization, cloud, and other related technologies.

The welcome pages for the administrator and user surveys were designed to describe the StratusLab project and goals as well as the surveys themselves. All of those visiting the surveys should at least be aware of the project and have seen a link to the project's web page. The administrator and user surveys were visited 248 and 607 times, respectively. These visits include the number partial and complete

survey responses, but do not include multiple visits from the same IP address. (Geographic information is not available from these visits.) Consequently, information about the StratusLab project has reached a significant number of people within our target communities.

As a result of these surveys, 59 people have agreed to be on the StratusLab announcement mailing lists (including both incomplete and complete responses). There were 16 people who responded to the administrator survey that opted-in and 48 people from the user survey. Five people responded to and opted-in on both surveys.

Overall, the project has increased the number of direct contacts interested in the StratusLab results and increased awareness of the project within the targeted communities through these surveys.

4 Collaboration & Standardization

4.1 Aims and Objectives

A large number of projects, companies, and standards bodies currently focus on cloud and virtualization technologies because of the promise and growing adoption. StratusLab must actively engage with those entities to ensure that the project's results are well represented, that we are aware of others' advances, and that we drive standardization in a direction consistent with our vision.

StratusLab will collaborate with other EU projects, particularly those in the Distributed Computing Infrastructures area, to identify common interests, and to avoid duplication of effort.

The project will identify and actively engage with relevant standardization bodies, both in terms of using existing standards where appropriate, and driving the standardization process.

4.2 Collaborations

During the project lifetime, collaboration will be pursued according to the plans and agreements in place (while allowing for flexibility if circumstances change) and collaboration activities will be reported as part of the project's regular reporting cycle. We expect collaboration to take place at both managerial and technical levels.

Many other concurrent European projects will be interested in the results of StratusLab or will be exploring complementary technologies. This task will identify those projects of interest, open discussions about collaboration with those projects, and conclude Memoranda of Understanding with those projects, if appropriate.

Note: An overall roadmap for collaboration with projects has yet to be established, beyond the desire to establish initial contact with major projects in Q3 2010, which has broadly been achieved.

We expect each bilateral collaboration to have its own roadmap, determined by the collaborating contacts and forming part of any agreement between projects.

4.2.1 Distributed Computing Infrastructure Projects

Preparations for collaboration with projects in the DCI area have already begun in the form of in-person meetings (where possible) as well as audio conferences between representatives from all related DCI projects. The purpose of these meetings is to build understanding of each project's objectives and to uncover possible collaborations.

A common deliverable *Distributed Computing Infrastructure Collaborative Roadmap* [26] detailing the proposed interactions between each of the DCI projects has been created with input from StratusLab and the other DCI projects.

Follow-up face-to-face meetings with each of the DCI projects are now being planned to further define the interfaces for interaction between project.

4.2.1.1 EDGI

EDGI operates a distributed computing infrastructure based on desktop computing technologies (BOINC and Xtremweb) with bridges that allow interoperability between the European Grid Infrastructure and the EDGI infrastructure. EDGI will extend this interoperability to infrastructures based on cloud technologies and looks to create a higher quality of service for users of their infrastructure through dynamic allocation of additional resources.

The respective roles of StratusLab and EDGI as providers and users of cloud technologies forms a natural focal point for collaboration between the two projects. StratusLab will offer its expertise to EDGI in the areas of virtual appliance creation and will also provide appliance-building tools, a test environment and the use of the StratusLab appliance repository.

EDGI will use the StratusLab test resources to test its virtual appliances and will upload them to the StratusLab appliance repository, providing feedback on these tools to StratusLab. The EDGI virtual appliances will form a valuable test of the StratusLab use-cases.

Common participants: CNRS

4.2.1.2 EGI-InSPIRE

The European Grid Initiative is the organization responsible for the coordination and operation of a Europe-wide grid infrastructure provided by resource centres residing inside National Grid Initiatives. EGI-InSPIRE is a four-year project that will conclude the transition from EGEE to EGI and instantiate the new models of NGI-based grid operations.

StratusLab and EGI-InSPIRE will have to collaborate closely in order to ensure the integration of sites running StratusLab into the EGI infrastructures.

Common participants: CNRS, GRNET, TCD

4.2.1.3 EMI

The European Middleware Initiative project is responsible for the support and development of the Unified Middleware Distribution (UMD) a term used to refer

to an integrated distribution of the three most popular European grid middleware, namely gLite, Unicore and ARC.

StratusLab will deploy grid sites running the UMD services on top of cloud infrastructures and interaction with EMI will be required in order to ensure that the two distributions are fully compatible.

Common participants: GRNET, TCD

4.2.1.4 IGE

IGE (Initiative for Globus in Europe) coordinates European Globus activities and supports European e-infrastructures in their use of the Globus Toolkit. EGI intends to develop grid services and tools that can take advantage of the cloud. The two projects will share requirements and experiences, and will also work together to cloud-enable Globus services and tools using the StratusLab distribution.

Common participants: UCM

4.2.1.5 Venus-C

StratusLab and VENUS-C (Virtual multidisciplinary Environments Using Cloud infrastructures) share a common interest in cloud computing. VENUS-C will develop and deploy a Cloud Computing service for research and industry communities in Europe. This will be presented to the interested communities in the form of a scientific Platform as a Service (PaaS). Both projects, at different levels, try to bring the benefits of Cloud Computing to the scientific community.

The StratusLab distribution can fit into the VENUS-C architecture back-end, and the projects will also share requirements and evaluation results, as well as working together on standards definition.

Common participants: The projects have no common participants, but the collaboration efforts will be led by UCM.

4.2.2 Related Projects

4.2.2.1 e-Science Talk

e-ScienceTalk (formerly GridTalk) is a project aimed at coordinating the dissemination outputs of European grid computing efforts, ensuring their results and influence are reported in print and online.

The project publishes or maintains a number of dissemination channels including GridCafe, International Science Grid This Week (iSGTW) and GridBriefings.

StratusLab will provide e-Science Talk with information, news, and announcements on its cloud/grid activities as well as expert opinion on cloud and grid interactions. To date there have been several items on StratusLab in iSGTW which were described more fully in section 3.5.3.

e-Science Talk can provide StratusLab with dissemination advice, contacts, channels and opportunities, for instance, iSGTW, the GridTalk blog, and Grid Briefings. A Grid Briefing on grid and cloud is due during the lifetime of the StratusLab project.

StratusLab and e-Science Talk may work together on some activities, such as producing dissemination material and reports. Collaboration will take place at an organizational level although there may be some technical collaboration on document production or design, and the interactions will be led by TCD.

4.2.2.2 HEPiX Virtualization Working Group

The goals of the HEPiX group are to see adoption of worker node virtualization at grid sites; and to enable images created at one site to be used at other sites, primarily within the HEPiX and WLCG community.

StratusLab will collaborate with the HEPiX group at a technical level. The initial plan is to provide the HEPiX group with experience in virtualization, cloud computing, and OpenNebula in particular. Later, the StratusLab project results including the StratusLab toolkit distribution, the appliances repository, use cases, and tools for automatic image generation will be made available to the HEPiX group.

The HEPiX group could provide some level of confidence in using virtualization on grid sites by acting as test sites (“guinea pigs”) for the StratusLab distribution. Of particular interest to StratusLab is the work of the HEPiX group on “Image Generation Policy” and StratusLab hope to work with the HEPiX group on developing this policy, and on issues of image contextualisation and support for various virtualization hypervisors: HEPiX identified KVM and Xen as hypervisors of interest.

StratusLab partners will take account of the HEPiX group’s policies wherever they are relevant to the work of the project.

Common participants: CNRS/LAL

4.2.2.3 Quattor

Quattor is a system administration toolkit providing a powerful, portable, and modular set of tools for the automated installation, configuration, and management of clusters and farms.

The Quattor Toolkit, heavily used by grid resource centers, has been chosen as the first example to demonstrate the StratusLab Distribution’s compatibility with automated site management tools.

Quattor is developed as a community effort and provided as open-source software and thus the interactions will largely take the form of StratusLab members contributing feedback, bug-reports, and possibly patches back to the Quattor project. StratusLab members will also attend the regular Quattor Workshops in order to keep up to date on new developments in Quattor, and to inform other Quattor users about the potential of the StratusLab distribution to integrate into their Quattor-managed sites.

4.2.2.4 RESERVOIR

RESERVOIR is a European Union FP7 funded project that will enable massive scale deployment and management of complex IT services across different admin-

istrative domains, IT platforms and geographies. The project will provide a foundation for a service-based online economy, where through the use of virtualization technologies resources and services are transparently provisioned and managed on an on-demand basis at competitive costs with high quality of service.

The RESERVOIR project is involved in developing new functionalities for OpenNebula, upon which the StratusLab distribution is based.

StratusLab members have attended the Research/Industry Collaboration on Open Source Cloud Middleware workshop organised by RESERVOIR and is in communication with the project to try to identify areas for collaboration.

As the RESERVOIR project will be drawing to a close during the lifetime of the StratusLab project, there has been some discussion about the possibility of transferring some code to StratusLab, but this is still under discussion.

4.2.2.5 ECEE

Enabling Clouds for eScience (ECEE) is a project devoted to collaboration between cloud projects that serve the eScience community. StratusLab is a member of the ECEE project and thus has access to this forum for collaboration between related projects, discussion of requirements, formulation of use cases, etc.

The other project members are NEON, BalticCloud, NGS, GRNET, SARA, UCM, VENUS-C, SEECI and CESGA. StratusLab will collaborate with these projects through the ECEE forum.

4.2.3 Standards and Policy Bodies

StratusLab will engage with standards bodies to drive standardization in a direction consistent with our vision. Possible standardization targets include: Cloud API, Contextualization, Cloud/Grid Interface, Security and Accounting.

The project will investigate the use of existing standards where appropriate, and will engage with standards bodies to extend existing standards or create new standards.

Individuals participating in the technical work packages are best placed to participate in the technical discussions surrounding the standardization of cloud and virtualization services, thus all partners will be involved in contributing to standards and policy bodies. This task will identify standards activities of interest to the project, encourage project participants to work in those standardization activities, and report on their impact within the standardization efforts.

4.2.3.1 OGF OCCI

OGF OCCI (Open Cloud Computing Interface) is a specification of a cloud REST interface for the remote management of compute, storage and networking resources. OCCI is the only open specification accepting contributions from the community, and so a main candidate to be exposed by the StratusLab distribution.

This OGF Working Group is now evaluating possible extensions to the specification. StratusLab will contribute in order to ensure that upcoming versions of this specification address the requirements of the use cases implemented in the project.

StratusLab may also be in a position to become a reference implementation for some of these new extensions.

The OCCI WG maintains several documents with use cases and requirements which could form inputs to StratusLab and help inform the choice of components to implement.

StratusLab's involvement with OCCI will take the form of contributions to discussions in the mailing list and meetings.

TID, as WP6 coordinator, will represent StratusLab in this group, and UCM is also deeply involved in this collaboration.

