



Enhancing Grid Infrastructures with
Virtualization and Cloud Technologies

Project Quarterly Report

Quarterly Report QR6 (V1.0)
15 December 2011

Grant Agreement Number	INFSO-RI-261552
Project acronym	StratusLab
Project title	Enhancing Grid Infrastructures with Virtualization and Cloud Technologies
Funding Scheme	CP/CSA
Date of latest version of Annex I against which the assessment will be made	2010-05-31
Periodic report	1st
Period covered	2010-06-01 to 2011-05-31
Name, title and organisation of the scientific representative of the project's coordinator	Dr. Charles Loomis, Research Engineer, Centre National de la Recherche Scientifique (CNRS)
Tel	+33 (0)1 64 46 89 10
Fax	+33 (0)1 69 07 94 04
E-mail	loomis@lal.in2p3.fr
Project website address	http://stratuslab.eu



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 © 2011, 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
Marc-Eliau Bégin	SixSq	WP3, WP4
Christophe Blanchet	CNRS-IBCP	WP2, WP3
Kathryn Cassidy	TCD	WP3
Vangelis Floros	GRNET	WP3, WP5, WP6
Javier Fontan	UCM	WP4
Eduardo Huedo	UCM	WP6
Stuart Kenny	TCD	WP5
Ignacio M. Llorente	UCM	WP3
Charles Loomis	CNRS-LAL	WP2, WP3, WP5, Summary, Mgt.
Louise Merifield	SixSq	WP4
Ruben S. Montero	UCM	WP6
Henar Muñoz	TID	WP2, WP4, WP6
David O’Callaghan	TCD	WP3

Document History

Version	Date	Comment
0.1	1 Dec. 2011	Initial outline.
1.0	15 Dec. 2011	Final version.

Contents

List of Figures	7
List of Tables	8
1 Publishable Summary	9
1.1 Project Context and Objectives	9
1.2 Summary of Work Performed and Achievements	10
1.3 Final Results and Potential Impact and Use	12
1.3.1 Improved Interdisciplinary Scientific Collaboration	12
1.3.2 Impact on DCI Evolution.	13
1.3.3 Improved Usability of DCI Platforms	13
1.4 Contact Information	13
2 Project Objectives for the Period	15
2.1 Objectives	15
2.1.1 WP2: Interaction with Users and Related Communities	15
2.1.2 WP3: Dissemination	16
2.1.3 WP4: Integration, Distribution and Support of Open-Source Cloud Distribution	17
2.1.4 WP5: Operation of a Production Grid Site Running Stra- tusLab.	17
2.1.5 WP6: Innovative Cloud-like Management of Grid Ser- vices and Resources	18
2.2 Detailed Objectives by Quarter.	18
2.2.1 Quarter 5.	18
2.2.2 Quarter 6.	19

2.2.3	Quarter 7	19
2.3	Review Recommendations	20
3	Progress and Achievements	25
3.1	WP2: Interaction with Targeted Communities	28
3.1.1	Summary	28
3.1.2	Task 2.1: Interactions with Resource Providers and End-users	28
3.1.3	Task 2.2: Intensive Evaluation of StratusLab Products	30
3.1.4	Issues and Corrective Actions	31
3.2	WP3: Dissemination.	32
3.2.1	Summary	32
3.2.2	Task 3.1: Dissemination	32
3.2.3	Task 3.2: Collaboration with Standards Bodies and Related Projects	34
3.2.4	Task 3.3: Development of Exploitation and Sustainability Plan	36
3.2.5	Issues and Corrective Actions	37
3.3	WP4: Software Integration and Distribution	38
3.3.1	Summary	38
3.3.2	Task 4.1: Definition of Reference Architecture	38
3.3.3	Task 4.2: Integration of Open-source Distribution	38
3.3.4	Task 4.3: Contextualization of Grid Services	39
3.3.5	Task 4.4: Technical Support	39
3.3.6	Issues and Corrective Actions	39
3.4	WP5: Infrastructure Operation	41
3.4.1	Summary	41
3.4.2	Task 5.1: Deployment and Operation of Virtualized Grid Sites	41
3.4.3	Task 5.2: Testing of the StratusLab Toolkit	42
3.4.4	Task 5.3: Virtual Appliances Creation and Maintenance	43
3.4.5	Issues and Corrective Actions	43

3.5	WP6: Innovative Cloud-like Management of Grid Services and Resources	45
3.5.1	Summary	45
3.5.2	T6.1: Dynamic Provision of Grid Services	45
3.5.3	T6.2: Scalable and Elastic Management of Grid Site Infrastructure	45
3.5.4	T6.3: Cloud-like Interfaces Specific for the Scientific Community	46
3.5.5	Issues and Corrective Actions	46
4	Project Management	47
4.1	Consortium	47
4.2	Management Tasks	47
4.3	Issues	47
4.4	Planning	48
4.4.1	Objectives for Next Quarter	48
4.4.2	Roadmap	48
5	Deliverables and Milestones	57
6	Use of Resources	61

List of Figures

2.1	Primary and supporting objectives	16
3.1	Visits for Q6.	33

List of Tables

1.1	StratusLab Information and Support	13
1.2	StratusLab Partners	14
3.1	Talks	35
4.1	Meetings (Q1).	50
4.2	Meetings (Q2).	51
4.3	Meetings (Q3).	52
4.4	Meetings (Q4).	53
4.5	Meetings (Q5).	54
4.6	Meetings (Q6).	55
4.7	Metrics	56
5.1	Deliverables (Year 1).	58
5.2	Deliverables (Year 2).	59
5.3	Milestones	60

1 Publishable Summary

1.1 Project Context and Objectives

The StratusLab project is aimed at service provisioning, networking, and research of technologies that will bridge cloud and grid infrastructures to simplify and optimize the use and operation of existing distributed computing infrastructures (e.g. European Grid Infrastructure) and to provide a more flexible, dynamic computing environment for scientists.

The European production grid infrastructure has had many notable successes. It has allowed scientists from all over Europe and indeed from all over the world to federate their computing resources to advance their scientific aims. More importantly, the infrastructure allows them to federate their data and expertise to accomplish more than they would be able to do singlehandedly. Common APIs and service interfaces make it possible to take advantage of these distributed resources without having to modify applications for each site.

Despite its success, the grid also has its limitations. The uniformity of service interfaces unfortunately does not extend to the underlying computing resources, where users are exposed to significant heterogeneities in the computing environment, complicating applications and increasing failure rates. Passive calculations are handled well by the grid, but many applications require active services to coordinate the distributed analyses. Either scientists must provide their own resources for such services or negotiate with a particular site to provide them. This reduces the speed at which new calculations can be done.

Virtualization technologies provide a mechanism for offering customized, uniform environments for users with negligible performance degradation. Using grid technologies combined with virtualization allows the grid to provide users with a homogeneous computing environment, simplifying applications and reducing failures. Emerging cloud technologies allow users to dynamically allocate computing resources (often in less than a minute) and to specify the characteristics for the allocated resources. The fusion of cloud and grid technologies provides a more dynamic and flexible computing environment for grid application developers.

Cloud and virtualization technologies also offer other benefits to administrators of resource centers, such as the migration of live services for load balancing or the deployment of redundant servers. Reduced costs for managing resources immediately benefit users by freeing money for additional computing resources or

by having better user support from administrators.

A combined computing infrastructure that uses grid technology's strengths for federating resources, virtualization's strengths in providing custom, uniform environments, and the cloud's strengths in dynamic resource allocation, maximizes the utility of European distributed computing resources to scientists.

The StratusLab project creates an complete, coherent, open-source private cloud distribution to allow administrators of grid resources centers to take advantage of virtualization and cloud technologies. It provides new ways of using existing distributed computing resources to make the infrastructure more adaptable and more useful for scientists.

1.2 Summary of Work Performed and Achievements

In the sixth quarter, the project worked to enhance the existing production release and to update the architecture for the second major release due at the end of the project. Unfortunately, a major cooling failure at LAL seriously perturbed the build and test infrastructure and consequently also the ability of the project to make incremental releases of the StratusLab cloud distribution. Nonetheless, one incremental release was made that was used to support a number of tutorials. Progress was made on all of the quarter's objectives, although many were not fully completed because of the above problem. The points below summarize the progress made.

Incremental Releases of the StratusLab Cloud Distribution Version 1.1 of the StratusLab cloud distribution was made publicly available. This version was quickly installed on the project's reference cloud infrastructure and on the LAL's test infrastructure. Two additional planned incremental releases unfortunately did not take place because of the perturbations to the build and test infrastructure described above.

Implementation of Use Cases In the previous quarter, seven use cases have been defined on which to focus the project's porting and support efforts. Work on three of these—a bioinformatics application, a commercial 3-tier prototype service, and a software development PaaS—have significantly advanced, although not yet to the stage where they have concrete results and can be publicized. All of these are expected to provide dissemination opportunities in the next quarter.

Support for OpenSuSE Support for a second operating system by the StratusLab cloud distribution is desired for a couple of reasons. First, it demonstrates portability of the code and services. Second, support of OpenSuSE will allow us to test GPFS, which should have much better performance than NFS, as a shared file system for distributing images on the cloud infrastructure. Although modifications to the build system have been made to facilitate this, support for OpenSuSE has not yet been demonstrated.

Tutorials and Training Materials A major effort and success for this quarter has been a series of well-attended tutorials. These took place as part of the EGI Tech-

nical Forum in Lyon, France; as part of the ACGRID 3 school in Hanoi, Vietnam; and as a standalone event in Orsay, France. Overall, around 90 people participated in these events. The materials developed for these tutorials will continue to be updated for future releases of the distribution.

Restructured Communication Channels The communication channels between the project and the targeted communities have been somewhat restructured. In October an additional “user announcements” (announcement only) mailing list was created, specifically for users of the StratusLab reference infrastructure to inform them of upcoming service outages, upgrades and other important news. A new public user forum was created and went live on 16 November, and has been running on a trial basis so far. The aim of the forum is to provide a contact point for StratusLab users and to build a community support mechanism. The group already has 12 members, and has seen seven messages posted so far in two topics. It will thus be officially launched in December 2011.

Updated Architecture The quarter began with a redefinition of the global architecture of the StratusLab cloud distribution. The major changes were the additions of monitoring services, accounting services, and an “inter-cloud connector”. The designs of components for advanced services and functionalities have also been updated, providing a clear roadmap for the second year developments.

Improved Image Management A major integration effort centered on the persistent storage service. The image management code was significantly refactored to take advantage of this service and to allow efficient caching of virtual machine images. With these changes, the primary delay for starting images is now the latencies associated with the scheduling procedure. The “create image” feature has also been refactored to take advantage of this service, drastically reducing the time necessary to create a new image.

Streamlined Release Procedures In order to better streamline the release and certification of new software releases, new procedures have been put in place, supported by changes in our continuous integration system. This should allow future releases to be done “with the click of a button” with nearly all of the test and certification tasks done automatically.

Reference Configuration and Performance Benchmarks Although an important topic, little time was spent on this in Q6 because it was more important to get the build and test system working correctly and refactored to support more fluid release procedures.

Implementation of Advanced Networking Services Work was done to provide advanced networking services such as the dynamic provisioning of VLANs and configuration of firewalls. These features are part of the OpenNebula 3.0 release and is in the process of being integrated into the mainline StratusLab distribution. Integration problems on both the StratusLab and OpenNebula sides are being resolved and these new features will appear in the StratusLab distribution once OpenNebula 3.0 has been integrated.

