Enhancing Grid Infrastructures with
Virtualization and Cloud Technologies

Project Quarterly Report

Quarterly Report QR5 (V1.0)
15 September 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.

Contributors

Name	Partner	Sections
Marc-Elian 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 Sep. 2011	Initial draft with new format.
1.0	15 Sep. 2011	Final version.

Contents

List of Figures

List of Tables

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 fifth quarter, the project worked to finish and to disseminate the first production version of the StratusLab cloud distribution (v1.0). This release contains all of the core functionality required for an IaaS cloud deployment. Following this successful release, the project prepared for its first periodic review which took place at the beginning of July. For the remainder of the summer, work continued to consolidate services within the StratusLab distribution and to plan the work for the coming year, in addition to the standard operation of the reference infrastructure and other services. Highlights of the activities are given below.

Definition of Use Cases The D2.3 deliverable "Survey of Targeted Communities Concerning StratusLab" contains a set of identified use cases that can demonstrate the benefits of cloud technologies for real users and can serve as subjects for dissemination activites. Seven use cases have been defined and will be scheduled for implementation over the next year.

General StratusLab Paper A general paper concerning StratusLab's goals and results has been prepared and will appear in the book "European Research Activities in Cloud Computing" in March 2012. This provides a good technical overview of the project and will be used for dissemination and training activities.

Dissemination for v1.0 The v1.0 release of the StratusLab cloud distribution was a critical milestone for the project, marking the first production release of our software. Articles were published by EGI.eu and in iSGTW and news of the release was carried by HPC in the Cloud and HPCwire.

Evolution of Distribution The distribution continues to evolve including more functionality and becoming more robust. Improvements to the distribution include better integration with the persistent storage, feedback from operations, and client tools for Claudia. Code for RESTful tests has been included in the Marketplace to allow for better testing of the service. Work on including OpenNebula v3.0 continues.

MapReduce To demonstrate the appliability of StratusLab to different use cases, Hadoop was deployed inside of a StratusLab cloud. An appliance was create to

support this and is generally available through the Marketplace. SlipStream from SixSq was used to deploy the system, demonstrating the integration of commercial tools with the StratusLab platform as well as the features of StratusLab itself.

New Scaling Policies The ability for autoscaling deployments via Claudia has been expanded with a number of new scale down policies added. This allows better control of the resources used in service deployements.

Most of the detailed objectives for Q5 have been achieved. The update of requirements from users and system administrators was not very successful because of the low number of survey responses. Nonetheless, we believe that the extensive requirements collected at the start of the project are still valid and form a reasonable basis for guiding the StratusLab work. The other objective not achived in Q5 was the full support of a second operating system. This work will be pushed into Q6.

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@grnet.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

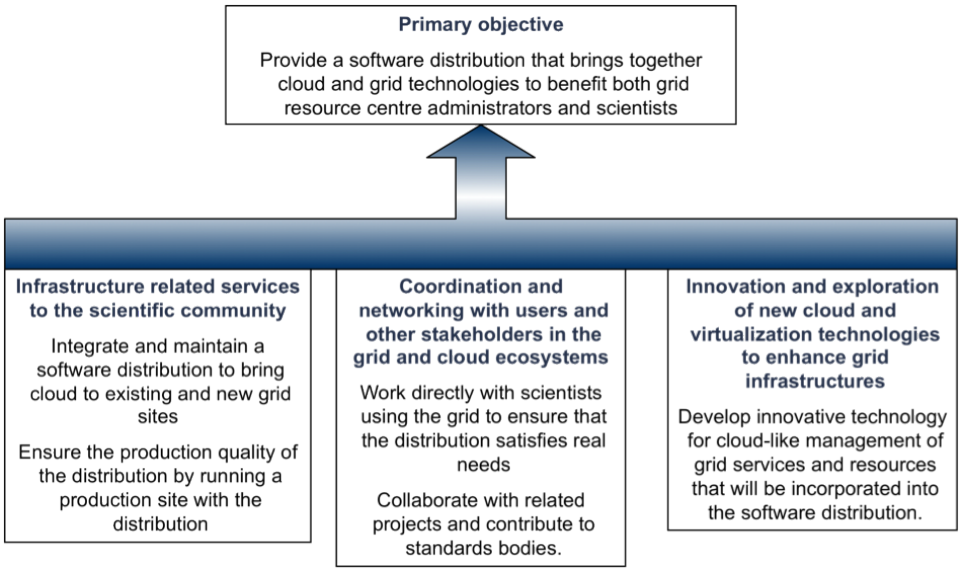

Figure 2.1: *Primary and supporting objectives*

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 pro-vision 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

2.2.4 Quarter 8

2.3 Review Recommendations

The responses given below to the reviewer recommendations are preliminary as the project personnel have had little time to react given the arrival of the formal recommendations in mid-August. The responses will continue to be updated as the recommendations are further analyzes and solutions are implemented.