4.2.3.2 e-IRG

The e-IRG (e-Infrastructures Reflection Group) is a body of national representatives who discuss Research e-Infrastructures and provide recommendations.

StratusLab intends to approach the organisers of the e-IRG workshops, which are held twice a year, and propose to have a presentation about StratusLab at the workshop in order to make e-IRG members more aware of the project, its aims and results.

The e-IRGSP (e-Infrastructures Reflection Group Support Program) is a project aiming to support the workings of the e-IRG, by aiding in the organisation of e-IRG events (workshops and meetings), preparing the e-IRG publications, supporting e-IRG task forces, etc. e-IRGSP is currently near the end of its second incarnation (e-IRGSP2, ending in the end of 2010) and preparing the third project in its series (e-IRGSP3, starting in December 2010).

Collaboration with e-IRGSP3 will be promoted by contacting its liaison activities and making sure that StratusLab is included in the knowledge base of e-Infrastructure related projects they maintain.

4.2.3.3 SIENA

StratusLab and SIENA share a common interest in collaboration with standards organizations on aspects of cloud computing.

StratusLab will provide SIENA with expert opinion on grid and cloud computing and requirements for standardization from the perspective of implementers and users and, indeed, a cloud software implementation of standards as an exemplar.

SIENA in turn, will provide StratusLab with contact to standardization bodies relevant to cloud computing, and references to existing standards, as well as expert industry opinion and requirements.

The SIENA project is developing a Roadmap on "Grid and Clouds for Research and Public Services" and StratusLab intend to work jointly with SIENA to contribute to this.

The interactions will be led by TCD.

4.2.3.4 DMTF

The Distributed Management Task Force (DMTF) [31] is the leading industry organization for the development, adoption and promotion of interoperable manage-

ment standards and technologies for IT systems and infrastructures.

Regarding the StratusLab vision and objectives, the key DMTF groups to monitor in the field of Cloud Computing technologies are:

Cloud Management Working Group (CMWG) carries on with the work developed in the DMTF Cloud Incubator group. The CMWG will develop a set of prescriptive specifications that deliver architectural semantics as well as implementation details to achieve interoperable management of clouds between service requestors/developers and providers.

System Virtualization, Partitioning and Clustering (SVPC) Work Group is developing DMTF standards for virtualization management. This includes the discovery, configuration, and deployment of virtual computer systems

Telefonica I+D will contribute to the consortium with its experience as active member of the DMTF, leading all the project's activities concerning DMTF standards in this field. Telefonica I+D takes actively part in the DMTF groups mentioned above and its contributions are focused in the following areas.

Firstly, Telefonica has released and submitted to the DMTF its API for Cloud Computing interoperability, the TCloud API [30]. TCloud API is based on vCloud API specification 0.8, as published by VMware, Inc. In essence, compatibility for the main operations and data types defined in vCloud are maintained in TCloud, but it provides extensions on advanced Cloud Computing management capabilities including additional shared storage for service data, network element provisioning (different flavors of load balancers and firewalls), monitoring, snapshot management, and so on. This submission took place in the Cloud Incubator (this group is now closed) and now it is one of the proposed APIs for guiding the standardization approach around the Cloud Computing interoperability. The CMWG is in charge of this work.

Secondly, the SVPC group is developing the Open Virtualization Format (OVF) [24], a standard which describes an “open, secure, portable, efficient and extensible format for the packaging and distribution of software (Virtual Appliance) to be run in virtual machines”. The OVF format is broadly supported by the main industry actors (e.g. VMWare, Citrix, etc). Telefonica is monitoring the evolution of this standard, proposing extensions and providing feedback basing on the experience and work developed in several internal and collaborative projects (e.g. Reservoir, Nuba, Vision, etc).

5 Conclusions

The project has established overall aims for dissemination, namely to raise awareness and knowledge of StratusLab among the core target audiences of system administrators and scientific users, and to further promote it to the general public with an interest in the area.

Similarly, StratusLab has set out aims for collaboration and standardization: to further the aims of the StratusLab project by finding common interests with related projects with the overall goal of enhancing the distributed computing infrastructures available without duplicating effort.

In this plan we describe the methods that we will use to achieve our dissemination aims in various arenas. At the present time, the project has established an on-line presence, and has been well covered by the online press, with a number of articles and very good media coverage of the project kickoff press release. The project has also had good visibility at events with posters, booths, presentations and discussions at a number of large grid and cloud conferences and workshops. The announcement mailing list has over 60 people subscribed, well over the target for this metric. Partners have delivered many talks and presentations on StratusLab at different venues throughout Europe.

A plan for future dissemination is in place with each software release, and a number of major European events acting as focus points for increasing awareness of the project within the target communities. As the project progresses, reporting of results will take on more importance with the publication of scientific papers in journals and conference proceedings.

This document also gives an outline of the plans for collaboration with related projects, which are underway. Particular attention must be given to the Distributed Computing Infrastructure projects and, indeed, StratusLab has contributed to a common vision for the DCI projects [26].

Other related projects cover areas such as dissemination, fabric management, and virtual machine image security requirements. These collaborations should provide benefits to all sides through work on common interests.

Standardization efforts are in the early stages, but a number of targets have been identified and these will be pursued by the partners. In addition, the project aims to base a number of technical developments on established standards, where it makes sense to do so.

References

- [1] CloudCamp. <http://www.cloudcamp.org/>.
- [2] CloudComp2010. <http://www.cloudcomp.eu/index.shtml>.
- [3] DSA Research Blog. <http://blog.dsa-research.org/>.
- [4] e-INIS. <http://www.e-inis.ie/>.
- [5] European Middleware Infrastructure. <http://www.egi.eu/>.
- [6] European Middleware Initiative. <http://www.eu-emi.eu/>.
- [7] European Space Agency. <http://www.esa.int/>.
- [8] France-Grilles. <http://www.france-grilles.fr/>.
- [9] GridGurus. <http://gridgurus.typepad.com/>.
- [10] GRISBI Bioinformatics School. <http://www.grisbio.fr/evenements/ecole2010/>.
- [11] Institut des Grilles (IDG). <http://www.idgrilles.fr/>.
- [12] International Science Grid This Week (iSGTW). <http://www.isgtw.org/>.
- [13] ISC Cloud2010. <http://www.isc-cloud.com/2010/>.
- [14] Morfeo. <http://www.morfeo-project.org/>.
- [15] Morfeo Cloud Technologies Blog. <http://cloudtechnologies.morfeo-project.org/>.
- [16] Networked European Software and Services Initiative. <http://www.nessi-europe.com/>.
- [17] OGF Thought Leadership Series. <http://www.ogf.org/TLS/>.
- [18] OGF30/Grid2010. <http://www.ogf.org/OGF30/>.
- [19] Open Grid Forum. <http://www.ogf.org/>.

- [20] Spanish Initiative in e-Science. <http://www.e-ciencia.es/>.
- [21] Spanish Technology Platform on Software and Services (INES). <http://www.ines.org.es/>.
- [22] Virtualization Conference and Expo. <http://www.virtualizationconference.com/>.
- [23] VMWorld. <http://www.vmworld.com/>.
- [24] Distributed Management Task Force. Open virtualization format specification. specification dsp0243 v1.0.0d, Sep 2008. <https://www.coin-or.org/OS/publications/optimizationServicesFramework2008.pdf>.
- [25] C. Loomis and M. Airaj. Review of the use of cloud and virtualization technologies in grid infrastructures, 2010. D2.1.
- [26] S. Newhouse, A. di Meglio, A. Papaspyrou, D. O'Callaghan, J. Kovacs, and A. Manieri. Distributed Computing Infrastructure (DCI) Collaborative Roadmap. online at: <https://documents.egi.eu/public/ShowDocument?docid=172>, September 2010.
- [27] StratusLab. Stratuslab calls for papers. online at: <http://www.stratuslab.eu/doku.php?id=cfp>.
- [28] StratusLab. Stratuslab presentations. online at: <http://www.stratuslab.eu/doku.php?id=presentations>.
- [29] StratusLab. Stratuslab press. online at: <http://www.stratuslab.eu/doku.php?id=press>.
- [30] Telefonica. TCloud API specification, version 0.9.0. Online resource, 2010. http://www.tid.es/files/doc/apis/TCloud_API_Spec_v0.9.pdf.
- [31] DMTF The distributed management task force webpage. Online resource, 2010. <http://www.dmtf.org>.

Distributed Computing Infrastructure (DCI) Collaborative Roadmap

Contributors on behalf of their projects: Steven Newhouse (EGI-InSPIRE), Alberto Di Meglio (EMI), Alexander Papaspyrou (IGE), David O'Callaghan (StratusLab), Jozsef Kovacs (EDGI), Andrea Manieri (VENUS-C)

Note: This material has been provided by the ‘Distributed Computing Infrastructure (DCI)’ projects funded in the November 2009 call to form a common text agreed by all the projects. This text is to be used *unaltered* by these projects in their respective deliverables. It is expected that the text will be placed into project specific document templates and into each project’s own review process to gain the endorsement of each project. If used as part of a document addressing other aspects it must either be referenced, or included in its entirety unaltered as an appendix in the document. This document has been produced with the co-funding of the DG Information Society & Media, European Commission. The content of this publication is the sole responsibility of the DCI project members and cannot be considered to reflect the views of the European Commission.

Executive Summary

As a result of an open call that closed in November 2009, six projects are being granted with nearly €50M by the European Commission’s 7th Framework Programme in the area of Distributed Computing Infrastructures. Together these projects will provide a pan-European production infrastructure built from federated distributed resources, ensure the continued support, maintenance and development of the middlewares (gLite, ARC, UNICORE and Globus) that are in common use in Europe, explore how grid sites and different applications can be hosted sustainably in commercial, public, publicly procured and private ‘cloud computing’ environments, and provide desktop resources to the European research community.

These projects are the result of over a decade of community building that has taken place in the area of European Distributed Computing Infrastructures – both in their operational provision to a multi-disciplinary user community and the research and associated software development to build such infrastructures. Together, these infrastructures face the challenge of evolving their services to the changing needs of their data-intensive user communities, and providing a sustainable service that will support their users today, tomorrow and the years to come. A vision is presented of moving from the current production infrastructure in Europe to one based upon federated virtualised resources. It is expected that this change will increase the flexibility of resource providers to meet the changing needs of the user communities they serve by adopting best practices from other sectors.

In this report, the individual interactions between the six projects are recorded. Many of the projects expect to define these interactions through Memorandum of Understanding, and where there is an operational relationship between the projects through a Service Level Agreement. Not all of the projects have identified concrete interactions at this point in time with any other, though these may develop during the course of the individual projects.