Sustainability As we reach the end of the StratusLab project, sustainability is an important question for the users of our software. It is a question that comes up at nearly every StratusLab presentation. The face-to-face meeting at the end of the quarter allowed a detailed discussion of the sustainability strategy that will serve as a good basis for the final plans that will be developed over the last six months of the project.

Despite the problems with the build and test infrastructure's availability, significant progress was made on all of the objectives for this quarter. Those that have not been completed will be moved forward into the next quarter. More emphasis will be placed on supporting a second operating system as this has been pushed forward two quarters now. The next quarter will see the beta of the 2.0 release with prototypes of all of the expected services.

1.3 Final Results and Potential Impact and Use

Most scientific and engineering research requires significant computing resources. Distributed computing infrastructures have brought unprecedented computational power to a wide range of scientific domains. Although, these architectures and the related software tools have been considerably improved over the years, they exhibit several difficulties, mainly due to limitations of physical platforms, which discourage adoption of grid technologies. StratusLab has the potential to profoundly change existing grid infrastructures.

1.3.1 Improved Interdisciplinary Scientific Collaboration

Cloud technologies are expected to have significant impact, both immediate and long-term, in the way scientific research is carried out. Grid infrastructures have provided a remarkable advantage over the past years offering access to vast amount of computing power and storage space, and most importantly by offering a sustainable platform for scientific collaboration enabling the sharing of computing resources and scientific data. Cloud computing is expected to take this one step further by facilitating the easy deployment of customized grid infrastructures. These infrastructures are expected to have further positive impact on the way interdisciplinary scientific research is taking place.

StratusLab focuses on the provision of scientific infrastructures over cloud computing, investigating in particular the provision of customized Virtual Machine images. This customization will be done on the user side, which means that the user can have more immediate influence on the infrastructure itself. In this way the infrastructure will adapt to the user requirements and not vice-versa. By easing the management of grid sites and the configuration of hosting services we expect to attract a broader number of scientific communities and further facilitate their collaboration.

Table 1.1: *StratusLab Information and Support*

Website	http://stratuslab.eu/
RSS Feed	feed://stratuslab.eu/feed.php?ns=news&linkto=page
Twitter	@StratusLab
YouTube	http://www.youtube.com/user/StratusLab
Support	support@stratuslab.eu

1.3.2 Impact on DCI Evolution

Currently, there is a big shift in all e-Infrastructure projects, and related efforts in Europe, to expand their activities in order to include cloud computing technologies. StratusLab will play a key role in this landscape by providing a focused environment for development, deployment and experimentation of cloud computing services.

The projects proposal reflects an evolutionary path from the existing large-scale monolithic grid e-Infrastructures to novel, beyond the state-of-the-art, cloud-based, grid-enabled ones. Through its expected collaborations with other projects, StratusLab will disseminate its findings and drive direct impact on the way e-Infrastructure provision is currently done.

1.3.3 Improved Usability of DCI Platforms

Virtualization is the cornerstone of cloud computing and a key for achieving optimal usability of DCI platforms. Moreover, virtualized environments have the ability to adapt to different hardware platforms enabling a quick transition from one environment to another.

StratusLab operates such a virtualized platform on a variety of hardware environments. By offering customized machine images, users will be able to set-up an environment that better suits their application requirements. This will dramatically improve the current situation where current infrastructures are forced to offer a common configuration—a common denominator—that tries to do its best to satisfy many users with different runtime requirements. Another aspect where StratusLab will contribute is on power consumption efficiency (Green Computing) and the increase reliability by incorporating failover mechanisms using virtual machine snapshots and migration.

1.4 Contact Information

More information about the StratusLab project can be obtained from the sources listed in Table 1.1. Individual partners can also be contacted to obtain more specific information about their contributions to the project. Table 1.2 contains the list of StratusLab partners and relevant contacts.

Table 1.2: StratusLab Partners

CNRS	Centre National de la Recherche Scientifique	Charles LOOMIS loomis@lal.in2p3.fr
UCM	Universidad Complutense de Madrid	Ignacio LLORENTE llorente@dacya.ucm.es
GRNET	Greek Research and Technology Network S.A.	Evangelos FLOROS efloros@gnet.gr
SIXSQ	SixSq Sàrl	Marc-Elian BEGIN meb@sixsq.com
TID	Telefónica Investigación y Desarrollo SA	Henar MUÑOZ henar@tid.es
TCD	The Provost Fellows and Scholars of the College of the Holy and Undivided Trinity of Queen Elizabeth Near Dublin	David O'Callaghan david.ocallaghan@cs.tcd.ie

2 Project Objectives for the Period

2.1 Objectives

The primary objective of the project is to provide a software distribution that brings together cloud and grid technologies to benefit both grid resource center administrators and scientists. In order to achieve this main objective, we have defined a set of interrelated objectives to be addressed in the project. The objectives are organized, for clarity of exposition, into three groups of objectives, corresponding to networking, service and research activities (see Figure 2.1):

- The first group represents coordination and networking with users and other stakeholders in the grid and cloud ecosystems. The project will work directly with scientists using the grid to ensure that the distribution satisfies real needs; and will collaborate with related projects and contribute to standards bodies.
- The second group represents infrastructure related services to the scientific community. The project will integrate and maintain a software distribution to bring cloud to existing and new grid sites and will ensure the production quality of the distribution by running two production sites with the distribution.
- The last group represents innovation and exploration of new cloud and virtualization technologies to enhance grid infrastructures. The project will develop innovative technology for cloud-like management of grid services and resources that will be incorporated into the software distribution.

These objectives are presented by work package below. Similarly, the work program is built around these objectives. There is a one-one correspondence between objectives and activities, so facilitating an easy cross-reference between objectives and activities throughout this document, and their verification during the project execution. The activity on project coordination has not been included here.

2.1.1 WP2: Interaction with Users and Related Communities

StratusLab targets two distinct communities: resource providers and end-users. The StratusLab software will simplify grid site administration and improve the reliability of the site. Later releases in the second phase of the project will provide

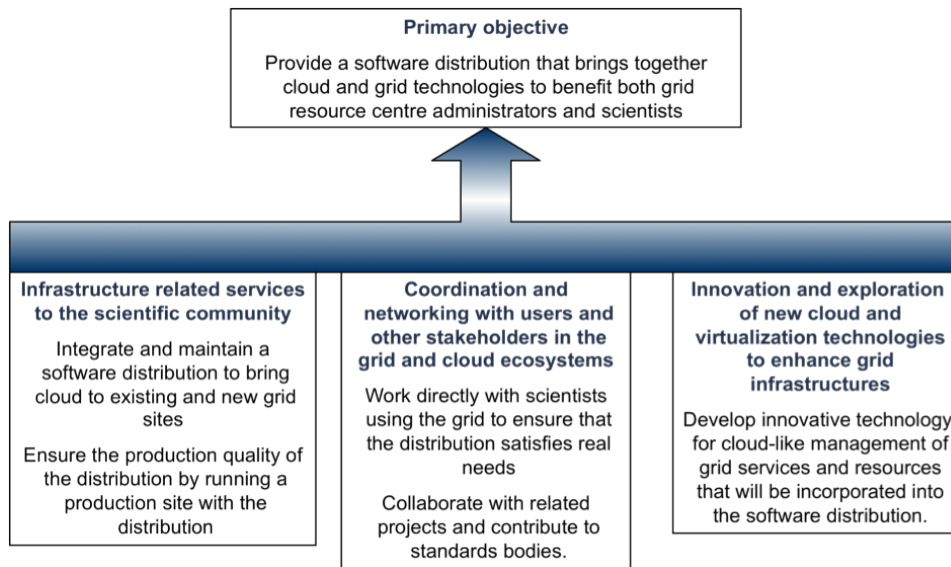


Figure 2.1: Primary and supporting objectives

direct cloud APIs that will be attractive for scientists porting applications to the grid. The communications between these communities and the project must be managed to ensure the project fully addresses their needs and any problems that arise. One community will work directly with the project to evaluate early releases of the software. Results of the project must be disseminated as widely as possible to those two communities as well as the general public. Scope of the objective.

- Manage communication with resource providers regarding their needs concerning virtualization and cloud technologies and their feedback on StratusLab software.
- Manage communication with end-users regarding their use of resources running StratusLab software and their needs for direct access to virtualization and cloud features.
- Training sessions will be organized to encourage dissemination of technical information and adoption of the StratusLab software.
- Evaluate early versions of StratusLab software from a users perspective with respect to utility and stability.

2.1.2 WP3: Dissemination

A large number of projects, companies, and standards bodies currently focus on cloud and virtualization technologies because of their promise and growing adoption. StratusLab must actively engage with those entities to ensure that the projects

results are well represented, that we are aware of others advances, and that we drive standardization in a direction consistent with our vision. Scope of the objective.

- 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.

2.1.3 WP4: Integration, Distribution and Support of Open-Source Cloud Distribution

StratusLab will integrate and support an open-source cloud distribution enabling grid site virtualization and dynamic scaling to remote clouds. This distribution will address the specific requirements of the grid resource providers and enable the deployment of science clouds, as well as addressing infrastructure cloud-like access demands from user communities, including industrial users. Scope of the objective. StratusLab will address the following topics:

- Selection of software components, from best of breed in open source software, to compose a robust and industry grade open source StratusLab toolkit. This distribution will integrate with typical administration tools in grid infrastructures and fabric management. This process will be driven by real needs and constraints in production infrastructures.
- Integration and management of open-source distribution, definition and maintenance of reference configurations and sustainability in the context of EGI and its official middleware distribution. The StratusLab toolkit will integrate the innovation developed in the research activity.
- Technical support for installation and configuration of the distribution, following industrial practices in term of quality, maintainability, testability and usability
- Definition of a process for automatic configuration of the virtual appliances

2.1.4 WP5: Operation of a Production Grid Site Running StratusLab

StratusLab will engage two resource centers that will be responsible for the deployment of middleware and tools developed in the project. One the main tasks of these resource centers will be the operation of two production grid sites running StratusLab toolkit. The sites should be able to pass the certification procedures imposed by EGI. The activity will demonstrate the security, performance, reliability and scalability of the distribution, and will provide support for the creation of

the virtual appliances for different user communities. The activity will also investigate the feasibility of offering a repository of reference images for cloud users, with demonstrated interoperability among the supported cloud infrastructures (including the private cloud deployed in the re-source centers, as well as a selected number of public clouds). Scope of the objective. StratusLab will address the following topics:

- Deployment and operation of virtualized grid sites
- Testbed for the StratusLab toolkit
- Support for the creation of virtual appliances for different user communities.

2.1.5 WP6: Innovative Cloud-like Management of Grid Services and Resources

StratusLab will conduct research on grid service automatic deployment and dynamic provision, including automatic elasticity mechanisms for scaling up and down to meet performance goals (typically defined by SLAs). StratusLab will also conduct research on novel infrastructure cloud-like resource provisioning paradigms, and dynamic and scalable management of virtualized infrastructures for grid services. The research will be performed to address technology gaps defined by the service activities according to user requirements collected by the networking activities. Scope of the objective. StratusLab will address the following topics:

- Framework for grid service elasticity and dynamic provision of grid services
- Grid specific virtual machine management techniques
- Infrastructure cloud interfaces for grid sites and its integration with existing Grid services

2.2 Detailed Objectives by Quarter

2.2.1 Quarter 5

- Solidify the v1.0 StratusLab cloud distribution through increased testing and hardening of existing services.
- Support for a second operating system to ensure the portability of the distribution.
- Survey of the users and system administrators to see if the requirements have evolved from those already collected in Y1.
- Update and expand the target reference architecture for the distribution.
- Continued dissemination of project results.

- Continued operation of reference infrastructure and support to users and system administrators.
- Expansion of the number of users and sites using StratusLab.