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.*
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) communitites 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. 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.

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.

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 at the same time as this report. 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.

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. Because of the arrival of these recommendations at the end of the quarter and the difficulty of obtaining financial information from the partners' administrations in August, the new budget will not be available until the end of September 2011.

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 sould 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 will also look at alternative cloud APIs like jclouds that are becoming defacto standards.

    The StratusLab technology is not grid-specific; the grid only serves as a concrete use case; nonetheless, the grid aspects will be de-emphasized in

preference to scientific and commercial use cases in the second year of the project.

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.*

    The response to this recommendation requires more analysis. Our initial 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.

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 StratusLab 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 solit relationships with stakeholders outside the traditional high energy physics (HEP) community.*

    The PMB is defining the profile of the person to fill this role and thinking of possible candidates.

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 sucessfully implemented.

18. *Future reports should contain less "fat" and should be more cripy 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 is a proposal for a booth at the upcoming Innovation Convention in Brussels. This will require materials targeted at a general, non-technical audience.

3 Progress and Achievements

In the fifth quarter, the project worked to finish and to disseminate the first production version of the StratusLab cloud distribution (v1.0). This release contains all of the core functionality required for an IaaS cloud deployment. Following this successful release, the project prepared for its first periodic review which took place at the beginning of July. For the remainder of the summer, work continued to consolidate services within the StratusLab distribution and to plan the work for the coming year, in addition to the standard operation of the reference infrastructure and other services. Highlights of the activities are given below.

Definition of Use Cases The D2.3 deliverable "Survey of Targeted Communities Concerning StratusLab" contains a set of identified use cases that can demonstrate the benefits of cloud technologies for real users and can serve as subjects for dissemination activites. Seven use cases have been defined and will be scheduled for implementation over the next year.

General StratusLab Paper A general paper concerning StratusLab's goals and results has been prepared and will appear in the book "European Research Activities in Cloud Computing" in March 2012. This provides a good technical overview of the project and will be used for dissemination and training activities.

Dissemination for v1.0 The v1.0 release of the StratusLab cloud distribution was a critical milestone for the project, marking the first production release of our software. Articles were published by EGI.eu and in iSGTW and news of the release was carried by HPC in the Cloud and HPCwire.

Evolution of Distribution The distribution continues to evolve including more functionality and becoming more robust. Improvements to the distribution include better integration with the persistent storage, feedback from operations, and client tools for Claudia. Code for RESTful tests has been included in the Marketplace to allow for better testing of the service. Work on including OpenNebula v3.0 continues.

MapReduce To demonstrate the appliability of StratusLab to different use cases, Hadoop was deployed inside of a StratusLab cloud. An appliance was create to support this and is generally available through the Marketplace. SlipStream from SixSq was used to deploy the system, demonstrating the integration of commercial tools with the StratusLab platform as well as the features of StratusLab itself.

New Scaling Policies The ability for autoscaling deployments via Claudia has been expanded with a number of new scale down policies added. This allows better control of the resources used in service deployements.

Most of the detailed objectives for Q5 have been achieved. The update of requirements from users and system administrators was not very successful because of the low number of survey responses. Nonetheless, we believe that the extensive requirements collected at the start of the project are still valid and form a reasonable basis for guiding the StratusLab work. The other objective not achived in Q5 was the full support of a second operating system. This work will be pushed into Q6.

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. Those 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

After preparations for the first EU review of the project, the primary activity in this quarter has been the work on the deliverable D2.3 "Survey of Targeted Communities Concerning StratusLab". This deliverable identified seven concrete use cases to be implemented during the second year of the project. Approximately half of them are "internal" use cases that can be implemented using only StratusLab personnel; the other half involve interaction with scientists and engineers outside of the project. Once implemented, the results will be disseminated to the wider community (in cooperation with WP3) to encourage broader usage of the StratusLab cloud distribution.

This deliverable represents the first action in refocusing the WP2 efforts to bring more users onto StratusLab cloud infrastructures. This was the principal recommendation for WP2 coming from the first European review in June. This refocusing of purpose will continue as concrete plans and schedules for the identified use cases are made.