1. Introduction

Call 7 (FP7-Infrastructures-2010-2) under the e-Infrastructures topic of the FP7 "Capacities" Specific Programme which closed in November 2009 called for proposals under the topic of 'Distributed Computing Infrastructures'. As a result of this call, six projects were funded, with an expected total EC contribution of nearly €50M, that together are referred to in this and related documents as the 'Distributed Computing Infrastructure Projects' or 'DCI Projects'. These projects are described in detail in Annex A and are summarised below:

- EGI-InSPIRE: Federation of national and domain specific resource providers into a European Grid Infrastructure for multi-disciplinary use.
- European Middleware Initiative (EMI): Continued support, development and harmonisation of the European middleware stacks from gLite, ARC, UNICORE and dCache.
- Initiative for Globus in Europe (IGE): Dedicated support for the European Globus community.
- European Desktop Grid Initiative (EDGI): To deploy desktop grids and cloud computing services for European user communities.
- StratusLab: Exploration of running production grid services in a cloud environment and providing cloud resources to research user communities.
- VENUS-C: Will explore and demonstrate the applicability of private and public cloud computing environments to different scientific applications to speeding up e-Science built on the sustainable public procurement of computing and storage resources on the cloud market.

Together the projects fund different activities in the area of distributed computing infrastructures ranging from the provision of production environments, the development, maintenance and support of the middleware used in Europe, and the exploration of the provision and use of virtualised computing resources. Due to the competitive nature of the funding model, many of these projects were developed in isolation in order to maintain confidentiality during the proposal phase. As a result, one of the goals required by the European Commission is that each project has to establish how they plan to collaborate with each other (if at all) and what the results of those collaborations may be within the scope of the project, and the impact that there may be long-term within the community.

This document shows how the provision of e-Infrastructures in Europe *could* evolve over the next 5 years and the contributions that each project may make towards this future by working with each other. As background, an overview of each of the six projects is provided in an Appendix.

It provides a record to the European DCI community as the potential results of the collaboration between the six distinct activities and the opportunities for collaboration it not only opens up between these projects but the wider community. It is essentially a technical document – describing the relationship between the projects and the technologies they will produce – and will it is expected become the basis for dissemination material to other interested stakeholders.

2. A Vision for European DCIs

2.1. What are DCIs?

Historically, a single data processing or generating resource (storage, computers, instruments, etc.) has been under the exclusive control of the administrative domain that owns it. However, some scientific, academic and research organisations, which already own these data related resources, increasingly need to securely share these resources with others. In order to federate their local resources into a production infrastructure, these organisations have had to establish mutual trust, adopted compatible middleware stacks and procedures integrated through operations teams to bring their resources together into a distributed computing infrastructure (DCI).

The recurring feature of the various DCIs that offer production resources (e.g. EGI, DEISA, PRACE, etc.) is that each one integrates multiple locally managed administrative domains into a usable environment. The middleware deployed by each DCI provides its users, according to their DCI credentials, consistent access rights to all resources managed by that DCI.

2.2. A decade of community building

The last decade has seen an unprecedented period of experimentation and prototyping in the collaborative use of distributed computing infrastructures. The EC funded European Data Grid (EDG) and Enabling Grids for E-science (EGEE) projects have built a collaborative infrastructure of primarily High Throughput Computing (HTC) resources to support intensive data analysis. The DEISA and other projects have focused on integrating an infrastructure of High Performance Computing (HPC) resources to support large-scale computing simulations. Together, with the provision of an integrated network of National Research and Educational Network (NREN) providers supported through the GEANT series of projects, these activities have been developing the core of a European e-Infrastructure service.

These EC funded projects have also provided a structuring effect in the geographical region around Europe. Infrastructure projects such as BalticGrid and SEEGrid linking the Nordic and Baltic states and South East Europe. As a result of this activity many of these countries are now part of the EGI-InSPIRE project. This structuring relationship between Europe and other regions around the world continues in the networking, computing and application space through several related projects.

The contribution from the European Commission to this activity has been a small but enabling contribution to the investments made by the national funding agencies. The EC investment has contributed towards the staff needed to bring these compute, storage and networking into a European infrastructure. The hardware and operating costs for these activities, in addition to funding the research undertaken on the e-Infrastructure, has all been funded outside the FP7 programme. For instance, in the 4 years of the EGI-InSPIRE project, the EC investment of €25M to the provision of a European Grid Infrastructure is a small proportion of the estimated €330M invested by the countries involved in the project in providing the European e-Infrastructure.

The provision of the European e-Infrastructure has been driven by the needs of the user communities that have needed access to large scale data analysis infrastructure to support their research needs as part of their pan-European research collaborations. Over the last decade, the European e-Infrastructure has benefited greatly from the growing maturity of the available open-source software solutions and where necessary have through middleware consortia such as gLite, UNICORE, ARC and Globus, and specialised technology providers like dCache, developed new, or extended existing solutions, in order to meet the needs of its user communities. These early adopting user communities have helped drive the development of the e-Infrastructure we have available today, which provides a production quality

federated resources, integrated through the middleware specifically developed to meet the demanding use cases coming from within the user and operations community.

2.3. Current Challenges

Even with the globally recognised achievements of the e-Infrastructure activities in Europe - delivering a production quality environment that supports a multi-disciplinary user community - broader adoption of e-infrastructures across the whole research computing community remains elusive. The reasons that other communities have not adopted the current e-Infrastructure offerings may range from:

- The data analysis challenges being faced by other communities have to date been within the scope of their current resources
- The usability and integration of non-local resources when compared to their desktop for solving problems is too high a barrier to overcome
- The service offering developed for the current user communities do not match the needs of other communities
- The future sustainability and governance of the e-Infrastructure to those communities that have not been actively involved in its development is not clear or assured
- The true cost of delivering a world-class data-intensive analysis infrastructure, regardless of the resources used to deliver it, needs to be exposed to the resource providers, the consuming end-user community and policy makers.

Recent activities within Europe are addressing these five issues.

The next generation of pan-European research infrastructures (the projects that are part of the European Strategic Forum on Research Infrastructures – <http://ec.europa.eu/research/esfri>) presents an opportunity for European e-Infrastructure providers to support a new wave of data-intensive research activities that will be highly dependent on a distributed computing and storage models. For these new communities, establishing and maintaining their own independent e-infrastructure is a diversion from their primary mission of doing science. Having access to a reliable European e-Infrastructure, available as a service, becomes an attractive option. It is also essential that any e-Infrastructure that they use in Europe be integrated with the e-Infrastructure used by their non-European research collaborations.

After a decade of investment in European e-Infrastructure, the production quality service offering now provided to the European Research Area are coalescing into two main areas:

- High Performance or Capability Computing provided currently by the DEISA and in the future the PRACE (Partnership for Advanced Computing in Europe) projects integrate high-end resources (generally of 10,000-100,000 cores) across Europe. Generally, these resources are used for closely coupled parallel applications for the few researchers with problems and applications able to benefit from them.
- High Throughput or Capacity Computing integrated into a European Grid Infrastructure (EGI) supported through projects such as EGI-InSPIRE and EDGI. These resources may include loosely or tightly coupled clustered, volunteer desktop or virtualised computing clusters contributed into a European infrastructure through national groupings of resource providers. Generally, these resources support the ‘bags of tasks’ applications where each task involves the execution of a program which needs minimal porting to run in such an environment.

The discussion for the remainder of this section will concentrate on the development of ‘high throughput’ computing resources both for single processor and parallel applications which are expected to provide the majority of resources and support the main stream of application communities in the years to come as new virtualised computational resources (currently available commercially on demand as cloud computing resources) are evaluated for integration into the publicly funded production infrastructure in Europe.

2.4. EGI and the DCI evolution

The goal of EGI is to provide a secure integrated federated production infrastructure constructed from national and domain specific resource providers, that is open to all users with potentially different computing models, needing access to different types of distributed resources (high-throughput, high-performance, desktop, virtualised, etc.), that are linked to physically remote data stores. Some of the high-performance computing resources may include some of those currently classed as DEISA resources.

Such an environment - a secure integrated sustainable production infrastructure - imposes constraints on those that produce software technology for deployment within it, and those that provide the resources to the infrastructure.

The EGI model is based around the contribution of resources from within different administrative domains where remote access is given to defined virtual organisations (groups of individuals coming together for a common goal) which will include users from different organisations. At the core of this model must be common mechanisms for establishing and entity’s identity (authentication) and the ability to control access to particular resources (authorisation). These mechanisms must be embedded into the access mechanisms for all resources to ensure a consistent reliable predictable security model.

Integration is necessary so that end-users are presented with consistent reliable secure interfaces to the same class of resource regardless of the resource provider and the implementation used to expose the resource to other users. To give resource providers the ability to deploy different software implementations to provide the same functionality it is necessary that the implementations demonstrate interoperability. The easiest route to achieve this is through the adoption of standards, and the verification of these interfaces through appropriate conformance tests. The ability of the interface deployed on a site to be available and to behave as expected is an aspect that is monitored remotely. High availability and the planned management of outages (reliability) is a vital aspect in defining a production (as opposed to a research) infrastructure.

In addition to monitoring the availability and reliability, an additional characteristic of a production infrastructure is the ability to account for its usage. This is important in order to understand current usage (of sites by virtual organisations) and to plan for changes that may occur in the future. Resources in a production infrastructure must therefore provide accounting information that allows usage by resource, resource type and virtual organisation to be tracked.

Against these technological constraints is a need for a sustainable operational model to be developed. The user communities planning to adopt the production infrastructure to support their research activities are doing so as part of a research programme that may persist for 10 or 20 years. Therefore the sustainability of the resources, the way they are funded, organised and operated, has to persist outside of any particular project. The European Grid Infrastructure (EGI) is now coordinated on behalf of the community by a dedicated organisation (EGI.eu) funded and responsible to the community it serves. It provides the centralised coordination necessary to bring together individual resources providers, either public or private, to deliver a secure integrated infrastructure, and as a means to gather and prioritise requirements from the resource providers and user communities as to how the infrastructure should develop.

2.5. The DCI common vision

The pressures (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 space. 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.

This ability to provision resources ‘on-demand’ to meet the the needs of particular research collaboration provides significant challenges to resource providers in the research space. In the commercial world ‘cloud computing’ has provided a ‘pay per use’ business model that has 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’. 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.