2.2.2 Quarter 6

- Release incremental production versions of the StratusLab cloud distribution.
- Support for a second operating system to ensure the portability of the distribution.
- Implementation of an identified use case.
- Provision of tutorials for finding and training new users.
- Update and expand the target reference architecture for the distribution.
- Dissemination of project results with emphasis on the general public.
- Definition of reference cloud configurations and implementation of performance benchmarks.
- Improved integration of image management and caching.
- Implementation of more advanced networking services (e.g. dynamic firewalls).

2.2.3 Quarter 7

- Solidify the v1.0 StratusLab cloud distribution through increased testing and hardening of existing services.
- Support for a second operating system to ensure the portability of the distribution.
- Produce initial beta release of v2.0 of the StratusLab cloud distribution.
- Dissemination of project results of a completed use case.
- Expanding the sites using StratusLab and completion of MS4.
- Continued operation of reference infrastructure and support to users and system administrators.
- Operate a pre-production elastic Grid site elasticity, verify the applicability of the technology and move results to a production site.
- Evaluate GPFS as a backend storage solution. Prioritize and evaluate additional file systems.

- Develop additional use cases similar to MapReduce (e.g. Matlab application showcase)
- Integrate caching sub-system within the production cloud service
- Integrate NFS persistent storage service in the reference cloud service
- Integration of OpenNebula 3.0 with initial tests of new networking functionality

2.3 Review Recommendations

The responses given below to the reviewer recommendations have been updated to reflect the situation at the end of Q6.

1. *Due to unscheduled availability of a physical production infrastructure experiments, testing and debugging have been affected. This was caused by problematic financial issues at the start of the project. The project should make provisions that such events are better mitigated in the future.*

(See next response.)

2. *Several WP have shown a “delayed” start due to various hiring issues at the beginning of the project. A better process should be put in place.*

These delays were largely related to the difficulty of hiring new personnel in the summer months and to lengthy administrative procedures. Most of the partners compensated to some extent with effort from permanent staff. Currently, all of the partners are fully staffed and no further perturbations are expected.

3. *The dissemination work is focussed towards the more technically oriented (system administrators) communities rather than the scientific user groups who could benefit from StratusLab. This is acceptable for year 1 but the focus should shift to the scientific users in year 2.*

The project agrees that this is a reasonable shift in focus for the dissemination work in year 2. The project proposed a demo for the EU Innovation Convention, which unfortunately was not accepted. However, work continues on a video demo for the general public. More effort will be put into creating general dissemination materials and demos, both of which will require effort from all activities within the project.

4. *The dissemination targets as mentioned by the related KPI metrics were not very ambitious. The project should establish more ambitious KPI metrics with respect to dissemination targets.*

The project management along with the activity managers are in the process of redefining the metrics and the targets for the second year of the project. More ambitious dissemination targets will be proposed.

5. *A clear and simple demonstration showing the benefits rather than the technology itself should be produced. This benefits should be illustrated focusing on one or two clear use cases. The solutions should clearly demonstrate what the real benefits are (for Scientific Users / System administrators) of the StratusLab toolkit.*

A major part of the D2.3 deliverable was to identify possible use cases and would serve as a good basis for demonstrating the benefits of cloud technologies and as good topics for focused dissemination efforts. Work on these use cases are advancing and we expect that in Q7 the first dissemination activities related to them will appear.

6. *Provide a clear map of the components of the toolkit. Which components are re-used, which are newly developed and which are adapted from existing components?*

Deliverable D4.4 that provides the reference architecture for StratusLab 2.0 was structured to provide a separate description for each service. That description includes details on where the component was developed, external dependencies, and interactions with other components. This should provide a clearer overview of the development activities of the project. In addition, more effort has been made in the deliverables to describe clearly the project's work with respect to individual components.

7. *The periodic report is in draft status. Please submit a final version.*

The final version of the periodic report for year 1 has been submitted. The delay was due the difficulty in getting the necessary financial information from the partner's institutes during the summer months. Through the PMB, the partners have been advised that this information must be provided in a more timely manner for the following periods. Forcing partners to provide the information quarterly will also help in preparing the final information at the end of period 2.

8. *Knowing year 1 budget under spending, a new forecast of the planned spending for year 2 including a recovery plan for the current under spending should be presented to the EC as soon as possible. A reasonable deadline is one month after the receipt of this review report. It should take into account the reasons of year 1 under spending and should introduce measures that allow the effective implementation of the recommendations for year 2.*

The project is in the process of revising the effort and spending plans for year 2. The work of revising the budgets is ongoing. It is hoped that the projected budgets along with figures for Q5 and Q6 will be available in January 2012.

9. *The Data Management layer should be improved. In particular, StratusLab should be able to use existing and robust parallel file-systems which have better scalability than NFS such as Panasas or GPFS.*

This recommendation refers to the use of shared file systems to make machine images available to the various computational hosts of the cloud infrastructure. Alternate technologies such as iSCSI that do not rely on a shared file system have been investigated and used successfully. Nonetheless, shared file systems provide a convenient alternative. Consequently, tests of GPFS will be done at LAL to gauge its performance when used as a machine image cache. Furthermore, additional distributed and parallel file systems will be evaluated (e.g. PVFS) but also file systems that have already been tested (e.g. GlusterFS) will be re-visited for further investigation and scrutiny.

10. *Testing and benchmarking in WP5 should be more detailed including performance aspects.*

A testing plan will be developed that defines various reference cloud configurations and measures of performance. The existing application benchmarks will be used as a starting point and then expanded as necessary. To the largest extent possible, these tests will be automated through the existing Hudson continuous integration infrastructure.

11. *More emphasis should be put on the Cloud API rather than the GRID.*

By moving work on the Cloud API from the second year of the project to the first year already shows that the project believes that cloud-like access to resources is important. The project has decided to support OCCI 1.1 and CDMI as formal standards in addition to the existing proprietary interfaces. The highest priority, however, will be a jClouds plugin as this has been the most requested feature related to service interfaces.

12. *Although security issues are taken very seriously, privacy issues should be taken seriously as well. For instance, in case when a “closed” Grid infrastructure is complemented/bridged with an external public Cloud infrastructure when it is processing for instance medical sensitive information that can be relayed back to an individual person.*

Our feeling is that StratusLab should remain a “neutral” carrier that allows users and administrators to implement their own mechanisms for ensuring privacy and confidentiality. Nonetheless, we see areas in which the services can improve to provide, for instance, better logging and auditing information that can complement user-level strategies for privacy and confidentiality. Those areas will be targeted for improvement.

13. *The security incident as reported in Q3 should be analysed thoroughly and measurements should be taken to prevent this to happen again on the live production system.*

These security incidents were taken seriously and analysed thoroughly. As a result of these incidents two additional features were added to the Stra-

tusLab distribution: 1) enhanced logging and 2) image policy enforcement. The enhanced logging makes it easier to trace the characteristics, ownership, and history of a particular machine image aiding forensic analysis. The distribution now also contains a policy enforcement engine that allows system administrators to define what images are authorized to run on a given infrastructure. This policy enforcement mechanism is closely tied to the information provided in the Marketplace. Cloud administrators will still need to monitor the cloud for suspicious activity and take corrective actions as necessary. From the operations point of view (WP5) we will remain alerted for potential future security incidents and will be ready to respond quickly in coordination with other infrastructure operations teams (NGI NOC, EGI CSIRT team etc.).

14. *The project should clearly define a small number of use cases and focus the project towards delivering real value to these user communities, targeting system administrators as well as scientific users. As also indicated later in section 5 “Use and dissemination of foreground” the project should adopt an attitude of someone who tries to offer services to the market and must convince someone to spend funds for the services. This way the project can maximize impact and the work can contribute to sustainability.*

The deliverable D2.3 defines a set of seven initial use cases. These will be scheduled for implementation over the second year of the project. A reorientation of the dissemination activities to provide better “marketing” is being planned by WP3 and will be supported by WP2 and the other activities.

15. *A person should be appointed in the Project Management Board that can help the project to move from technology towards real end user solutions and benefits. This person should also help the project to establish solid relationships with stakeholders outside the traditional high energy physics (HEP) community.*

The project agrees that having a “user champion” within the project is a good idea. However, the PMB is not the proper body as it meets infrequently and isn’t in direct contact with the daily technical advances of the project. The project is still evaluating what is the best option for having stronger relationships with people outside of HEP.

16. *While it is clear what StratusLab could offer to the scientific community, the impact of StratusLab would be much bigger when the toolkit could also be used for users in the commercial world. Through collaboration with the Spanish TID private Cloud project, a large number of potential additional use cases (Telefonica’s customers) could be developed and should be taken seriously, including dissemination towards other DCI projects.*

Better marketing of the StratusLab distribution will include both commercial and scientific communities as described above. TID will be more involved

in presenting the project's work to enterprises with several possible venues being considered.

17. *Dissemination of the StratusLab Toolkit should become more marketing driven and should target both ICT press to reach potential industry users as well as to potential VOs beyond the current bioinformatics users. A demo centred on a use case could help.*

As stated in other responses, we agree that the dissemination should become more marketing driven. Deliverable D2.3 has identified initial use cases for implementation. Specific dissemination activities and possible demonstrations will be considered as each use case is successfully implemented.

18. *Future reports should contain less "fat" and should be more crispy and to the point. Executive summaries should be self-contained and should answer: (a) why should I read the deliverable, (b) the benefits for my company/organisation, (c) aspects addressed in this deliverable, (d) summary of recommendations/findings. The report should clearly describe if components are newly developed, improved or reused and integrated by StratusLab.*

The project will make year 2 deliverables more concise with better executive summaries following the guidelines given above.

19. *Ensure the project periodic and final reports are available at least two weeks prior to the review meeting.*

The review for year 2 will be scheduled to ensure that the final periodic report is available at least two weeks prior to the review. All of the deliverables for year 1 were available at least two weeks prior to the review and we intend that to also be the case for year 2.

20. *On top of the internal reports, consider publishing for the general press, like a newspaper and possibly to organise a public demonstration for the less technical audience, e.g. at the European Parliament.*

In August, a general paper was prepared that describes the goals of the project and the StratusLab distribution. Although still technical, this gives a good overview of the project. This and other project documents will be further generalized to appeal to a wider audience. A first step in this direction was a proposal for a booth at the EU Innovation Convention in Brussels, which was unfortunately not accepted. Nonetheless, work on general materials and demos for non-technical audiences continues.

3 Progress and Achievements

In the sixth quarter, the project worked to enhance the existing production release and to update the architecture for the second major release due at the end of the project. Unfortunately, a major cooling failure at LAL seriously perturbed the build and test infrastructure and consequently also the ability of the project to make incremental releases of the StratusLab cloud distribution. Nonetheless, one incremental release was made that was used to support a number of tutorials. Progress was made on all of the quarter's objectives, although many were not fully completed because of the above problem. The points below summarize the progress made.

Incremental Releases of the StratusLab Cloud Distribution Version 1.1 of the StratusLab cloud distribution was made publicly available. This version was quickly installed on the project's reference cloud infrastructure and on the LAL's test infrastructure. Two additional planned incremental releases unfortunately did not take place because of the perturbations to the build and test infrastructure described above.

Implementation of Use Cases In the previous quarter, seven use cases have been defined on which to focus the project's porting and support efforts. Work on three of these—a bioinformatics application, a commercial 3-tier prototype service, and a software development PaaS—have significantly advanced, although not yet to the stage where they have concrete results and can be publicized. All of these are expected to provide dissemination opportunities in the next quarter.