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

Survey of Targeted Communities (D2.3) The primary activity of the last quarter has been the deliverable D2.3 "Survey of Targeted Communities Concerning StratusLab". In keeping with the recommendations coming from the first European review of the project, the focus of this deliverable shifted slightly to identify specific use cases that should be implemented during the second year of the project. Seven different use cases have been identified; three of which are "internal" use cases that can be done using only effort within the project and the remaining four involve scientists and engineers outside of the project. An attempt was also made to update the requirements collected in D2.1. However, this was largely unsuccessful because of the summer holidays. Instead the project will continue to rely on the requirements listed in D2.1 and feedback from people as we work to fulfill the identified use cases in D2.3.

General Overview of StratusLab A general overview of the project's goals, developments, and results has been prepared based on a presentation at the Second Workshop on Software Services (WoSS) in Romania. This overview will appear as a chapter in the book "European Research Activities in Cloud Computing" (Dana

Petru and Jose Luis Vasquez Poletti, editors) that should appear in March 2012.

Bioinformatics Use Cases CNRS IBCP has participated to the "Survey of Targeted Communities Concerning StratusLab" described in the project Deliverable D2.3. We have proposed two use cases: one about Bioinformatics Web Services and the other, TOSCANI, dealing with the determination of protein structures based on Nuclear Magnetic Resonance (NMR) information. The goal of the first bioinformatics use case is to create bioinformatics appliances containing the usual tools that scientists and engineers can deploy on demand. To meet researcher's expectations, these appliances must present a standard programmatic, public, web service interface, permitting users to combine the different bioinformatics methods in useful analysis pipelines. The second bioinformatics use case has been defined in collaboration with the M. Nilges group at the Institut Pasteur Paris. TOSCANI, "TOwards StruCtural AssignmeNt Improvement" concerns the scientific disciplines around (i) molecular and structural biology (determination of biomolecular structures up to atomic resolution) and (ii) bioinformatics, which includes the ensemble of computer algorithms for treating the data from biological systems. We propose there to demonstrate the flexibility of the cloud to deploy the different bioinformatics tools required to improve the determination of protein structures based on NMR information.

Interaction with French Bioinformatics Community Members of IBCP have participated to a meeting of the RENABI GRISBI community on June 21-22 in Lyon. Different points were discussed about the distributed infrastructure GRISBI serving the national community and the reference bioinformatics data and tools that are deployed. This is designed to foster acceptance and deployment of the StratusLab Toolkit within this community and to gather their requirements.

Telefónica Private Cloud Contact Telefónica is offering a Private Cloud for the Spanish regional administration, where TID is developing part of this solution. In fact, they are developing the product from ideas and the service manager prototype, which is released in StratusLab. In the quarter, some contacts have been done with the product leader in order to obtain feedback with end-users of the service.

Commercial Use Case Identification In order to provide requirements and feedback from a commercial perspective, TID has been working on the identification of a commercial use case to be used in order to validate the StratusLab distribution and show their advantages.

3.1.3 Task 2.2: Intensive Evaluation of StratusLab Products

Evaluation of StratusLab Installation in a representative Bioinformatics Lab CNRS IBCP has evaluated the manual and the quattor-based installation modes of the StratusLab framework. Several issues have been identified and reported to the project's developers, who have corrected most of them. Deploying a cloud site on CNRS IBCP hardware resources aims at identifying site constraints usually encountered with national bioinformatics sites (not enough public IPs, NAT,

etc.), to provide bioinformatics scientists and engineers with computing access and usage suitable to their customs (simple interface, Web if possible, Shibboleth credentials), and to deploy specific cloud services devoted to Bioinformatics. Some developments are being made to adapt the StratusLab components to these specific constraints: PAT forwarding (port address translation) to satisfy the lack of public IP addresses, ”snapshot mode” to create a new virtual machine from a previous base appliance, evaluation of the integration of Shibboleth authentication procedure. Manual installation is potentially the most realistic way of deploying the StratusLab system in most bioinformatics laboratories that do not have enough system administrators. Quattor-based installation mode is also a good perspective for the sites of the French bioinformatics distributed infrastructure RENABI GRISBI who have begun to use the Quattor system these last years. The main goal of this evaluation was to fix bugs and to confirm the reliability of these procedures for the public release 1.0 of StratusLab.