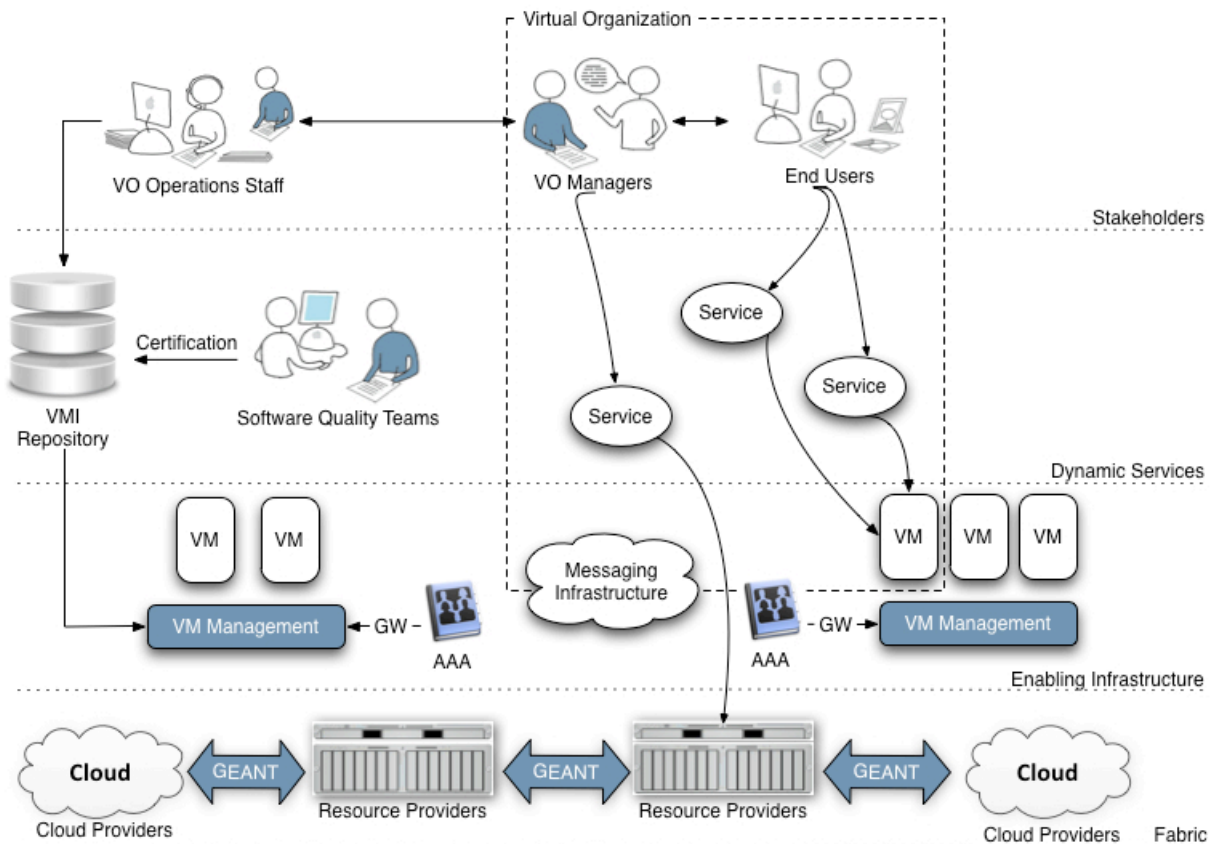


Figure 1 - Virtualised Federated Resources

In the research world the trend to consolidate data centres will also be continued to be balanced with the need to collaborate to share resources. This will lead to federated grids of virtualised resources (See Figure 1), in many ways similar in architecture to today's federated grids of computing resources, but providing truly generic infrastructure that can be accessed by any authorised research collaboration, as an alternative or alongside commercially provided resources. This virtualised infrastructure will be used to instantiate a platform to support particular research collaborations. These platforms (which will be comparable to the current gLite, Globus, ARC or UNICORE environments) may be deployed directly by the research collaboration using pre-defined images, by using bespoke images created from within the collaboration, or provided as a service by third-parties within the ecosystem. To many of the end-users within the research community who take no interest in the details as to how their infrastructure is provided, the result will appear as just a set of services available 'out there' for them to use.

2.6. Implementing the vision

A grid of virtualised resources, with federation taking place within a region, national borders, or across the European Research Area has many potential benefits, opportunities and challenges for end-users in the research community, service providers within the research community, commercial organisations wishing to engage and provide services and resources to the research community, and other organisations for establishing the policy environment for such a federated infrastructure to operate within.

The following sections discuss some of the challenges and potential benefits of the proposed model for some of the participants within this vision for DCI provision within Europe.

2.6.1. For Infrastructure Providers

This roadmap for European DCI provision provides many benefits. The alignment of infrastructure provision in the research e-Infrastructure community with models used in the commercial world - provision of end-user environments through virtualisation - allows tools and techniques used in industry to be adopted in academia. This approach has already demonstrated increased server utilisation, better energy utilisation and greater flexibility in the commercial world. Tools developed to meet the management challenges used in these commercial environments may also prove effective in the research environment.

This virtualised environment will allow resource providers to deploy virtual images on demand to meet the needs of different user communities. The flexibility provided through a trusted repository of virtual images would allow resource providers to support a greater number of different environments and therefore a greater number of different user communities. This provisioning activity may be undertaken directly by the local resource provider or by authorised third-parties, e.g. from other resource centres, by representatives of the user communities or by other authorised entities. Such a model requires a trust model between the local resource provider and the generator of the virtual image. Policy and technical discussions around this area are ongoing and conceptually such a trust model is similar to that currently used for the pilot job frameworks used within the High Energy Physics community where the resource provider delegates the actual payload executed in their machine to a trusted third party. Implementing such a model requires work within the community to manage the distribution of virtual machine images, mechanisms for image signing and site policies for accepting images based on signed images to create trusted image repositories.

A virtualised infrastructure that allows environments to be deployed on demand by authorised groups allows a different security model to be used for the provisioning activities than is used within the virtual machine image. Clearly, there is a need for the infrastructure provider to be reassured as to the activity that will take place within the instantiated virtual machine, depending on what that activity might be. The level of authorisation and logging that may be needed within the virtual machine might be conditional on

the end-user control of what takes place within the virtual machine. A similar conditional policy on the levels of authorisation is in place for portal access to the e-infrastructure.

2.6.2. For the Software Provider

Large-scale adoption of virtualisation by the infrastructure, and the effective management of the software deployed within the virtual images, imposes operational requirements on the software services. Once the virtual machine is running consistent service management and monitoring interfaces are needed to configure the services within its instantiated environment (e.g. What services should be run? What certificate should be used? Who is allowed to access these services? etc.) and to monitor their operation and health. Providing consistent standardised interfaces enables third-party management tools and protocols to be used to support manual intervention by the operations staff.

The loosely-coupled dynamic nature of this infrastructure needs a flexible system for linking the virtual machine hosting environments, the transient virtual machine images that run on the hosting environments and the services within the virtual machine image itself. Modern messaging systems have been designed for use in just such a distributed environment through the ability to have different messaging queues and provide persistent message delivery. A messaging infrastructure will underpin the future DCI and should be used as the basis for messaging and management by the deployed services.

The messaging infrastructure provides a basis for higher level applications to build upon. This includes existing operational functions such as accounting and service monitoring, and provides a basis for research into new operational tools such as autonomic management of the infrastructure. As the scale and complexity of the infrastructure increases, autonomic management functions become essential - to recognise when virtual machine instances or services have stopped working and to restart or redeploy these instances to ensure the required services remain available to the user communities.

2.6.3. For the end-user

For end-users in the research community, easy usage of e-Infrastructures is essential – regardless as to who operates it or the technology used to deliver these services. A federated virtualised infrastructure presents many additional benefits. It provides a means for the user communities to deploy within the infrastructure the services that they wish to use when they wish to use them. These services will need to be encapsulated into a virtual machine image and be able to meet the policy requirements imposed by the infrastructure relating to security, configuration, management, monitoring, etc.

Provisioning of this infrastructure for end-user communities may come from many sources. For communities that have the required technical knowledge they may generate their own customised virtual machine images to the appropriate specifications and deploy these to the virtualised resources that they need to use. Other communities may work with experts outside the community to have an environment created, deployed, managed and monitored on their behalf. The resource providers may provide a basic generic environment over some resources to provide a resource for communities that do not need a customised environment. The execution environments required by an end-user could also be made available through commercial cloud infrastructures alongside other resources that are able to provide dynamic scale-out capabilities, which do not require long-term resource and organisational commitments. This model provides much greater flexibility to the end-user community as to the environment that is available to them - if they need to exploit such flexibility.

The use of messaging as a fundamental part of the infrastructure provides flexibility in how the end-user interacts with the distributed resources. For instance, it provides the ability for the user to easily subscribe to events that they are interested in - when an application starts running, when it stops, or if it fails to complete. Results can be sent back to the user through the messaging system. As the messaging system is

DCI Collaborative Roadmap – September 2010

capable of asynchronous delivery, it allows results to be stored and then delivered when the user is ready - for instance when they reconnect their laptop to the network in the morning.

3. Individual Interactions

To support the development of e-infrastructure provision in Europe and the vision outlined in the earlier sections the DCI projects will be working on various collaborative aspects. These are summarised in the table below and detailed in the remainder of this section.

3.1. EGI-InSPIRE

The EGI-InSPIRE project's main focus is to deliver a production infrastructure for the European Research Area. In order to deliver this, it needs to deploy a software environment that brings together software components provided from both within and from outside the DCI community.

Two projects are seen to initially provide these software components:

- European Middleware Initiative (EMI)
- Initiative for Globus in Europe (IGE)

A Memorandum of Understanding (MoU) will be established with each project to describe common plans around dissemination, representation to ensure the exchange of requirements and the development roadmaps. Specific Service Level Agreements (SLAs) will be defined to govern the expected operational interactions on the provision of third line support and security incident handling.

It is envisaged that future DCIs will make extensive use of virtualisation technology. An MoU will be established with StratusLab to ensure effective joint dissemination and events, where applicable, and for the operational staff within EGI to ensure that the software environment being released from StratusLab will be the needs of the production infrastructure. This may include requirements on reliability, scalability, monitoring and accounting.

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 mean 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. EGI-InSPIRE would like to provide input to VENUS-C on the criteria for the experiments that are selected from the open call, to ensure they are of relevance to the EGI user community.

It is expected that the EDGI project will build a desktop resource across Europe. EGI-InSPIRE would like to ensure that this computing resource can be integrated alongside the resource types offered within the production infrastructure. For this end, EGI-InSPIRE will collaborate with EDGI through an MoU that will establish the monitoring, accounting and functional integration of desktop resources into EGI. As a result of this the reliability and use of this resource can be established alongside the other provided resource types and these desktop resources can then be offered up to the EGI user community alongside the others.

3.2. EMI

As one of the major providers of middleware services for the Distributed Computing Infrastructures, the EMI project will establish interactions with both infrastructure providers deploying the services and other middleware and application developers complementing or extending the EMI services. In particular, within the group of DCI projects described in this document the following interactions have been identified after the first four months of operations of EMI in numerous discussions with the relevant projects.

DCI Collaborative Roadmap – September 2010

Table 1 DCI Interactions summary

		Consuming Project					
		EDGI	EGI-InSPIRE	EMI	IGE	StratusLab	VENUS-C
Providing Project	EDGI		Dissemination and integration of desktop resources.	Technology bridging through interoperability and standardisation.	Interest from UK NGI, EDGI-SW, GRAM gateway development, requirements.		Adapting desktop grids to be run over cloud resources
	EGI-InSPIRE	Dissemination.		Deployment, requirements, dissemination, feedback & usage.	Deployment, requirements, dissemination, feedback & usage.	Requirements for integrating virtualisation into the operational infrastructure.	Promoting Open Call toward EGI user communities
	EMI	Technology previews.	SLA defined middleware maintenance and support. Dissemination and training.		Requirements for Globus, standardisation and interoperability.	Middleware able to run on OpenNebula.	Middleware able to run in VENUS-C (under defined constraints).
	IGE	Globus support and GRAM gateway development.	SLA defined middleware maintenance and support. Dissemination and training.	Standardisation and interoperability. Support for Globus components.		Support of StratusLab Cloud Platform, support globus.eu, investigate dynamic Grid deployment.	Adapting globus-enabled application to run over cloud resources
	StratusLab	Virtual appliances for Grid services	Requirements dynamic deployment of virtualised grid sites.	Requirements for Virtual appliances for Grid services, access to virtual testbeds	Hosting of globus.eu (using Amazon API) & requirements for dynamic Grid deployment		Adopting Stratuslab toolkit as IaaS middleware in one or more VENUS-C sites.
	VENUS-C	Supporting desktop grids on top of cloud resources.	Best practices of scientific communities using clouds. Opportunity for EGI communities to experiment with the VENUS-C platform through the Open Call	Requirements for Virtual appliances for Grid services, access to virtual testbeds, Sharing information on accounting data formats and approaches	Supporting Globus users in experimenting on cloud resources.	Providing feedback on usage requirements and user experience on the IaaS approach.	