Support for OpenSuSE Support for a second operating system by the StratusLab cloud distribution is desired for a couple of reasons. First, it demonstrates portability of the code and services. Second, support of OpenSuSE will allow us to test GPFS, which should have much better performance than NFS, as a shared file system for distributing images on the cloud infrastructure. Although modifications to the build system have been made to facilitate this, support for OpenSuSE has not yet been demonstrated.

Tutorials and Training Materials A major effort and success for this quarter has been a series of well-attended tutorials. These took place as part of the EGI Technical Forum in Lyon, France; as part of the ACGRID 3 school in Hanoi, Vietnam; and as a standalone event in Orsay, France. Overall, around 90 people participated in these events. The materials developed for these tutorials will continue to be

updated for future releases of the distribution.

Restructured Communication Channels The communication channels between the project and the targeted communities have been somewhat restructured. In October an additional “user announcements” (announcement only) mailing list was created, specifically for users of the StratusLab reference infrastructure to inform them of upcoming service outages, upgrades and other important news. A new public user forum was created and went live on 16 November, and has been running on a trial basis so far. The aim of the forum is to provide a contact point for StratusLab users and to build a community support mechanism. The group already has 12 members, and has seen seven messages posted so far in two topics. It will thus be officially launched in December 2011.

Updated Architecture The quarter began with a redefinition of the global architecture of the StratusLab cloud distribution. The major changes were the additions of monitoring services, accounting services, and an “inter-cloud connector”. The designs of components for advanced services and functionalities have also been updated, providing a clear roadmap for the second year developments.

Improved Image Management A major integration effort centered on the persistent storage service. The image management code was significantly refactored to take advantage of this service and to allow efficient caching of virtual machine images. With these changes, the primary delay for starting images is now the latencies associated with the scheduling procedure. The “create image” feature has also been refactored to take advantage of this service, drastically reducing the time necessary to create a new image.

Streamlined Release Procedures In order to better streamline the release and certification of new software releases, new procedures have been put in place, supported by changes in our continuous integration system. This should allow future releases to be done “with the click of a button” with nearly all of the test and certification tasks done automatically.

Reference Configuration and Performance Benchmarks Although an important topic, little time was spent on this in Q6 because it was more important to get the build and test system working correctly and refactored to support more fluid release procedures.

Implementation of Advanced Networking Services Work was done to provide advanced networking services such as the dynamic provisioning of VLANs and configuration of firewalls. These features are part of the OpenNebula 3.0 release and is in the process of being integrated into the mainline StratusLab distribution. Integration problems on both the StratusLab and OpenNebula sides are being resolved and these new features will appear in the StratusLab distribution once OpenNebula 3.0 has been integrated.

Sustainability As we reach the end of the StratusLab project, sustainability is an important question for the users of our software. It is a question that comes up at

nearly every StratusLab presentation. The face-to-face meeting at the end of the quarter allowed a detailed discussion of the sustainability strategy that will serve as a good basis for the final plans that will be developed over the last six months of the project.

Despite the problems with the build and test infrastructure's availability, significant progress was made on all of the objectives for this quarter. Those that have not been completed will be moved forward into the next quarter. More emphasis will be placed on supporting a second operating system as this has been pushed forward two quarters now. The next quarter will see the beta of the 2.0 release with prototypes of all of the expected services.

3.1 WP2: Interaction with Targeted Communities

This activity manages the relationships with the communities targeted by the project, notably scientists from diverse fields and system administrators interested in deploying a cloud infrastructure. Through those interactions, it provides requirements and feedback to the other activities within the project while at the same time evaluating the StratusLab distribution from the points-of-view of users and system administrators. The activity also helps provide support to the targeted communities.

3.1.1 Summary

The major effort for this period has been giving tutorials for the StratusLab cloud distribution. Three large tutorials were given (Lyon, France; Hanoi, Vietnam, and Orsay, France) that included around 90 students. In addition, a cloud module featuring StratusLab has been added to an software engineering masters programs in Marrakech, Morocco with 18 students. The developers of the Narval data acquisition framework ran a tutorial of their software using the LAL StratusLab cloud infrastructure, allowing them to provide easily temporary resources for the students. This activity by an outside institute is an important validation of the StratusLab software. In addition to these training activities, the work package has worked on specific use cases (commercial and academic) to demonstrate the capabilities of the StratusLab cloud distribution.

3.1.2 Task 2.1: Interactions with Resource Providers and End-users

Tutorials People from this activity have worked intensively this quarter to update existing teaching materials and develop new material for recent releases of the StratusLab distribution. These materials were used in three tutorials, collectively with around 90 students. The first was at the EGI Technical forum in Lyon, France. The second was in Hanoi, Vietnam (1–3 November) as the cloud part of the ACGRID-III school. The third took place in Orsay, France (17–18 November). All tutorials used the LAL test cloud infrastructure. All were successful with good feedback from the students, although the Vietnam tutorial was affected significantly by a lack of resources because of air conditioning problems at LAL. The agendas for these tutorial are listed in the dissemination section.

Narval Tutorial From the 15-21 October 2011 there was the first “Narval School” organized by the Nuclear Physics Institute (IPN) in Orsay, France. Narval is a highly-distributed and modular framework for data acquisition. This tutorial took place on the StratusLab cloud infrastructure at LAL. The use of this infrastructure allowed the organizers to provide a flexible test infrastructure for the students with a minimum of effort. This tutorial, by an outside institute, was an important validation of the StratusLab software and shows the cloud’s utility for supplying temporary infrastructures.

Cloud Module in Engineering Curriculum A new module on “Cloud Computing and StratusLab” was added to the Information System and Networking degree program at the Faculty of Science and Technology in Marrakech, Morocco. This course has 18 students and a total duration of 24 hours. The goals of the course are to allow the students to identify the basic components, types, and modules of Cloud computing and understand how they function. StratusLab played a key role by providing a concrete set of IaaS services and an open infrastructure for experimentation.

CernVM CernVM is an initiative from CERN to provide virtual machine images adapted to the needs of the high-energy physics community. In particular, these images contain the software of the four Large Hadron Collider (LHC) experiments. LAL personnel have worked with the CernVM developers to allow these images to be used directly on a StratusLab cloud infrastructure. Differences in the contextualization procedure have been resolved (on both sides) to allow future versions of these images to run on StratusLab resources.

Alternate Schedulers The scheduling algorithm provided natively by OpenNebula is functional but quite basic. We have been approached by two groups with more elaborate schedulers that are interested to integrate them with the StratusLab distribution: first, to show that such integration is possible and second, to understand how they behave with real users. One group is from INRIA Lille Nord¹ in Lille, France; the other is based at Distributed Systems Research Lab (DSRL)² in Romania. Work continues with these groups on the integration.

E-business Applications Use Case An e-business application use case is being defined and implemented inside WP2 in order to validate the StratusLab distribution and show its advantages. This e-business application use case involves the deployment of a typical enterprise application (composed by front-end, business logic and database) in the StratusLab infrastructure. Their main requirements are simplicity (automatic deployment), multi-tier management, scalability and security (by LAN and firewall). Four images have been developed and are registered in the StratusLab Marketplace.

Bioinformatics Web Service Use Case We have added new bioinformatics tools to the “biocompute” appliance and updated existing ones. We have also added mechanisms for mounting automatically biological databases repository at start-up. Users wanting to run a biocompute instance connected to the biodatabases repository of the cloud infrastructure simply set the appropriate options when launching the machine. For example on IBCP’s cloud by setting the BIO_DB_SERVER parameter to the value ‘idb-databases.ibcp.fr’. We have also instrumented the appliances to use the StatusLab persistent disk functionality.

TOSCANI Use Case TOSCANI deals with the determination of protein structures based on Nuclear Magnetic Resonance (NMR) information. In collaboration

¹<http://www.inria.fr/centre/lille/>

²<http://dsrl.coned.utcluj.ro/>

with the M. Nilges group at the Institut Pasteur Paris, we matched the ARIA requirements with the StratusLab services. The main requirements are the management of the input and output data (at the scale of GB), the deployment of a complete infrastructure, including one master and up to 100 compute nodes for usual experiments. Shared storage, between the ARIA master and CNSsolve compute nodes, is required; additionally, the data produced need to be kept for further analysis. The persistent disk utility corresponds well to this requirement.

Interaction with French Bioinformatics Community Members of IBCP have participated in a workshop of the RENABI GRISBI community on 2-4 November in Rennes. We have discussed the requirements about representative bioinformatics databases and tools in the context of the distributed infrastructure GRISBI serving the national community. We have also demonstrated the StratusLab facilities using IBCP's cloud infrastructure. This is intended to foster acceptance and deployment of the StratusLab cloud distribution within this community and to gather feedback concerning their requirements.

3.1.3 Task 2.2: Intensive Evaluation of StratusLab Products

Refactored Tests To streamline the user-oriented testing of new StratusLab releases, the existing tests have been refactored to provide feedback on particular features and use cases. These refactored tests have been integrated into the Hudson continuous integration system, like their more monolithic predecessors. These provide systematic and rapid feedback for changes in the distribution. Moreover, these tests now use cloud resources to reduce contention for the available, physical testing resources and to avoid interference between different tests running in parallel.

Grid'5000 Initial investigations have been done on using the Grid'5000³ infrastructure to perform scalability tests of the StratusLab services. Use of the infrastructure has turned out to be more difficult than expected and unfortunately StratusLab services have not yet been deployed within Grid'5000 successfully. Work on this will continue in the next quarter.

StratusLab operations in a Bioinformatics Lab CNRS IBCP has continued its evaluation of the StratusLab distribution in our bioinformatics context. The release 1.1 is now in production on our local resources. We have deployed the new persistent disk component and have integrated it in our infrastructure. We integrated the PAT (Port Address Translation) implementation with the command line interface. We have begun an evaluation of specific machines with "big memory" responding to specific requirements of bioinformatics tools. We have run bioinformatics instances of 64GB and 96GB with success. We have also evaluated the portability of Intel-defined appliances on AMD CPUs. The "biocompute" appliance, defined on Intel-CPU server, was run on this server without notable differences.

³<https://www.grid5000.fr/mediawiki/index.php/Grid5000:Home>

SLCS Authentication and Authorization To fulfill a strong requirement from our scientific community of simple but robust access to our cloud by bioinformaticians, we have deployed in IBCP a authentication solution based on the SLCS framework⁴. This work has been done in collaboration with the Swiss SWITCH team maintaining SLCS software. SLCS is a service to create short live certificates based on Shibboleth authentication and a certificate authority (CA). Certificates are built with user attributes provided by Shibboleth identity provider and should be directly usable on all clouds accepting the CA. No modifications of StratusLab authentication layer are required. We have a working prototype and this is now in production on IBCP's cloud infrastructure.

3.1.4 Issues and Corrective Actions

Further Delay of MS4 The MS4 milestone has been further delayed. This is largely because we are waiting for the StratusLab 1.2 release that contains a large number of new features. That release has been delayed by serious air conditioning problems at LAL in this quarter and because of the need to streamline the build and test procedures. We expect production use of the 1.2 release by external sites in the next quarter.

⁴<http://www.switch.ch/grid/slcs>

3.2 WP3: Dissemination

Work Package 3 coordinates the project's activities in dissemination, collaboration, exploitation and sustainability. Its objectives are to disseminate results of the project to resource providers, end-users, and the general public; identify project contributions to standards bodies and standardization efforts; and coordinate interactions with related projects, developing Memoranda of Understanding between projects where appropriate.

3.2.1 Summary

Dissemination efforts have continued this quarter and though press coverage and website visits are down on Q5, the number of training events increased with three very well-attended hands-on tutorial sessions and a 89 people trained in the StratusLab software this quarter. A large number of talks and presentations featuring StratusLab were also given.