3.2.1. EGI-InSPIRE

The European Grid Infrastructure, supported by the EGI-InSPIRE project, is the main user of the EMI services. EGI represents therefore the major source of requirements and the primary target for the delivery of software and support services from EMI. EMI and EGI have discussed in several occasions the points of interactions and a common vision and common plans for maintaining and evolving the European research infrastructures. The relationships between EMI and EGI will be formalised with the establishment of a Memorandum of Understanding and the negotiation and signature of a commercially-oriented Service Level Agreement. In particular the SLA will describe in details how the two parties commit to provide and access the EMI services and what levels of service quality are expected. It is EMI intention to establish together with EGI a prototype of possible future professional service provision relationships that could extend to commercial providers.

In summary three major collaboration points have been defined:

Requirements: the collection of requirements is an essential part of the middleware development lifecycle. EMI must be sure that what is provided by its Product Teams is relevant, usable and able to support the EGI roadmap vision for the research infrastructures. EMI will therefore actively take part in the definition of the EGI UMD Roadmap by means of mechanisms provided by EGI, like the Technology Collaboration Board (TCB). EMI will provide clear deadlines for the releases of its software services with particular attention to the existing and new functional requirements discussed with EGI and its user communities. In addition, EMI will work with EGI on the definition and implementation of the Infrastructure Roadmap and the integration of existing middleware services with emerging computing and data storage technologies.

Maintenance and Support: Although EMI has to evolve the middleware services towards the implementation of the DCI vision, the continuous operational efficiency of the infrastructures has to be guaranteed. EMI together with EGI will continuously monitor and assess the quality of the software developed by EMI and deployed by EGI in order to react to any incident or user request in the most professional manner. Clear criteria for transitioning the services from EMI to the EGI roll-out service will be defined and periodically revised. Clear support policies, service lifetime policies and migration paths from old to new services will be defined by EMI and EGI and formalised in commercial quality Service Level Agreements, which will be periodically revised and improved. EMI and EGI together will also explore possible alternative model for supporting the middleware involving commercial partners whenever feasible and desirable.

Dissemination and exploitation: EGI and EMI together represent a large part of the European and international infrastructures in support of scientific research communities. An efficient and timely dissemination of information and the expansion of the user base are keys to the correct exploitation of the infrastructures. EGI and EMI have therefore engaged in establishing common dissemination strategies to provide coherent and complete perspectives on the various components of the infrastructures and their applications. A first important result of this engagement will be the joint organization of a major international event during Spring 2011 to bring together existing and new scientific user communities, presents achievements and results and collect trends and ideas.

3.2.2. IGE

The Initiative for Globus in Europe is a provider of middleware aiming at creating an official link between European users of Globus and the US developers maintaining it. A number of Globus components are currently used within the EMI services to provide specific functionality. In addition, a number of services developed by EMI and Globus are providing similar, but not always interoperable

functionality. EMI and IGE have therefore identified two major areas of collaboration, which will be formalised with an appropriate MoU:

Standardisation and Interoperability: EMI and IGE will work together and as part of other standardization and interoperability bodies to defined and evolve standards for the distributed computing middleware, especially in the areas of Compute and Data Management.

Support and maintenance: EMI needs to rely on continuous support from the Globus developers and maintainers in case of software issues. EMI will also provide requirements for Globus to IGE as needed and will monitor together with IGE the implementation of those requirements in the Globus releases. In exchange EMI will gradually move from its current usage of Globus, distributed as part of the EMI middleware services, to a more standard use of official Globus packages maintained by IGE and distributed as part of the major Operating Systems distributions, like Fedora or Ubuntu.

3.2.3. EDGI

EMI and the European Desktop Grid Initiative are software providers for EGI with strongly complementary roles, since they maintain and promote different types of distributed computing middleware for different sets of use cases. There are a number of interesting common points that link EMI and EDGI at the boundary where the two technologies meet. Essentially part of the EDGI services makes use of middleware services provided by EMI to bridge standard grids and desktop grids. The work of bridging the two types of grid was already started in previous project, but EMI and EDGI are now planning to work together in completing such bridges and providing access to resources not only via gLite, but also via ARC and UNICORE and future standard-based resource management clients. EDGI will work with EMI as part of the ‘Works with EMI’ technical collaboration program that allow technology providers and consumers to have direct access to technical previews and dedicated support for complementing and extending the EMI services.

3.2.4. VENUS-C

VENUS-C is providing both platform APIs and resources to enable scientific users to access commercial cloud providers or (public or private) data centers. As part of the overall DCI vision for how the research infrastructures will be shape in the coming years, it is clearly acknowledged that cloud or similar dynamic service provision models will be more and more used. EMI is therefore fully committed to understand how the existing distributed services can be improved and evolved to exploit such service provision models while retaining their existing flexibility and security. EMI and VENUS-C are discussing on the possible integration paths across grid and cloud platforms. EMI will put effort in introducing any modifications in it services to make them fully compatible with the VENUS-C infrastructure. At the same time VENUS-C will provide EMI with access to technology and resources to validate and test the EMI services on virtualized environments. Common work on security and accounting formats enabling interoperability will also be considered.

3.2.5. StratusLab

StratusLab is providing software to setup distributed computing infrastructure based on the emerging cloud technology. As part of the overall DCI vision for how the research infrastructures will be shape in the coming years, it is clearly acknowledged that cloud or similar dynamic service provision models will be more and more used. EMI is therefore fully committed to understand how the existing distributed services can be improved and evolved to exploit such service provision models while retaining their existing flexibility and security. EMI and StratusLab are discussing on how existing EMI grid services can run on virtualized environments based on OpenNebula and which modifications are needed in the services configuration capabilities to make them able to be instantiated as on-demand services or pre-configured appliances. StratusLab will provide EMI with requirements and with access to testbeds, while

EMI will incorporate and support in its releases the functionality needed to exploit virtual environments based on StratusLab technology.

3.3. IGE

The IGE project strives to integrate as tightly as possible with the other DCI projects in order to deliver a convincing user experience to the scientists within the European Research Area. To this end, interactions are to be established as follows:

3.3.1. EGI-InSPIRE

The main goal is to ensure collaboration for the integration activities with respect to Globus-contributed infrastructure, wherever appropriate. This includes acting as a software provider towards EGI by setting up reasonable SLAs, delivering Globus and Globus-related components to the UMD, and contributing training and support where necessary. Moreover, IGE will ensure the appropriate representation of European Globus-based research communities within the Virtual Research Environments.

As a first concrete action, a Memorandum of Understanding (MoU) will be established that describes common plans around dissemination, representation, and exchange of requirements and development roadmaps.

3.3.2. EMI

IGE will strive to become the main provider of Globus components within the EMI software stack. To this end, a close collaboration with EMI will be setup in order to ensure continuous support for the US-based Globus package distribution, and to collect additional requirements from EMI, such as more standard installation and deployment procedures (through major Operating System distributions such as Fedora and Ubuntu). Moreover, both projects will collaborate in the area of standards and interoperability in the areas of Compute and Data Management for DCI middleware.

As a first concrete action, the modus operandi of this interaction will be detailed in a MoU between EMI and IGE.

3.3.3. StratusLab

One major goal will be to work with StratusLab on hosting the anticipated globus.eu branch of the newly developed Globus.org SaaS platform on StratusLab infrastructure. In this context, the dynamic deployment within a Grid environment (i.e. submitting VMs instead of traditional jobs through a Compute Service interface such as GRAM). Moreover IGE will interact with StratusLab to ensure that (a) the StratusLab Cloud Platform by itself and (b) the new delivery paradigms are supported within Globus; in particular, it is to be ensured that Globus is compatible with the IaaS interfaces and the creation of VMs.

To this end, IGE will frequently test the StratusLab innovations with Globus and regularly provide feedback throughout the whole development.

3.3.4. EDGI

Recently, NGI-UK has formulated the requirement to utilise Desktop Grid resources as part of the national infrastructure. Since Globus is part of the middleware stack here as well, IGE will collaborate with EDGI to develop, contribute, and maintain a Desktop Grid Bridging Service for Globus to cater the need for EDGI integration.

3.3.5. VENUS-C

IGE offers to support VENUS-C in understanding the special requirements of traditional Grid users for enabling their applications to use Cloud infrastructures, how such use cases can to be deployed, and by

providing access to Globus, Grid application, and test resources. IGE will also provide technical assistance and expertise to VENUS-C where required on pertinent aspects of platform interoperability with Globus, e.g., concerning Security and Authentication.

3.4. EDGI

As a collaboration activity between EGI.eu and EDGI, the EGI.eu dissemination and training channels will be used to reach the existing EGI user communities. In order to get the highest possible impact EDGI will organize dissemination events in the framework of events organized by EGI.eu. International Desktop Grid Federation and EDGI will work in strong collaboration with EGI, EMI, NorduGrid, UNICORE Forum and interested NGIs in order to reach the widest possible user and resource provider communities. Most of the dissemination work of EDGI will be done in the framework of the European Chapter of the International Desktop Grid Federation.

The Desktop Grid Federation is set-up to be long-lived, i.e. after the EDGI project has finished. We try to align it as much as possible with existing e-Infrastructure organisations, such as EGI, so it could also be possible that (part of) the Federation could become a user group in these e-Infrastructure organisations. The International Desktop Grid Federation will organise the grid operators and application developers in the European Union. It will strongly collaborate with the International Desktop Grid Federation run by DEGISCO in order to organise the grid operators and application developers outside the European Union, especially the ones in the ICPC countries.

Standardization activities will be carried out through several channels. The EDGI Bridge will use the HPC profile job submission mechanism (an OGF standard) in order to guarantee the interoperability with every Service Grids that follow this standard. In particular, the UMD developed in EMI currently follows this standard that makes sure that the EDGI Bridge will be compatible with the middleware supported by EGI.eu. In order to maintain this compatibility for the whole duration of the project and beyond, EDGI will strongly collaborate with EMI.

EMI will concentrate on the major Service Grid middleware systems and will further develop ARC, gLite and UNICORE towards making them interoperable and based on them will create a unified middleware distribution, but will not cover any Desktop Grid extension of these middleware systems. EDGI will cover this important area in the e-science infrastructure eco-system. The objectives of EDGI and EMI are complementary (both want to further develop middleware) but technologically orthogonal (the middleware to develop are different). EDGI will carefully follow any improvements and further developments of ARC, gLite and UNICORE created by EMI in order to make sure that the Service Grids → Desktop Grids bridge middleware developed by EDGI will be compatible with any new versions of the ARC, gLite and UNICORE middleware.

EDGI will strongly collaborate with DEGISCO that is a support action project to disseminate the results of the EDGeS project outside the EU countries. Since EDGI is also a follow-up project of EDGeS and aims at disseminating desktop grid related knowledge in EU countries there are many commonalities between the two projects.

EDGI is furthermore open for any DCI project to use Desktop Grid resources. One possibility is to provide a solution for Globus users to transparently and seamlessly utilise Desktop Grid resources through the EDGI Bridge. To support this idea, EDGI will investigate the possible alternatives together with the IGE DCI project. One potential user for the Globus → Desktop Grids Bridge is the UK NGS.

3.5. StratusLab

StratusLab is open to collaboration with all DCI projects with an interest in using cloud resources.

3.5.1. EGI-InSPIRE

StratusLab will deploy grid sites over cloud infrastructures that will join the EGI infrastructure. The operation of virtualized sites may require adjustments in the way grid resources are certified, managed and operated. A virtualized grid site will expose elasticity and volatility at a level not previously experienced in operational sites.

The two projects will have to collaborate closely in order to ensure that the operational models implemented by EGI will be cloud-friendly and flexible enough in order to take advantage of the merits brought by cloud computing. Additionally cloud software should be enhanced in order to enable dynamic provisioning and configuration of grid-resources permitting the provision of grid-sites-on-demand.

3.5.2. EMI

The European Middleware Initiative project is responsible for the support and development of the Unified Middleware Distribution (UMD) a term used to refer to an integrated distribution of the most popular European grid middleware, namely gLite, UNICORE and ARC. Apart from the existing high-level collaboration among DCI projects, there is also room for collaboration on a technical level. StratusLab will be using UMD to deploy grid services on top of cloud infrastructures. In many cases technical restrictions may impede the efficient installation and operation of these grid services. A channel of interaction among the two projects would be important in order to convey problems and requirements. The final goal of the above collaboration will be at the end of the projects to have UMD and the StratusLab software distributions to be fully compatible.

3.5.3. IGE

IGE will develop grid services and tools that can take advantage of the Cloud. StratusLab could provide to IGE requirements and an architecture for Grid sites taking advantage of Cloud concepts and technologies, and later cloud-enabling Globus-based infrastructures using StratusLab architecture and tools. IGE could provide to StratusLab requirements from Globus-based infrastructures, and later Globus services (e.g. GRAM for job execution) and tools (e.g. GridWay Metascheduler) able to operate on a StratusLab-based infrastructure and to take advantage of Cloud concepts. StratusLab and IGE could work together (probably at a technical level) on cloud-enabling of Globus services and tools using the StratusLab distribution.

3.5.4. VENUS-C

StratusLab and VENUS-C share a common interest in cloud computing. VENUS-C will develop and deploy a Cloud Computing platform service for scientific communities in Europe. StratusLab can provide VENUS-C with the StratusLab toolkit, as a comprehensive, open-source private cloud distribution for the VENUS-C IaaS backend to form part of the VENUS-C infrastructure on one or more sites. StratusLab will be interested in the requirements of VENUS-C for the IaaS platforms used in their infrastructure. Also, some partners of VENUS-C plan to evaluate OpenNebula as an IaaS solution, their feedback would be valuable. VENUS-C partners may provide adaptations of OpenNebula components made for their deployments. (e.g. development or tuning of the storage/network/virtualization plugins) StratusLab and VENUS-C may work together on the definition of extensions of current IaaS API's.

3.5.5. EDGI

StratusLab will produce and maintain a repository of virtual appliances for grid services. EDGI may use these virtual machine images and supply related requirements,

3.6. VENUS-C

In scientific computing, there are user communities which traditionally have no strong need to utilize grid and supercomputer facilities. Instead, they utilized local smaller HPC compute clusters to do their

simulations. In recent years, these user communities are tackling more and more complex computational problems. The VENUS-C project's goal is to equip these users with a tool set to easily scale their existing scientific workloads into the publicly available cloud resources. The expectation is that cloud resources are instantly provisioned, and that the users can scale out their workloads quickly at predictable costs. In order to keep the barrier of entry as low as possible, the project aims to keep the software dependencies and requirements as minimalistic as possible. A scientist should be able to test a scientific executable or script on his laptop, and then easily scale out the job in a map/reduce fashion to his resources in a public cloud.

The project will enable seven existing e-Science applications as part of the original workplan, as well as an additional of up to twenty selected applications from an open call, to run in the cloud. Several VENUS-C project partners donate a significant amount of both IaaS- and PaaS-based cloud resources (compute, storage, transfer) to the project's scenario partners and the open call participants.

3.6.1. EGI-InSPIRE

As part of its open call ready for communication in the latter part of 2010, the VENUS-C project will financially support up to 20 selected applications to run their applications 'in the cloud', using both the VENUS-C tool set and the allocated cloud computing and storage resources. In order to obtain feedback from a broad and heterogeneous set of scientific applications, the VENUS-C project invites and encourages the EGI-InSPIRE project and the EGI user communities to engage actively in the open call. Broad participation in the open call will enable the VENUS-C project to assess the applicability of cloud computing to scientific communities, and a diverse set of applications will enable the identification of specific communities which may particularly benefit from the adoption of cloud computing. From an econometrics point of view any legal, and socio-economic findings on the suitability of adopting the cloud model for the EGI service provisioning could be shared between the projects.

3.6.2. EMI

EMI will provide a reference implementation for grid middleware in Europe. The trends in virtualised infrastructure and cloud computing will allow EMI to experiment such technologies and to better understand the cloud model. One opportunity for EMI and VENUS-C is to assess how traditional grid middleware (as supported by EMI) relates to a computational model put forward by VENUS-C. One concrete technological area for collaboration between EMI and VENUS-C is security and accounting. In particular, it would be interesting for VENUS-C to learn from EMI about the EMI project's consolidated vision for accounting scientific users across different middleware platforms.

On the other hand, VENUS-C can support EMI in adopting both the cloud model and understanding the cloud technology, by providing access to best practices and resources to experiment and test the EMI services, at a cost of developing any bridge or slight adaptation of their services.

3.6.3. IGE

VENUS-C can support IGE in allowing the project, as a facilitator to Globus users, to understand how to adopt the cloud model and understanding the cloud technology, by providing access to best practices and resources to experiment and test Globus. With budget and resources permitting, IGE may look at developing an interface of Globus and VENUS-C. In addition, IGE may, as in EMI, also share information and experiences on the aspect of Authentication and Accounting in order to speed-up the interoperability amongst the different platforms.

3.6.4. EDGI

VENUS-C and EDGI may explore in the coming months the opportunity and feasibility of running desktop grids ontop of Cloud services, providing feedback both on the desktop grid applications and on the Platform APIs to increase the adoption of the Cloud environments by the scientific communities.

3.6.5. StratusLab

VENUS-C aims to deploy scientific workloads both on IaaS and PaaS cloud offerings. The StratusLab project aims to develop an IaaS StratusLab Cloud Distribution. Depending on the schedule and availability of the StratusLab toolkit, it might be possible to conduct an experiment to utilize a StratusLab-based grid resource from within VENUS-C. Given that OpenNebula is a relevant cloud platform in both VENUS-C and in StratusLab, it is desirable to exchange user and development experiences between the two projects. Therefore, VENUS-C and StratusLab may jointly promote the evolution and the adoption of standardised API at Infrastructure level, sharing best practices and interfacing with addressed standard bodies.

4. Conclusions

Even with the committed investment from the European Commission through the FP5, 6 and 7 programmes and the member states of the European Union, a clear challenge remains for European e-Infrastructure providers. To ensure their longer-term sustainability they need to be seen as providing a reliable and efficient service to all user communities in Europe needing to use research computing and storage services. This is essential in order to be able to attract a broad base of European and national research funds. However, in order to attract these user communities a broader range of services need to be provided for these individual communities within the same (and probably reduced) operational costs.

New technology offers a route for resource providers in the research sector to deliver these services with greater reliability, scalability and efficiency. However, such a route is not without its challenges. Firstly, it requires fundamental changes in how services are provisioned in the research community that builds on the experiences gained in the commercial space by using virtualised infrastructure to produce in some sectors a so-called cloud business model. Secondly, it requires open-source software community to adapt their software to be managed, monitored and deployed within a federated virtualised environment by focusing on delivering services that are not available elsewhere. Thirdly, that the user communities find the services offered to them attractive and easy to use so that they can be incorporated into their data analysis workflows. Finally, it provides some challenges to the computer science community to provide solutions that enable the reliable and effective management of such a highly distributed infrastructure.

Together, the DCI projects are able to address some of these issues and move the community towards the presented vision of an integrated virtualised infrastructure for the Europe Research Area. EGI-InSPIRE will provide 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. EMI and IGE provide a source of innovation in the short-term, and it is expected this will be expanded over time to include the technology and procedures developed within the StratusLab project. VENUS-C will 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.

The vision presented in this document transitions the DCI community to providing an integrated infrastructure as a service and for EGI to help bring the technology innovations being developed within the DCI community and elsewhere through to use in the research and public sectors. Interoperability and integration between different e-Infrastructures and technologies is fundamental to the DCI projects and the work that will be undertaken between them. This activity will help contribute to the broader vision described in the recent Digital Agenda for Europe (DAE¹) communication that “Europe should also build its innovative advantage in key areas through reinforced e-Infrastructures”.

¹ A Digital Agenda for Europe : Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions - COM(2010) 245, 19.05.2010 a digital agenda for Europe

5. Annex A: The Distributed Computing Infrastructure Projects

Each of the six Distributed Computing Infrastructure (DCI) projects are summarised in the following table and described in more detail in the following section.

Project	EDGI	EGI-InSPIRE	EMI	IGE	StratusLab	VENUS-C
Website	edgi-project.eu	www.egi.eu	www.eu-emi.eu	www.ige-project.eu	www.stratuslab.eu	www.venus-c.eu
Start Date	01/06/2010	01/05/2010	01/05/2010	01/10/2010	01/06/2010	01/06/2010
Duration (months)	24	48	36	30	24	24
Total Budget (€)	2,436,000	72,000,000	24,000,000	3,693,000	3,137,221	8,803,046
Funding from the EC (€)	2,150,000	25,000,000	12,000,000	2,350,000	2,300,000	4,500,000
Total effort (person months)	281	9241	2319	277	340	639

5.1. European Grid Infrastructure - Integrated Sustained Pan-European Infrastructure for Researchers in Europe (EGI-InSPIRE)

EGI-InSPIRE will support the establishment of a sustainable model for a European Grid Infrastructure (EGI) that integrates resources contributed by national and domain-specific resource providers. Key to this process is a new organisation, EGI.eu, coordinator on behalf of the European resource provider community of the EGI-InSPIRE project, which is also more broadly dedicated to coordinating the EGI community on behalf of its stakeholders.