The project has put in place a new announcement mailing list and a user discussion forum which it is hoped will form the hub of a community support mechanism.

The project had great success at the EGI Technical Forum in September, with a booth and a number of presentations and participation in technical sessions. In addition the StratusLab poster "Virtualization of Bioinformatics Applications on Cloud Infrastructures" won the runner-up prize in the Best Poster competition.

A Memorandum of Understanding was signed with the IGE project, and a draft MoU has been drawn up with ERINA+ which will give StratusLab access to the impact assessment process of that project.

Work on Exploitation and Sustainability has increased, with participation in the DCI Sustainability meeting in October and continuing work within the project to identify avenues for sustainability.

3.2.2 Task 3.1: Dissemination

Demos The project is developing a number of demonstrations of the StratusLab software, use cases and features. These include a live interactive user-focused demo as well as several video demos of the functionality of the StratusLab software.

A proposal was submitted for a demo at the EU Innovation Convention to be held on 5–6 December 2011 in Brussels. The proposal was not accepted, but work on use cases to demonstrate continues and it is hoped that these will be used at future events.

Media & Publications StratusLab was featured in a piece in International Science Grid This Week (iSGTW) on 11 November entitled "European initiatives strengthen cooperation on cloud computing" which described how DCI projects are cooperating to enhance European Cloud infrastructures for research.

It was intended to produce some material for the popular press, but this has not yet been done and will happen instead in the next quarter.




Figure 3.1: Visits for Q6.

Website, Mailing Lists and Fora Figure 3.1 shows the number of visits to the website. The number for Q6 (4,985 visits) is down slightly from Q5 (5,472 visits).

The StratusLab Twitter feed now has 87 followers.

The announcements mailing list, which allows interested members of the public to keep up-to-date with news from the StratusLab project such as new releases, now has 72 members.

In October an additional “user announcements” mailing list was created, specifically for users of the StratusLab reference infrastructure. This is an announcement only mailing list used to inform users of upcoming service outages, upgrades and other important news. The list has 46 members, all of whom are users of the reference infrastructure.

A new public user forum was created and went live on 16 November, and has been running on a trial basis so far. The aim of the forum is to provide a contact point for StratusLab users and to build a community support mechanism. The idea of the forum appears to be feasible and the group already has 12 members, and has seen seven messages posted so far in two topics. It will thus be officially launched in December 2011 and disseminated more widely to StratusLab users and related communities.

EGI Technical Forum 2011 StratusLab had a very visible presence at the EGI Technical Forum which was held in Lyon, France from 19–23 September. The project had a booth displaying posters on Agile Methodology, Grid-Cloud Integration, Hadoop on StratusLab, The StratusLab Marketplace, StratusLab Reference Infrastructure, SlipStream and Cloud Storage. StratusLab-branded pens and post-it notes were printed and distributed from the booth. The project was also well represented in the technical programme of the conference with partner TCD’s David O’Callaghan co-chairing sessions in the Cloud and Virtualisation track.

Project members also gave four presentations at the conference covering topics such as cloud deployment, Agile Development, Bioinformatics on the cloud and Standards in StratusLab.

Christophe Blanchet of partner CNRS/IBCP won the runner-up prize in the poster competition for his poster “Virtualization of Bioinformatics Applications on

Cloud Infrastructures” and a successful hands-on tutorial session was also held.

Talks and Events Q6 saw an increase in the number of training events, with three hands-on tutorials held. The first at the EGI Technical Forum in September attracted about 25 participants. StratusLab participated in the 3rd School on Advanced Computing and GRID Technologies for Research (ACGRID3) in Hanoi, Vietnam in November, where 35 students were trained in using StratusLab over three days. Finally a two-day training event dedicated to StratusLab was held in November at the Linear Accelerator Laboratory in Orsay, France. This was attended by 25 people and the event was also streamed live on the internet with three participants following online. In total 89 people were trained in using StratusLab this quarter.

A large number of dissemination talks mentioning StratusLab were also given. A full list of talks and events is given in Table 3.1.

3.2.3 Task 3.2: Collaboration with Standards Bodies and Related Projects

Memoranda of Understanding

VENUS-C As part of the MoU signed between StratusLab and VENUS-C project, VENUS-C is evaluating OpenNebula with support from StratusLab, and StratusLab is evaluating storage and accounting solutions provided by VENUS-C.

EGI Collaboration with EGI continues in line with the MoU established with that project. An EGI GGUS support team to support sites running StratusLab was established in September which the projects hope will encourage more sites to run the StratusLab software.

StratusLab provided training and technical talks, as well as co-chairing a session at the EGI Technical Forum in September 2011 and the project has submitted a number of session abstracts (including another tutorial session) for the EGI Community Forum 2012. StratusLab has also committed to participate in technical integration sessions organised by EGI at the Community Forum.

Talks are in progress with EGI on the possibility of EGI using or running their own instance of the StratusLab Marketplace. This may provide a possible avenue for sustainability of the Marketplace.

StratusLab participated in a meeting on Sustainability of DCI projects organised by EGI in October. StratusLab is also providing input to the UMD roadmap, in particular EGI have asked for input in relation to messaging/notification integration.

EMI The technical collaboration with the EMI project has continued in Q6. StratusLab provided input to the EMI presentation on Cloud computing at the EGI Technical Forum in Lyon, France in September. In addition, the two projects worked together to prepare Torque images which were requested by EMI and have now been made available in the StratusLab Marketplace. StratusLab also provided support to EMI on the utilization of these VM images.

Table 3.1: Talks

Title / Event	Date
“Challenges in Federated and Hybrid Cloud Computing” at the 2011 International Conference on Parallel Computing, Ghent, Belgium	2011-08-30 – 2011-09-02
“Challenges in Federated and Hybrid Cloud Computing” at Cloud Day 2011, Stockholm, Sweden	2011-09-14
“StratusLab: status update and elastic deployment of services” at EGI Technical Forum, Lyon, France	2011-09-20
“StratusLab, Bioinformatics and the cloud” oral presentation at the EGI Technical Forum	2011-09-20
“Agile in StratusLab” at EGI Technical Forum, Lyon, France	2011-09-20
“StratusLab project: Standards, Interoperability and Asset Exploitation” at SIENA Initiative Workshop, EGI Technical Forum 2011, Lyon, France	2011-09-21
“Next Generation IaaS Cloud Computing with OpenNebula 3.0” at ISOD Workshop, OGF33, Lyon, France	2011-09-21
“StratusLab Tutorial” at EGI Technical Forum, Lyon, France	2011-09-22
“Virtualization of Bioinformatics Applications on Cloud Infrastructures” poster presentation and winner of runner-up prize in Poster Competition, EGI Technical Forum, Lyon, France	2011-09-19 – 2011-09-23
“StratusLab Booth” at EGI Technical Forum, Lyon, France	2011-09-19 – 2011-09-23
“Building Private Clouds for HPC with OpenNebula: Reference Deployments & Lessons Learned” at ISC Cloud 2011, Mannheim, Germany	2011-09-27
“StratusLab: Cloud Federation” at Federating IaaS Cloud for UK Research Workshop during the UK e-Science All Hands Meeting, York, UK	2011-09-28
“StratusLab Marketplace” at Atelier France Grilles, Lyon, France	2011-10-21
“OpenNebula Interoperability and Portability” at 5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management: Standards and the Cloud, Paris, France	2011-10-24
“StratusLab Tutorial” at 3rd School on Advanced Computing and GRID Technologies for Research (ACGRID3) Hanoi, Vietnam	2011-11-01 – 2011-11-03
“StratusLab Tutorial” Orsay, France	2011-11-17 – 2011-11-18
“Virtualisation de serveurs l’aide d’un logiciel de cloud” JRES 2011 in Toulouse, France	2011-11-23

IGE StratusLab signed an MoU with the IGE project at the EGI Technical Conference in Lyon, France in September. The MoU allows for the deployment of Globus services and tools on the StratusLab platform and creation of Globus Appliances for the StratusLab Marketplace. IGE will provide feedback on their use of the StratusLab tools as well as requirements for future developments. In addition the two projects will collaborate to deliver common training events.

SIENA StratusLab participated in the SIENA workshop at the EGI Technical Forum in Lyon, France in September, where Vangelis Floros of partner GRNET gave a presentation describing StratusLab's use of Standards and Interoperability.

The SIENA project also coordinated input from EU projects to the US National Institute of Standards and Technology (NIST) Cloud Computing Technology Roadmap, and StratusLab provided input in October to SIENA on how the projects solutions map to the NIST use cases and also participated in a number of phone conference meetings on this topic.

ERINA+ A draft MoU has been prepared between StratusLab and the ERINA+ project (which was signed in December). The collaboration will give StratusLab access to the impact assessment tools and process developed by ERINA+, in order to define a common methodology for the evaluation of impact of e-Infrastructures investments. It will allow StratusLab to measure the impact of the project's work in a methodologically sound way in collaboration with experts.

ERINA+ will provide a self-assessment online tool which is expected to be available from February 2012.

Mantychore The Mantychore project has expressed continued interest in liaising with StratusLab and Mantychore grid/cloud integration use cases will employ the StratusLab distribution.

3.2.4 Task 3.3: Development of Exploitation and Sustainability Plan

StratusLab participated in a meeting on Sustainability of DCI projects organised by EGI in October where the project presented its existing sustainability plan and the assets produced by the project.

Work continued at the StratusLab face-to-face meeting in Dublin from 30 November – 2 December on identifying the sustainability path for each of the project's assets. The discussions at this meeting will form the basis of the final Exploitation and Sustainability plan.

In addition to the work above, the commercial partners within the project have continued to exploit the outputs of the project. SixSq now provides a version of its SlipStream product which runs on StratusLab, partner Telefónica I+D has developed a use case for StratusLab based on cloud deployment of an e-commerce site.

Partner TCD is in discussions with the School of Business in Trinity College Dublin about the possibility of using StratusLab as a use case for business students

who would create a business plan for StratusLab as part of their final-year projects.

3.2.5 Issues and Corrective Actions

The project review in July 2011 called for an increase in dissemination to a wider audience, including the general public. Work has begun on writing articles for the popular science media, however, no such material has been published to date.

Indeed the overall press coverage for this quarter was quite low. This is partly due to the fact that only a single mainly bug-fixing release was made. The upcoming release of version 1.2 will be accompanied by a press release and it is expected that there will be more press coverage in Q7.

3.3 WP4: Software Integration and Distribution

This activity integrates and supports the StratusLab open-source cloud distribution. It integrates components required for grid site virtualization and dynamic scaling to remote clouds, addressing the specific requirements of the grid resource providers, and for the deployment of science clouds, addressing infrastructure cloud-like access demands from user communities, including industrial users.

WP4 works in close collaboration with WP5 for production deployment, WP6 for new service and component integration and all other work packages.

WP4 is also responsible to the execution of the project agile process, which includes active participation from all work packages.

3.3.1 Summary

During QR6, the integration and testing effort focused on the Persistent Disk Service, part of the StratusLab storage strategy. This important feature provides a significant boost in performance and scalability. The inclusion of OpenNebula 3.0 to the StratusLab distribution is also a focus, including several adaptations required on both sides to complete the integration.

The build and test infrastructure was improved by adding a certification staging step, with correspondent new infrastructure resources and jobs in our continuous integration server.

3.3.2 Task 4.1: Definition of Reference Architecture

The reference architecture of the StratusLab 2.0 distribution was discussed during the face to face meeting in Dublin, as well as during several TSCG conference calls, to ensure that it is still meeting the current project requirements and priorities. The addition of a jClouds client interface was added to the existing architecture.

3.3.3 Task 4.2: Integration of Open-source Distribution

In order to improve the pace at which we can release StratusLab, while reducing the effort in upgrading reference infrastructures and improve quality, a certification stage was introduced between the snapshot package generation and testing and the production infrastructure. To support this, WP5 deployed new machines dedicated to the certification tasks, such that stable packages can be deployed and tested, without being impacted by new snapshot packages being produced. To improve the deployment of StratusLab on this new testbed, new Hudson jobs were created in collaboration with WP5 such that these are automated.

A set of jobs is under development to generate, on-demand, new release package repositories, which contain release packages.

The integration of the Persistent Disk Server was significantly more complex than expected. Several important updates were required to the service implementation. Further, new OpenNebula instantiation scripts were required to integrate the Persistent Disk Server as a cache engine and quarantine. The image creation process was also significantly updated. The result is a significant boost in per-

formance, with for example near instantaneous image instantiation and a five fold speed-up factor for image creation.

Development of StratusLab client tools for Claudia The StratusLab client tools for Claudia has been integrated in the StratusLab distribution in this quarter. They are a set of python scripts, integrated with the StratusLab command line interface (CLI) tools, used to invoke the Claudia API, to deploy a service, undeploy it and obtain status information.

Integration of monitoring systems components in StratusLab Monitoring components software have been integrated in StratusLab in the monitoring-code git repository. Two new RPM packages have been created for installing the collector service and the probes. Finally, new jobs in Hudson have been created for the inclusion of the software in continuous integration.

Development of admin tools for Monitoring systems The development of admin tools for monitoring systems, which are being developing in WP6, has started in this quarter. They are a set of python scripts, integrated with the StratusLab command line interface tools, used to install the different components involved in monitoring systems (database, API, collector, probes and so on).

Integration of OpenNebula 3.0 to the StratusLab distribution OpenNebula 3.0 is being included to the StratusLab distribution. Several adaptations are required on both sides to complete the integration, including changes in the command line interface, changes in OpenNebula to refer to objects by name, as well as configuration generation by StratusLab tools.

3.3.4 Task 4.3: Contextualization of Grid Services

WP3 provided support to WP5 in the creation of automatic deployment of Torque-based clusters using SlipStream on the StratusLab reference infrastructure at GR-NET. This work provides the EGI/EMI community with an easier access to cloud resources for common computing patterns.

3.3.5 Task 4.4: Technical Support

Technical support continued to be provided, coordinated with WP5, on a wide range of issues, including the Persistent Disk Service, contextualisation and command-line tools.

3.3.6 Issues and Corrective Actions

LAL experienced a severe loss of cooling in its datacenter for a period of two weeks. This impacted seriously the StratusLab build and test infrastructure. This meant that continuous build and test jobs could not run during that time. During this time, technical tasks continued with regular updates to code and configuration, with however without feedback from the build and test jobs normally immediately alerting team members of problems after commit. This is in large part responsible for the fact that only one release of StratusLab was performed during this quar-

ter. The build and test jobs are now running without problems, but this came at a significant effort, to stabilize the build and test jobs. In the process, we improved the job structure, simplifying it and improving its performance (e.g. faster build and test) and introduced the certification and release jobs to reduce manual tasks in certifying new StratusLab release candidates. This should therefore improve the release rate for the rest of the project.

3.4 WP5: Infrastructure Operation

WP5 is responsible for the provision and operation of the project's computing infrastructure. It serves as a beta-tester of the software integrated by WP4 and WP6, deploying it in a production environment in order to verify its applicability for real life applications. In addition WP5 offers daily support to external users, either system administrators or those exploiting the project's public cloud services. Finally, WP5 is contributing with targeted development activities, related to the improvement of the cloud tools and services, as well as to showcase the capabilities of the StratusLab distribution to satisfy different use cases.

3.4.1 Summary

During Q6 work on Operations focused on the redesign the pre-production cloud services and the expansion of the resources providing the reference cloud service. In order to better streamline the adoption of the software developed by WP4 (an issue that has already been identified in the past reports) a new software certification process has been put in place and is being finalized. The main tools of this process will be the new certification pre-production site, the certification yum repository and a set of Hudson jobs that automate the deployment and test of release candidate versions of StratusLab software packages. Finally on the problems side, there have been some issues with the availability and reliability of the production grid site and the physical infrastructure hosted at LAL.

3.4.2 Task 5.1: Deployment and Operation of Virtualized Grid Sites

Expanded physical resources for reference cloud service The physical infrastructure providing the StratusLab reference cloud service that is located in GR-NET has been expanded with the addition of two more physical nodes. Thus the total resources provided now are 18 hosting nodes plus the front-end.

Reference cloud service upgrade to StratusLab 1.1 Shortly after the public release of v1.1 the reference service was upgraded to the latest version. In principle the set of services provided has not changed with the only exception of Claudia Service Manager which is now properly configured and available from the service front-end.

Expanded IP address range The range of public IP addresses that are offered to VM instances from the DHCP server has been expanded by 135. This means that now all the 62.217.122.0/32 subnet has been dedicated for VMs. As a result of this upgrade, lack of public IPs is no longer a limitation for the reference cloud service utilization.

Migration to new DNS domain The infrastructure has moved from the rather obscure **.one.yepeth.grnet.gr* domain to a more clean **.grnet.stratuslab.eu* domain. The primary domain is managed by LAL and the "grnet" subdomain is

managed by GRNET itself. This led to a change of the official hostname used for accessing the cloud front-end which now is `cloud.grnet.stratuslab.eu` (was `cloud-grnet.stratuslab.eu`). All the registered users had been informed of the change.

Redesigned support infrastructure The pre-production service has now been replaced by two individual testbeds: a development testbed and a certification testbed. The development testbed comprises two physical nodes and will be used for testing, experimentation of new features and overall development of new services. The development testbed typically runs a development version of the StratusLab distribution. The certification testbed also comprises two physical nodes (one fronted and one hosting node). The latter testbed will be used to test and certify the release candidate versions of StratusLab distribution. An automated procedure based on Hudson is currently under development (see testing section).

Development of the Torque cluster use case In the context of our collaboration with EMI there was a request to provide Torque ready VM appliances; essentially a stripped-down version of grid cluster without the rest of the grid software. As a result a new appliance has been developed and tested and is available from the Marketplace⁵. A single image is used to instantiate both the server and worker part of the cluster. A script that is available inside the image can be used to automate the installation process. Instructions on how to instantiate the cluster using the `stratus-run-cluster` command is available from the project web site⁶. In addition a StratusLab deployment module was developed as an alternative, more generic solution for cluster instantiation.

Refactored Persistent Disk Storage The persistent disk storage service was significantly refactored to make it more efficient, reliable, and testable. This was done to make it easier to implement image caching features developed by WP4.

Service Certificates for stratuslab.eu Domain The `stratuslab.eu` domain was transferred officially to CNRS allowing Terena certificates to be obtained for the `stratuslab.eu` domain. This benefits users of the StratusLab services because these certificates are trusted by default in all major browsers, avoiding spurious security warnings when using our services. These certificates will be rolled out progressively to all of the StratusLab public and private services.

3.4.3 Task 5.2: Testing of the StratusLab Toolkit

Automated certification of release candidate versions The goal of this task is to streamline the interaction between WP4 and WP5 and in particular the development, certification, and deployment process for new versions of the StratusLab distribution. In order to enable this, three components have been created:

- A certification yum repository that contains release candidate versions of rpm packages comprising the next version of the software distribution,

⁵Identifier: AViUGHkhd5NytPAKnIp03hnsVWG

⁶<http://stratuslab.eu/doku.php/tutorial:torquecluster>

- A certification testbed, described above. Provides the necessary physical resources for performing the deployment and verification tests, and
- A set of Hudson jobs that execute the necessary installation steps automating essentially the work that was done manually up to now in the former pre-production testbed.

The procedure is currently being finalized and will be put into action for the certification of version 1.2 expected near the beginning of Q7.

Storage Resources for LAL Cloud Service Over this quarter various improvements to LAL's cloud services have been done. In anticipation of opening the resources as a production service, a large amount of storage (40 TB) was acquired, half as a standard SAN and the other as a iSCSI server. The SAN has been integrated into the platform; the iSCSI server will be integrated in the next quarter. This will allow performance comparisons between the two systems.

Improved Monitoring for LAL's Resources LAL has upgraded all of its machines to use IPMI. This will allow more detailed monitoring of the resource utilization. In addition, a subset of machines has plug-level monitoring of the electricity consumption, allowing a cross-check of the IPMI values. This should be fully functional in Q7 and be able to provide feedback on resource utilization with different cloud configurations and services.

3.4.4 Task 5.3: Virtual Appliances Creation and Maintenance

Marketplace During this quarter development of the Marketplace has continued. The focus has been on improving testing, scalability and usability. Unit tests and resource-level testing have been provided. Processing of the metadata display table has been moved to the server-side. The implementation of the deprecation mechanism has been completed. A new client command *stratus-deprecate-metadata* has been provided, which allows users to easily deprecate Marketplace entries.

3.4.5 Issues and Corrective Actions

Degraded availability and reliability numbers for the production Grid site The availability and reliability numbers have been considerably affected in the past quarter due to problems with the site BDII service on the CE. The latter has been hanging regularly with no apparent reason making the site unusable for the rest of EGI infrastructure. Most probably this is a problem of the grid software and not of the underlying cloud service but the issue is still under investigation. For the time being we have tried to overcome the problem with closer monitoring of the service and the quick intervention from administrators to bring the service back online.

Slow adoption of new features by the reference cloud service Many of the technologies developed by WP4 still have not been adopted by WP5 including persistent storage and image caching. Will try to catch up in release 1.2. New features are difficult to adopt immediately since in many cases they bring radical

changes to the cloud service and require service redesign and long downtimes. Overall, this is the typical development-operations problem that we try to solve with the automated certification procedure.

Delay in delivery of D5.4 The deliverable D5.4 titled “*Economic analysis of infrastructure operations*” which was planned for PM18 has to be delayed until PM21. The delay has been requested in order to collect more data from the operation of the production grid site which has been running only a few months in full operation (since July 2011). We believe that this extension of the delivery date will serve better the purposes of the deliverable since it will give us more time to perform a thorough analysis of the data and thus provide better conclusions regarding the actual economic impact of grid site operation over cloud services.

Cooling Outage at LAL There was a significant outage of cooling in the LAL machine room for a few weeks around the beginning of November. This negatively impacted the standard build and test services and delayed anticipated releases of the StratusLab toolkit. It also greatly reduced the resources available for a StratusLab tutorial in Vietnam. This has since been corrected and hopefully there will not be a recurrence of this problem.

3.5 WP6: Innovative Cloud-like Management of Grid Services and Resources

The Joint Research Activity (JRA), carried out in WP6, develops advanced technology and features for deployment on existing Cloud infrastructures through automatic deployment and dynamic provision of grid services as well as scalable cloud-like management of grid site resources. More specifically, the objectives to be accomplished can be expressed as: i) the extension of currently available open-source service-level frameworks which provide elasticity on top of cloud infrastructures, ii) the invention of new techniques for the efficient management of virtualized resources for grid services and iii) the inclusion of novel resource provisioning models based on cloud-like interfaces.

3.5.1 Summary

The work done in QR6 mainly has involved the definition of the work plan for WP6 as documented in the deliverable D6.4. Monitoring system components have been included and extended in StratusLab in order to obtain metrics from different layers. Some work related to networking has been done in this period, mainly network isolation through VLAN and the firewall management. Monitoring, accounting and authentication have been improved in OpenNebula and some support has been integrated for user data injection.