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

The production infrastructure supports Virtual Research Communities – structured international user communities – that are grouped by specific research domains. Virtual Research Communities are formally represented within EGI at both a technical and strategic level. Direct support is coordinated through a central helpdesk, that brings together operational, technology and other support teams from within the EGI-InSPIRE project and other partner projects.

The EGI-InSPIRE project focuses principally on the European production infrastructure, it needs to support the collaborative research needs of its user communities, for their resources to be integrated with infrastructures around the world. In addition to over 40 partners located within geographical Europe, EGI-InSPIRE includes 8 unfunded partners from the Asia Pacific region. Strong collaborations are also expected with infrastructures in North America, and the emerging infrastructures in Latin and South America. EGI-InSPIRE will support and develop the European DCI community in three important ways:

- Integrate resource providers within the National Grid Initiatives (NGIs) and European International Research Organisations (EIROs).
- Support the development of policies to ensure effective technical management, integration and operation of the EGI for its user communities.
- Coordinate the development and support of structured Virtual Research Communities currently using the production infrastructure within the European Research Area.

EGI-InSPIRE's will provide services to the community through the partners within the project to:

- Operate a secure, integrated, reliable pan-European infrastructure that can support diverse science communities through the deployment of different technology solutions.
- Work with external technology providers (initially EMI and IGE) to ensure that their solutions meet the needs of the operational and user community in terms of reliability, scalability and functionality.
- Provide support to the communities that rely heavily on the infrastructure by supporting the shared services and tools that are common to many of them

Together these services will support a virtuous feedback circle – starting with a set of integrated services on the production infrastructure that meet the needs of its users, working with external technology providers to define new or improved services based on these existing services, assessing the quality of the new delivered services, followed by their deployment in to the production infrastructure.

Additional effort within the project will develop the operational tools to fully devolve these to an national rather than a central operational model, while ensuring that resources such as HPC, desktop grids and virtualised resources are fully integrated into the monitoring and accounting infrastructure.

5.2. European Middleware Initiative (EMI)

The European Middleware Initiative is a close collaboration of the three major middleware providers, ARC, gLite and UNICORE, and other software providers. It will deliver a consolidated set of middleware components for deployment in EGI (as part of the Unified Middleware Distribution or UMD), PRACE and other DCIs, extend the interoperability and integration between grids and other computing infrastructures, strengthen the reliability and manageability of the services and establish a sustainable model to support, harmonise and evolve the middleware, ensuring it responds effectively to the requirements of the scientific communities relying on it.

European scientific research has benefited in the past several years from the increasing availability of computing and data infrastructures that have provided unprecedented capabilities for large scale distributed scientific initiatives. A number of major projects and endeavours, like EGEE, DEISA, WLCG, NDGF, OSG, See-Grid, BalticGrid and others, have been established within Europe and internationally to share the ever growing amount of computational and storage resources. This collaborative effort has involved hundreds of participating research organizations, academic institutes and commercial companies. The major outcome is a number of active production infrastructures providing services to many research communities, such as High Energy Physics, Life Sciences, Material Science, Astronomy, Computational Chemistry, Environmental Science, Humanities and more.

At the core of these rich infrastructural facilities lies the grid middleware, a set of High Throughput Computing (HTC) and High Performance Computing (HPC) software services and components that enable the users to access the distributed computing and data resources, execute jobs, collect results and share information. Middleware like gLite from the EGEE project, ARC from the NorduGrid Collaboration, UNICORE, VDT, Globus and other specific services for computing and data management have allowed thousands of scientific researchers to access grid-enabled resources and produce scientific results.

After the necessary initial period of research and consolidation that took place in the past 6 to 8 years, the growing usage of distributed computing and data resources by scientific communities and individual researchers requires now the stabilization of the computing infrastructures and a simplification and standardization in the use of the associated software tools. It is of strategic importance towards the

establishment of permanent, sustainable research infrastructures to lower the barriers that still prevent potential communities of tens of thousands of scientists and researchers to consider grids as a commodity tool serving their daily research activities. The ultimate vision is that establishing distributed scientific collaborations and using distributed computing and data resources should be as easy as opening a web application, entering simple identification information, entering a few clear parameters to define the task to be executed and its requirements and then waiting for the results to be made available in a well known, easily accessible place.

The EMI project will make the realization of this vision possible by addressing and solving a number of problems that today still prevent users from easily accessing and using the existing computing infrastructures:

- Usability will be enhanced by removing redundancy and consolidating the services, simplifying the security management without compromising its strengths, adding integrated support for high level gateways and portals and transparently making use of virtualization to increase resource availability and management.
- Compatibility will be improved by removing proprietary interfaces in the middleware services and ensuring true interoperability through the adoption of agreed community standards.
- Manageability will be improved by providing standard service configuration, monitoring and instrumentation interfaces and making accounting and other operational information more readily accessible.
- Interoperability between grids, supercomputers and emerging computing models like clouds and desktop grids will be extended to address scalability and accessibility requirements.
- Sustainability will be improved by establishing collaboration programs with commercial companies, adopting off-the-shelf components to reduce maintenance costs and to facilitate easier adoption by wider user communities. The definition together with the resource providers of measureable Service Level Agreements will provide the base for establishing more standard service provision business models.

5.3. Initiative for Globus in Europe (IGE)

The Initiative for Globus in Europe, IGE, serves as a comprehensive service provider for the European e-infrastructures regarding the development, customisation, provisioning, support, and maintenance of components of the Globus Toolkit, in close collaboration with the European Grid Initiative (EGI), Distributed Computing Infrastructure (DCI) projects, and Standard Development Organisations (SDOs).

By coordinating the European Globus activities, IGE drives forward Globus developments according to the requirements of European users and strengthen the influence of European developers in the Globus Alliance. This strengthens the representation of European topics such as security and privacy, data privacy protection, compatibility with Grid standards used in Europe to enable interoperability, and aspects of multi-nationality within Globus and the Globus Alliance.

5.3.1. Objectives

The overarching objective is to help the European researchers by lessening their hassle with using DCIs and allowing them to harness greater computing power already available (such as DEISA or PRACE). More specifically, IGE

- Adapts Globus to better fulfil European requirements by coordinating European input from both users, developers, and infrastructure providers and thereby strongly impact the open source progress of the Globus Toolkit,
- Adds the European perspective to Globus by delivering tailored software development, operation, support, training, and documentation services to the European communities, and act as a central hub for Globus within Europe, and
- Broadens the adoption of Globus in Europe through coordinated dissemination, standardisation, and test infrastructure operation to foster seamless use of Grid infrastructures in other parts of the world.

5.3.2. Action plan

Over the past years, Europe has heavily invested in building e-Infrastructure for science. Especially in the area of DCIs, the Globus Toolkit is widely adopted as middleware solution and many scientific communities already contributed large efforts into using their application on top of the Globus middleware. Therefore, it is crucial to protect these investments during and after the transition to EGI.

To this end, IGE connects the European efforts on Globus usage, development, and operation by providing a single focal point: Through the European Globus Hub, major stakeholders will be able to learn about and use Globus, get involved in development and training, and contribute to the overarching goal of advancing Globus according to European needs. This includes various areas of concern:

Networking: IGE aggregates, consolidates, and provisions experiences in usage, development, and training from European Grid communities with the European Globus Hub. The visibility, presence, and adoption in Europe and maintaining close cooperation with the international community—including the Globus Alliance—is increased through coordinated efforts of the European Globus Liaison Office.

Services: IGE supports the definition and implementation of Grid infrastructures on the basis of Globus Toolkit. To this end, a comprehensive reference installation and test environment is provided which specifically caters the needs of European user, developer, and infrastructure provider communities.

Research: IGE delivers Globus components and tools with a particular European focus which fulfil the specific needs of the European e-Infrastructures' user communities. In this context, the provision of components that are eligible for inclusion into the UMD provided by EGI and interoperable with other middlewares through agreed standards (e.g. OGSA-BES and JSDL) is paramount. In addition, IGE collaborates with other Globus developers to foster reliability, usability, and stability of the Globus and cooperates with the Globus Alliance to add missing functionality, increase the manageability, and to introduce improvements into the core distribution.

User Integration: The substantial demand for Globus in Europe is shown by over thirty active supporters of IGE, ranging from industry, academia, e-Infrastructures, NGIs and international Grid projects. The seamless and progressive transition in services delivered through IGE will deliver a user-friendly, well-integrated Globus distribution and thus present a transparent and cost-effective way forward for current and emerging user communities in the European Research Area.

Internationality: Especially in the Americas and Asia-Pacific area, Globus is often the solution of choice for building Grid infrastructures. In order to support cooperation with international researchers, Globus is provided in a coherent way to European researchers, thus promoting close collaboration and interoperability with already established research infrastructures worldwide. These transatlantic relations in DCI research, development, and operation is further strengthened by the increase in international

collaboration. As a grassroots movement, a reliable link with the Globus development team is established within IGE by the incorporation of University of Chicago as a full partner.

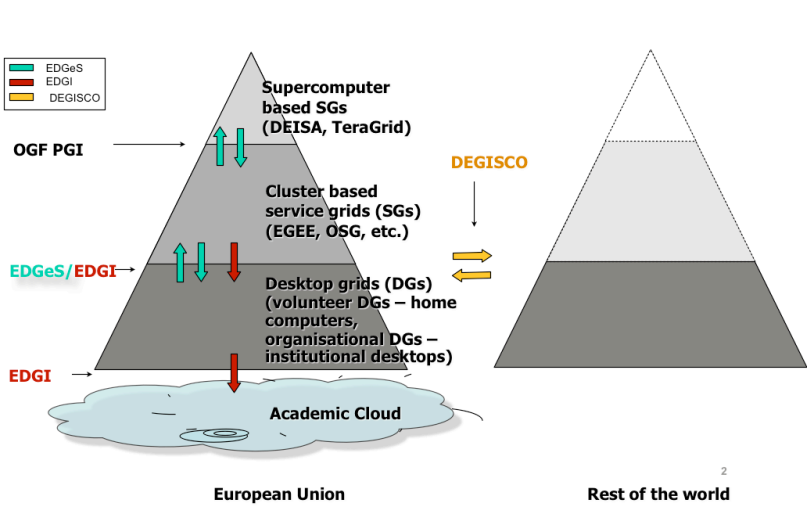
5.4. European Desktop Grid Initiative (EDGI)

EDGI (European Desktop Grid Initiative) is aimed at deploying Desktop Grid and Cloud Computing services for the European Grid Initiative (EGI) research user communities that require large-scale distributed computing resources for multi-national projects. In order to achieve this goal EDGI will develop middleware for extending Service Grids (SG) (ARC, gLite, UNICORE) with Desktop Grids (DG) (BOINC, XtremWeb, OurGrid) enhanced by Academic Clouds (Eucalyptus and OpenNebula). Software components of ARC, gLite, UNICORE, BOINC, XWHEP, Attic, 3GBridge, OpenNebula and Eucalyptus will be integrated into a SG → DG → Cloud platform for service provision and as a result EDGI will extend ARC, gLite and UNICORE grids with volunteer and institutional DG systems. In this way, the whole European e-science ecosystem will benefit from Desktop Grid extensions, since parameter sweep applications that run millions of sequential jobs can be directed from the expensive cluster and supercomputer resources to cheap desktop resources.