3.5.2 T6.1: Dynamic Provision of Grid Services

Design of the Cloud-like Management of Grid Sites The deliverable D6.4 has updated the component design defined in D6.1 with the work to be done in Y2. The main contributions of this report involve the inclusion of networking and image management functionality as well as the inter-cloud broker. In addition, Cloud APIs have been described in the interoperability and some updates have been included for the monitoring and accounting functionalities.

Monitoring systems The monitoring component in Year 1 has been updated to take into account multi-tier application requirements. That is, monitoring takes metrics from both VMs and hosts and incorporates probes inside the VMs to measure different parameters. Besides hardware information (RAM, CPU, etc.), the monitoring framework gathers information for Key Performance Indicators and software metrics. A collector component acquires the metrics and stores them in a database. Finally, monitoring information is accessible via TCloud monitoring API.

3.5.3 T6.2: Scalable and Elastic Management of Grid Site Infrastructure

Support for User Data Injection in VMs A common requirement for VM contextualization is the ability to push user specific data into the VM, most notably

access keys. This requires first the ability to store arbitrary data associated with each user, and then a flexible method to select and include user data in the context device. Both requirements are now met by OpenNebula.

Network Isolation through VLAN tagging OpenNebula provides support for host-managed VLANs to restrict network access through VLAN tagging. This mechanism is compliant with the IEEE 802.1Q standard, but it requires support from the hardware switches. Also, OpenNebula allows administrators to restrict network access through VLAN tagging with Open vSwitch, a production quality, multilayer virtual switch.

Linux Bridge Firewall Management Alternatively to the use of VLAN tagging, it is possible to restrict network access through ebtables rules. The ebtables program enables transparent filtering of network traffic passing through a Linux bridge. This complements the previous automatic setup of simple firewall rules for TCP/UDP ports and ICMP traffic.

Accounting and monitoring The OpenNebula accounting module tracks information of the resource utilization that will be stored in predefined intervals of time. This information can be retrieved by User, Virtual Machine and Host. The OpenNebula statistics module keeps a predefined number of samples containing statistics for Hosts and Virtual Machines. These samples are built from the information that is retrieved for each resource by the OpenNebula Information Manager.

3.5.4 T6.3: Cloud-like Interfaces Specific for the Scientific Community

Improvements in the Auth Module The OpenNebula authentication and authorization system has been extended in three areas. First, it now avoids some potential security holes when the end-user may choose the driver to authenticate with OpenNebula, specially when using X509 certificates. Second, the security of public cloud API has been strengthened by including special server accounts to run the services, these server processes can use symmetric cryptographic ciphers or X509 to authenticate with OpenNebula. And third, there have been improvements in the X509 and SSH authentication methods like native support for proxies or better support for DN strings.

3.5.5 Issues and Corrective Actions

The document D6.4 was delayed and instead sent at the end of Q6. Some additional analysis and discussion have taken place to clarify and define the work to be done in WP6 in the coming months. Moreover, the document has been reviewed and rewritten in order to have a better quality.

4 Project Management

4.1 Consortium

The project consortium consisting of six partners (CNRS, UCM, GRNET, SIXSQ, TID, and TCD) has not changed since the start of the project. There have been no changes in the legal status of those partners. The representatives for TCD and TID have changed because of retirements and internal reorganization of activities.

4.2 Management Tasks

Meetings Tables 4.1–4.6 contain a list of the meetings by quarter that have been planned to foster collaboration between the project participants. Not listed are the planning meetings for each development sprint and the daily standup meetings.

Metrics Table 4.7 contains the metrics for the project. The table groups related metrics together.

Deliverables and Milestones Tables 5.1, 5.2, and 5.3 list all of the documents. In addition, these are available from the project website.

Deliverable D4.4, originally foreseen for PM15, was delivered in PM16. Deliverable D6.4 that updates the design for advanced services in WP6 has been produced in this quarter, although also with a slight delay. Deliverable D5.4 has been moved from PM18 to PM21 to allow more time to collect data from the cloud infrastructure.

Memoranda of Understanding The project has signed MoUs with VENUS-C, EMI, IGE, and ERINA+ in this quarter.

4.3 Issues

Year 2 Budget In response to the resource utilization reported for the first reporting period, a new budget is being developed for the second year, taking into account the underspending from the first year. Unfortunately compiling this budget and collecting the necessary information from the partners has been more difficult than foreseen. Complete information now is not expected until January 2012.

Review Recommendations The responses to the review recommendations (included in this report) have been updated to reflect the situation as of the end of Q6. They will continue to be updated over remainder of the project.

4.4 Planning

4.4.1 Objectives for Next Quarter

- Solidify the v1.0 StratusLab cloud distribution through increased testing and hardening of existing services.
- Support for a second operating system to ensure the portability of the distribution.
- Produce initial beta release of v2.0 of the StratusLab cloud distribution.
- Dissemination of project results of a completed use case.
- Expanding the sites using StratusLab and completion of MS4.
- Continued operation of reference infrastructure and support to users and system administrators.
- Operate a pre-production elastic Grid site elasticity, verify the applicability of the technology and move results to a production site.
- Evaluate GPFS as a backend storage solution. Prioritize and evaluate additional file systems.
- Develop additional use cases similar to MapReduce (e.g. Matlab application showcase)
- Integrate caching sub-system within the production cloud service
- Integrate NFS persistent storage service in the reference cloud service
- Integration of OpenNebula 3.0 with initial tests of new networking functionality

4.4.2 Roadmap

The roadmap remains essentially the same as decided in the Lyon Face-to-Face meeting. The PMB in Q3 gave its formal approval of the following changes to the overall work program:

1. The tasks regarding having a public (user-visible) cloud and an associated cloud API have been moved from Y2 to Y1, largely because of interest from scientific communities and resource centers wanting to provide public clouds.
2. The tasks about hybrid clouds will be expanded to include also cloud federation models. This will be moved to Y2 to balance the change above. Also having a solid release will make these investigations easier.

3. As foreseen in the TA, the appliance repository consists of a single service that contains appliance metadata, appliance storage, and services for changing appliance formats. This has been split into different services. The Marketplace will handle appliance metadata. Storage will take place with normal cloud storage or outside of the cloud. Instead of providing a service for appliance format changes, client tools will be provided instead.

These changes have been made and followed at the technical level for sometime; they are now also agreed at the management level.

The architecture and roadmap have been re-evaluated in D4.4 along with D6.4. The project will continue to make incremental changes to the existing distribution. The priority for Year 2 is the demonstration of hybrid cloud functionality.

Table 4.1: Meetings (Q1)

Title	Date	Venue	Comments
StratusLab	14-15/06/2010	Orsay, FR	Kick-off of project. Detailed planning for accomplishing objectives.
Kick-Off Meeting			http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1129
Technical Meeting	22/07/2010	Madrid, ES	Detailed technical discussions for StratusLab development.
			http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1189
Sprint 1 Demo	30/07/2010	Phone/EVO	Sprint 1 demonstration meeting.
			http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1191
Sprint 2 Demo	20/08/2010	Phone/EVO	Sprint 2 demonstration meeting.
			http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1192

Table 4.2: Meetings (Q2)

Title	Date	Venue	Comments
Project Management Board	03/09/2010	Phone	PMB meeting to decide IPR policies. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1203
Sprint 3 Demo	10/09/2010	Phone/EVO	Sprint 3 demonstration meeting. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1203
Technical Meeting (TSCG)	21/09/2010	Phone/EVO	Shaping StratusLab distribution. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1213
WP6 research lines meeting	27/09/2010	Madrid, ES	Discussion about the main gaps identified in WP4 and some technologies to solve them. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1318
WP6 kickoff meeting	07/10/2010	Phone	Presentation of the lines to work on WP6 and distribution of work. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1320
Sprint 4 Demo	08/10/2010	Phone/EVO	Sprint 4 demonstration meeting. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1232
WP6 monitoring and accounting	26/10/2010	Phone	Audioconference about monitoring and accounting in StratusLab. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1321
Sprint 5 Demo	08/11/2010	Phone/EVO	Sprint 5 demonstration meeting. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1255
Face-to-Face Technical Meeting	15-16/11/2010	IBCP, Lyon, France	Discussion of StratusLab roadmap. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1243
Project Management Board	22/11/2010	Phone	Project overview; LoS policy. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1263

Table 4.3: Meetings (Q3)

Title	Date	Venue	Comments
Sprint 6 Demo	09/12/2010	Phone/EVO	Sprint 6 demonstration meeting. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1310
Sprint 7 Demo	17/12/2010	Phone/EVO	Sprint 7 demonstration meeting. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1323
Technical Meeting (TSCG)	27/01/2011	Phone/EVO	Feedback from EGI; priorities for distribution. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1213
Sprint 8 Demo	31/01/2011	Phone/EVO	Sprint 8 demonstration meeting. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1423
Technical Meeting (TSCG)	17/02/2011	Phone/EVO	Error reporting; priorities for next sprint. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1213
Sprint 9 Demo	18/02/2011	Phone/EVO	Sprint 9 demonstration meeting. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1442
Project Management Board	24/02/2011	Phone	Project status; MoUs; effort utilization; review planning. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1440

Table 4.4: Meetings (Q4)

Title	Date	Venue	Comments
Sprint 10 Demo	03/03/2011	Phone/EVO	Sprint 10 demonstration meeting. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1448
Technical Meeting (TSCG)	03/03/2011	Phone/EVO	Review of developments and priorities. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1460
Sprint 11 Demo	31/03/2011	Phone/EVO	Sprint 11 demonstration meeting. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1470
Metadata & Marketplace Demo	08/04/2011	EVO	Demo for HEPiX Virtualization Working Group. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1477
Sprint 12 Demo	29/04/2011	Phone/EVO	Sprint 12 demonstration meeting. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1492
Grid site deployment with Claudia (TID, GRNET)	09/05/2011	Phone	Discussion about how to use Claudia for the deployment of a grid site. http://indico2.lal.in2p3.fr/indico/conferenceTimeTable.py?confId=1530#20110509
Technical Meeting (TSCG)	10/05/2011	Phone	Persistent storage and cloud interfaces. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1526
Interproject Collaboration	11/05/2011	Amsterdam	StratusLab, HPC Cloud, and Mantychore discussions. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1510
Sprint 13 Demo	16/05/2011	Phone/EVO	Sprint 13 demonstration meeting. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1513
Integration Meeting	23-24/05/2011	Geneva	F2F meeting for 1.0 release. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1503
Interproject Collaboration	27/05/2011	Phone	Discussion with Contrail project. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1527
Grid site deployment and scalability (TID, GRNET)	27/05/2011	Phone	Discussion to align the work. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1529

Table 4.5: Meetings (Q5)

Title	Date	Venue	Comments
Sprint 14 Demo	10/06/2011	Phone/EVO	Sprint 14 demonstration meeting. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1533
StratusLab Face-to-Face Meeting	21-23/06/2011	Geneva, CH	Integration of software. Update of roadmap. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1502
Sprint 15 Demo	23/06/2011	Phone/EVO	Sprint 15 demonstration meeting. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1541
StratusLab First Periodic Review	04/07/2011	Brussels, BE	External review of project's progress. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1501
Sprint 16 Demo	29/07/2011	Phone/EVO	Sprint 16 demonstration meeting. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1557
Technical Meeting (TSCG)	25/08/2011	Phone/EVO	Priorities for upcoming sprints. Architecture for StratusLab 2.0. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1560
Sprint 2 Demo	20/08/2010	Phone/EVO	Sprint 2 demonstration meeting. http://indico.lal.in2p3.fr/conferenceDisplay.py?confId=1192

Table 4.6: Meetings (Q6)

Title	Date	Venue	Comments
Project Management Board	12/09/2011	Phone	Project status; Reviewer Feedback; Y2 Effort and Budgets. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1561
Sprint 17 Demo	16/09/2011	Phone/EVO	Sprint 17 demonstration meeting. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1574
Technical Meeting (TSCG)	26/09/2011	Phone/EVO	Architecture review; priorities for upcoming sprints. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1575
Technical Meeting (TSCG)	17/10/2011	Phone/EVO	Use Cases for Y2 Review; priorities for upcoming sprints. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1613
Sprint 18 Demo	19/10/2011	Phone/EVO	Sprint 18 demonstration meeting. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1617
Sprint 19 Demo	16/11/2011	Phone/EVO	Sprint 19 demonstration meeting. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1638
StratusLab Cloud Training	17–18/11/2011	Orsay, France	General training for users and administrators. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1565
Technical Meeting (TSCG)	21/11/2011	Phone/EVO	F2F agenda; priorities for upcoming sprints. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1642
Project Management Board	30/11/2011	Phone	Project status; ERINA+ MoU; Reviewer Recommendations; Sustainability. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1650
StratusLab Face-to-Face Meeting	30/11/2011–02/12/2011	Dublin, Ireland	Integration of software; Update of roadmap. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1621

Table 4.7: Metrics

Metric	Q2	Q3	Q4	Y1 Target	Q5	Q6	Q7	Q8	Y2 Target
No. of people trained on StratusLab software	N/A	N/A	~ 25	–	~36	~89			–
No. of people on StratusLab announcement list	67	67	67	25	70	72			75
Registered users on StratusLab discussion site	N/A	N/A	N/A	50	N/A	12			100
No. of views of website	2922	4623	4579	–	5472	4985			–
No. of completed sprints	5	5	4	–	3	3			–
No. of releases	1	1	1	–	2	1			–
No. of open user stories	38	72	101	–	118	107			–
No. of implemented user stories	69	40	67	–	50	48			–
No. of open bugs	6	15	22	–	28	51			–
No. of fixed bugs	7	11	27	–	14	20			–
No. of prod. sites running StratusLab dist.	1	1	1	5	1	3			10
Availability of hosted grid sites	N/A	N/A	100%	80%	91%	74%			95%
Reliability of hosted grid sites	N/A	N/A	100%	80%	92%	78%			95%
No. of VOs served via StratusLab hosted grid sites	0	1	1	10	21	18			30
No. of sci. disciplines served via StratusLab hosted grid sites	0	0	0	3	11	9			15
Delivered computing resources through hosted grid services	N/A	16 cores	16 cores	–	32 cores	32			–
Delivered computing resources through hosted cloud services	N/A	256 cores	256 cores	–	256 cores	288			–
Storage provided through cloud service	N/A	N/A	N/A	–	0	0			3 TB
No. of jobs run in hosted grid site	N/A	N/A	N/A	–	13,960	16,916			–
Norm. CPU time consumed in the hosted grid site (hrs)	N/A	N/A	N/A	–	26,202	14,231			–
No. base machine images	5	7	8	5	8	13			10
No. of base machine image downloads	783	2628	7072	–	7225	6657			–
No. appliances	0	6	7	5	7	7			15
No. of appliance downloads	0	252	687	–	1010	426			–

5 Deliverables and Milestones

Tables 5.1 and 5.2 show the deliverables for the first and second years of the project. Table 5.3 lists all of the milestones. All of the deliverables and milestones for the first year of the project have been produced and submitted as foreseen in the project's roadmap. All of these are available from the project's website¹.

Two technical notes have also been produced during the first year: "StratusLab Marketplace" describing the technical specification of the Marketplace and "Installing and operating a production grid site in the StratusLab cloud: Experience and issues" providing feedback to developers and advice to administrators running grid services within a cloud. These notes are also available from the project website.

¹<http://stratuslab.eu/doku.php/deliverables>

Table 5.1: Deliverables (Year 1)

No.	Title	Version	WP No.	Lead Beneficiary	Nature	Diss. Level	Due Date	Actual Date	Status	Contractual	Comments
D2.1	Review of the Use of Cloud and Virtualization Technologies in Grid Infrastructures	1.2	WP2	CNRS	R	PU	PM2	11/08/2010	Done	Yes	
D4.1	Reference Architecture for StratusLab Toolkit 1.0	1.0	WP4	SIXSQ	R	PU	PM3	14/09/2010	Done	Yes	
D5.1	Infrastructure Specification	1.0	WP5	GRNET	R	PU	PM3	14/09/2010	Done	Yes	
D3.1	Initial Plan for Dissemination, Collaboration and Standardization Activities	1.0	WP3	TCD	R	PU	PM4	18/10/2010	Done	Yes	
D6.1	Cloud-like Management of Grid Sites 1.0 Design Report	1.0	WP6	TID	R	PU	PM5	16/11/2010	Done	Yes	
D5.2	Infrastructure Tool and Policy Specification	1.0	WP5	GRNET	R	PU	PM6	15/12/2010	Done	Yes	
D6.2	Cloud-like Management of Grid Sites 1.0 Software	1.1	WP6	TID	P	PU	PM11	13/05/2011	Done	Yes	
D2.2	Report on Evaluation of StratusLab Products	1.0	WP2	CNRS	R	PU	PM12	15/06/2011	Done	Yes	
D3.2	Report on Dissemination, Collaboration and Standardization Activities	1.1	WP3	TCD	R	PU	PM12	16/06/2011	Done	Yes	
D3.3	Exploitation and Sustainability First Plan	1.1	WP3	TCD	R	PU	PM12	16/06/2011	Done	Yes	
D4.2	StratusLab Toolkit 1.0	1.0	WP4	SIXSQ	P	PU	PM12	15/06/2011	Done	Yes	
D4.3	First Year Software Integration Report	1.0	WP4	SIXSQ	R	PU	PM12	15/06/2011	Done	Yes	
D5.3	First Year Infrastructure Operations Report	1.1	WP5	GRNET	R	PU	PM12	16/06/2011	Done	Yes	
D6.3	First Year Cloud-like Management of Grid Sites Research Report	1.0	WP6	TID	R	PU	PM12	15/06/2011	Done	Yes	

Table 5.2: Deliverables (Year 2)

No.	Title	Version	WP No.	Lead Beneficiary	Nature	Diss. Level	Due Date	Actual Date	Status	Contractual	Comments
D2.3	Survey of Targeted Communities Concerning StratusLab		WP2	CNRS	R	PU	PM14	12/08/2011	Done	Yes	
D4.4	Reference Architecture for StratusLab Toolkit 2.0		WP4	SIXSQ	R	PU	PM15	03/10/2011	Done	Yes	Delayed PM15-16
D6.4	Cloud-like Management of Grid Sites 2.0 Design Report		WP6	TID	R	PU	PM17	15/12/2011	Done	Yes	Delayed PM17-18
D5.4	Economic Analysis of Infrastructure Operations		WP5	GRNET	R	PU	PM18		Delayed	Yes	Delayed to PM21
D6.5	Cloud-like Management of Grid Sites 2.0 Software		WP6	TID	P	PU	PM23			Yes	
D2.4	Final Report on StratusLab Adoption		WP2	CNRS	R	PU	PM24			Yes	
D2.5	Report on Evaluation of StratusLab Products		WP2	CNRS	R	PU	PM24			Yes	
D3.4	Final Review of Dissemination, Collaboration and Standardization Activities		WP3	TCD	R	PU	PM24			Yes	
D3.5	Exploitation and Sustainability Final Plan		WP3	TCD	R	PU	PM24			Yes	
D4.5	StratusLab Toolkit 2.0		WP4	SIXSQ	P	PU	PM24			Yes	
D4.6	Software Integration Final Report		WP4	SIXSQ	R	PU	PM24			Yes	
D5.5	Infrastructure Operations Final Report		WP5	GRNET	R	PU	PM24			Yes	
D6.6	Cloud-like Management of Grid Sites Research Final Report		WP6	TID	R	PU	PM24			Yes	

Table 5.3: Milestones

No.	Title	WP No.	Lead Beneficiary	Due Date	Achieved	Actual Date	Comments
MS1	Establishment of Management Infrastructure and Metrics Definition	WP1	CNRS	PM3	Yes	1/09/2010	
MS6	Website Operational	WP3	TCD	PM3	Yes	6/09/2010	
MS2	Contact Procedures and Supporting Tools for Targeted Communities	WP2	CNRS	PM4	Yes	10/12/2010	
MS7	StratusLab Development, Certification and Release Procedures in Place	WP4	SIXSQ	PM6	Yes	10/12/2010	
MS3	Creation of Virtual Appliances for Bioinformatics Community	WP2	CNRS	PM9	Yes	14/03/2011	
MS10	Initial virtual appliance repository	WP5	GRNET	PM9	Yes	4/03/2011	
MS14	Release of Cloud-like Management of Grid Services and Resources 1.0 Beta	WP6	TID	PM9	Yes	14/03/2011	
MS8	Release of StratusLab 1.0 Beta	WP4	SIXSQ	PM10	Yes	05/04/2011	
MS11	Operation of Site Running StratusLab toolkit v1.0	WP5	GRNET	PM10	Yes	04/04/2011	
MS4	Adoption of StratusLab Software by External Grid Sites	WP2	CNRS	PM14			Delayed
MS12	Delivery of Virtual Appliance Repository	WP5	GRNET	PM18	Yes	15/12/2011	
MS5	Opening of Virtual Appliances Repository to External Application Communities	WP2	CNRS	PM20			
MS15	Release of Cloud-like Management of Grid Services and Resources 2.0 Beta	WP6	TID	PM21			
MS9	Release of StratusLab 2.0 Beta	WP4	SIXSQ	PM22			
MS13	Operation of Site Running StratusLab Toolkit v2.0	WP5	GRNET	PM22			

6 Use of Resources

The effort and spending plan for year 2 is being revised. Collecting this information from the partners has been much more difficult than foreseen. Complete information is expected in January 2011. However, it is clear that the underspending continues and the project will have to develop strategy for dealing with this.

Glossary

APEL	Accounting Processor for Event Logs (EGI accounting tool)
Appliance	Virtual machine containing preconfigured software or services
CDMI	Cloud Data Management Interface (from SNIA)
CE	Computing Element in EGI
DCI	Distributed Computing Infrastructure
DMTF	Distributed Management Task Force
EGEE	Enabling Grids for E-sciencE
EGI	European Grid Infrastructure
EGI-TF	EGI Technical Forum
GPFS	General Parallel File System by IBM
Hybrid Cloud	Cloud infrastructure that federates resources between organizations
IaaS	Infrastructure as a Service
iSGTW	International Science Grid This Week
KPI	Key Performance Indicator
LB	Load Balancer
LRMS	Local Resource Management System
MoU	Memorandum of Understanding
NFS	Network File System
NGI	National Grid Initiative
OC CI	Open Cloud Computing Interface
OVF	Open Virtualization Format
Public Cloud	Cloud infrastructure accessible to people outside of the provider's organization
Private Cloud	Cloud infrastructure accessible only to the provider's users
SE	Storage Element in EGI
SGE	Sun Grid Engine
SNIA	Storage Networking Industry Association
TCloud	Cloud API based on vCloud API from VMware
VM	Virtual Machine
VO	Virtual Organization
VOBOX	Grid element that permits VO-specific service to run at a resource center
Worker Node	Grid node on which jobs are executed

XMLRPC	XML-based Remote Procedure Call
YAIM	YAIM Ain't an Installation Manager (configuration utility for EGI)