EDGI will create novel QoS support for the DG systems and will explore new service provision models in order to ensure harmonised DG→Cloud interfaces to ARC, gLite, UNICORE resources. The developed DG→Cloud bridge middleware has the goal to get instantly available additional resources on demand if the application has some QoS requirements that could not be satisfied by the available resources of the Desktop Grid system. New scheduling algorithms will be developed that will be able to take into consideration QoS requirements and will enable a more flexible allocation of task and resources in the Desktop Grid systems.

EDGI will further develop the support for data-intensive applications and not only in the context of gLite but also in the context of ARC- and UNICORE-based Grid systems. The ADICS P2P data management system and its bridge support developed in EDGeS at prototype level will be extended for ARC- and UNICORE-based Grid systems and will be deployed as production service in the EDGI project.

The figure shows the place of Desktop Grids in the well-known pyramid of computational resources for e-Science. At the top are the supercomputers, the large optimized systems located in supercomputer centres. A number of these machines are connected into a supercomputer Grid, pioneered by DEISA and continued by PRACE. For many applications clusters perform just as well, but they are less expensive and easier to manage. There are many more clusters than supercomputers. Clusters can be connected to cluster



Grids, managed typically by EGI/NGI's. At the lowest level are Desktop Grids. Especially when one looks at volunteer desktop Grids, the number of computers can be even large. Desktop grids are suited for a subset of cluster Grid applications. Clouds can fit in at many levels, but are placed outside the computing

pyramid, because of their specific function in the EDGI project. The figure also shows how the different levels of the pyramid are connected, or will be connected by the EDGI or DEGISCO project

EDGI does work closely together with the DEGISCO project. DEGISCO is a support project that supports extension of the European DCIs into countries outside the European Union, with a focus on Desktop Grids.

EDGI and DEGISCO did start the International Desktop Grid Federation to support Desktop Grid operators and developers for Desktop Grids. Integration of Desktop Grids into the European DCIs is an important goal.

EDGI will support the European Chapter of the International Desktop Grid federation with the aim of advancing and promoting Desktop Grid technology in Europe both by sharing and mutually leveraging experience and technological solutions acquired while independently operating Desktop Grids (such as Ibercivis, SZTAKI Desktop Grid, AlmereGrid, EDGeS@home, and many others). The federation will work at both technical and dissemination level. At technical level, it brings together Desktop Grid administrators who together will provide best practices and common solutions to common problems and share knowledge (instead of independently coming up with different and incompatible solutions to these problems). This forum is also used to disseminate the bridge middleware knowledge among European Desktop Grid system providers including companies. At the dissemination level, the federation is a key player in reaching European citizens to provide Desktop Grid resources that would not be possible for the individual European Desktop Grids alone. The International Desktop Grid federation will significantly contribute to the sustainability of the EDGI production infrastructure created in the project.

5.5. StratusLab

5.5.1. Summary

StratusLab is aimed at service provisioning, networking and research of cloud and virtualization technologies to simplify and optimize the use and operation of existing distributed computing infrastructures like the European Grid Infrastructure (EGI). The project is developing the StratusLab Toolkit, an open source cloud distribution. It incorporates cloud and virtualization innovation into existing grid infrastructures by integrating cloud technologies and services within grid sites. Further, it enriches existing computing infrastructures with “Infrastructure as a Service” (IaaS) cloud-like delivery paradigms.

5.5.2. Objectives

StratusLab brings several benefits to the e-Infrastructure ecosystem, in terms of simplification, added flexibility, increased maintainability, quality, energy efficiency and resilience of the sites. The new StratusLab Toolkit cloud distribution complements existing grid middleware services: the aim is for the cloud layer to be fully transparent to layers above. Existing grid middleware continues to provide the glue to federate the distributed resources and the services for high-level job and data management. StratusLab will help to improve the usability of distributed computing infrastructures, to attract scientific user communities, to appeal equally to industrial users, to keep European research infrastructures at the technological forefront, and to strengthen the know-how in virtualization and cloud computing of European industry.

5.5.3. Action plan

StratusLab will integrate, distribute and maintain a sustainable open-source cloud distribution to bring cloud to existing and new grid sites. The StratusLab toolkit will be composed of existing cutting-edge open source software and the innovative service and cloud management technologies developed in the project. It will also include the required additions to turn the software elements into a production grade

distribution to support production quality and operational systems, as will be demonstrated with the operation of production level grid sites in the project.

StratusLab is a two-phase project. In the **first phase**, the project will focus on **cloud computing for resource provisioning in grid sites**. This will entail development and integration of the initial StratusLab cloud platform, incorporating the components required for the virtualization of grid sites; and creation of virtual appliances for the scientific application domains in the project

In the **second phase** the emphasis will shift towards developing **new cloud-like delivery paradigms in grid sites**. This will build on the first phase, including new IaaS cloud interfaces and support for creation of new virtual appliances, which will be stored in a repository

Efforts to achieve both goals will start from the beginning of the project: the expectation is that the second goal will be achieved in the longer term.

Networking activities: The project's networking activities have been designed to foster collaboration over the complete spectrum of actors, from project participants, through our targeted user communities, to the ensemble of related European projects.

StratusLab will undertake extensive dissemination activities, targeting the user communities listed below, as well as the general public. Awareness of the project will be achieved through participation in relevant meetings, forums, workshops and conferences. StratusLab aims to publish in relevant journals and magazines. The project will also be active online through web presence. In-depth knowledge transfer will take place through demonstrations and training sessions.

Service activities: In order to certify the StratusLab toolkit, the project will deploy and maintain a small yet representative infrastructure. This 'pre-production' environment will provide the required platform for deploying incrementally the results of the cloud integration activity, but also provide a test-bed for joint research activities to deploy and test their research results.

The StratusLab infrastructure will also serve as an important platform for assessing the economic 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.)

Joint Research activities: In StratusLab the research activity consists of very specific and focused actions to achieve the main goal of the project that is to integrate a toolkit for offering cloud and grid services. The research activity will be targeted to extend current grid site management functionality, providing or enhancing tools and components to define and dynamically support service elasticity and SLA-powered scalability, optimize site provisioning, placement heuristics, virtual images management and resource sharing capabilities.

User communities: StratusLab benefits a wide variety of users: scientists, software scientists and engineers, community service administrators, system administrators and hardware technicians.

5.6. Virtual multidisciplinary EnviroNments USing Cloud Infrastructures (VENUS-C)

Goals: VENUS-C is aimed at developing and deploying a Cloud Computing service for research and industry communities in Europe by offering an industrial-quality service-oriented platform based on virtualisation technologies, with the aim of:

- Creating a platform that enables user applications to leverage cloud computing principles and benefits.
- Leveraging the state of the art to bring on board early adopters quickly, incrementally enable interoperability with existing Distributed Computing Infrastructures (DCIs) and push the state of the art where needed to satisfy on-boarding and interoperability.
- Creating a sustainable infrastructure that enables cloud computing paradigms for the user communities inside the project and new communities recruited through an Open Call.

Operation and Services: 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 underpinned by the Engineering data centre, Microsoft through the Windows Azure and its European data centres, and 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-premises environments or the Internet. From an Open Source perspective, the Eucalyptus and OpenNebula solutions will be evaluated, while the Emotive middleware for clouds will be offered by the Barcelona Supercomputing Centre, thus demonstrating interoperability and ultimately portability to the VENUS-C users.

Technical challenges addressed include virtualisation, service orientation and digital convergence, which are at the heart of the cloud model, as well as current open issues on interoperability with existing DCIs (e.g. Supercomputers), Data Management, Programming models, Application Security, Monitoring and Accounting, Networking and Network Security.

Action plan: In the first 12 months, the project will focus on the delivery of an end-to-end prototype which delivers immediate value to scientific partners: the first release will focus on dynamic job submission and workload dispatch into multiple underlying DCI and cloud providers. Subsequent milestones will enable integration with data management, security and programming models, working in synergy with our scientific users, primarily focusing on the functionality of directly usable application-level. Less-visible infrastructure work will start after the initial delivery of the core platform.

User scenarios: VENUS-C draws its strength from a joint co-operation bringing together industrial partners and scientific user communities through an innovative approach in the drive towards world-class research, and competitive edge for the European research community. The infrastructure will be initially tested across four thematic areas comprising seven applications: Biomedicine (integrating widely used tools for Bioinformatics, System Biology and Drug Discovery); Civil Protection and Emergencies (focusing on early fire detection), Civil Engineering (construction information management for environmental compliance), and data for science (Marine Biodiversity). To broaden the scope of the current user scenarios, VENUS-C will co-ordinate an Open Call, which will fund up to twenty new experiments in order to address the advanced needs of user communities, in some instances handling complex workflows and data-intensive scenarios. VENUS-C aims to empower these communities through the easy deployment of end-user services, in order to make e-Infrastructures more widely valuable across a spectrum of research fields without the complexity of existing grids and high up-front costs.

New Business Scenarios: An important goal of VENUS-C is to assess new business models as part of the drive to foster the shift away from the use of credit cards on a pay-per-use basis and placing more emphasis on a spirit of entrepreneurship through the involvement of pioneering European enterprises and

outreach to clusters and start-ups. The feasibility of different follow-on scenarios, such as Public-Private Partnerships (PPP), integration with pertinent on-going initiatives, or service provision through open tenders will be investigated and will bring on board the value-add of each VENUS-C partner, whether public or private, underpinned by co-operation and synergies on multiple levels.

Co-operation with external experts: Provisioning, deployment, sustainable growth and cost-effective investment at EU level will also be addressed by drawing on the advice and insights of a select group of experts in and outside Europe, recruited from the distributed computing and service-oriented technology arena. To this end, VENUS-C co-ordinates an External International Advisory Committee providing timely input on pertinent initiatives in and outside Europe, coupled with insights on technical and business-level developments and the broader, international landscape that will help position VENUS-C as an EC-funded initiative, in this landscape and help support potential integration, partnerships and synergies across the distributed computing arena.

Co-operation at EU and International level: VENUS-C partners have an extensive network of relations with other countries and initiatives with which strategic alliances will be established. VENUS-C objectives on the Open and generic APIs for scientific platform can only be fully reached if it succeeds in liaising with any actor in this context. U.S. initiatives like FutureGrids and Magellan, other R&D EU projects like VISION-CLOUD, Contrail, R&D projects on Experimental test-beds, like TEFIS, Bonfire, etc, projects on Scientific Data repositories like the D4Science Ecosystem, as well as initiatives on impact assessment of e-Infrastructure technologies like ERINA.

VENUS-C in the European Landscape: VENUS-C aims to broaden inter-disciplinary scientific collaboration in Europe and to address the following issues in the European landscape.