3.1.4 Issues and Corrective Actions

Focus on Use Cases One of the primary recommendations from the reviewers was to refocus the WP2 activities in Y2 to really bring users onto the reference infrastructure by implementing and then publicizing their use cases. The work done for D2.3 was the first step in refocusing of the WP2 activities. Further work will need to be done to define a concrete schedule for the identified use cases and corresponds well with the evolution of the StratusLab distribution.

Delay of MS4 The milestone MS4 concerns the deployment and use of the StratusLab cloud distribution by external sites, that is, sites not associated with one of the partner institutes. There are external sites evaluating the StratusLab distribution (IRES in Strasbourg, FR and RAL in Didcot, UK) but they are not yet running these services in production. Consequently, this milestone has been delayed while we continue to support them during their testing period.

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 with the release of a new updated website, press releases and articles in relevant publications relating to the release of StratusLab version 1.0 as well as the StratusLab Marketplace. Two hands-on training workshops were held along with a large number of talks, presentations and panel discussions where the StratusLab project and software was described.

Two new Memoranda of Understanding have been signed this quarter, with CYFRONET and VENUS-C. A draft MoU with IGE has also been drawn up, and progress continues to be made in line with the existing MoUs. Collaborations with other projects have also continued.

The project's exploitation and Sustainability plan has now been developed and progress towards these goals has begun, focussing on the work of the project's commercial partners SixSq Srl and Telefnica I+D, as well as initiating contact with other commercial companies.

3.2.2 Task 3.1: Dissemination

Release Dissemination The release dissemination plan was updated for the major release 1.0. EGI.eu published a short article 'The Grid on a Cloud'[1] on the project and new release, and International Science Grid This Week (iSGTW) carried an annoucement of the release.[2] CNRS produced a press release in French 'Cloud computing et grilles informatiques : lancement de Stratuslab 1.0' was picked up and reprinted by a number of french technical media outlets,[3] with an article appearing in the French technology outlet L'Atelier BNP Paribas[4]

StratusLab v1.0 was featured on the SixSq website[5] in conjunction with a commercial support offering.

Further afield, the news of the release was carried by HPC in the Cloud[6] and HPCwire. [7]

[1] http://www.egi.eu/about/news/news_0068_StratusLabSoftwareV1.html

[2] http://www.isgtw.org/announcement/stratuslab-releases-open-source-cloud-solution-ready-grid

[3] http://www.in2p3.fr/recherche/nouvelles_scientifiques/2011/9_stratuslab.htm

[4] http://www.atelier.net/articles/cloud-aider-organismes-de-recherche-gerer-leurs-pics-de-charge

[5] http://www.sixsq.com/front-page/stratuslab

[6] http://www.hpcinthecloud.com/hpccloud/2011-07-05/stratuslab_releases_open-source_cloud_solution_designed_for_the_grid.html

[7] http://www.hpcwire.com/hpcwire/2011-07-05/stratuslab_releases_open-source_cloud_solution_designed_for_the_grid.html

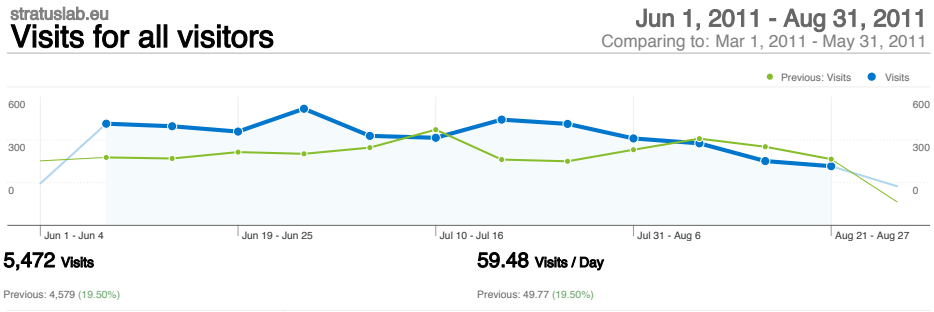

Figure 3.1: *Visits for Q5.*

Release 1.1 is due in early September and release dissemination will coincide with the EGI Technical Forum 2011.

Media & Publications StratusLab publications in the media in Q5 have been associated with the release of the production-ready StratusLab v1.0. As mentioned above, both the English and French language press releases were widely picked up and reprinted, while articles about the release were published by EGI.eu and L'Atelier BNP Paribas.

The release of the StratusLab Marketplace was highlighted and publicized as a news release by the Trinity College Dublin communications office.[8] which was picked up by the Irish technical online media outlet TechCentral.ie[9], as well as being covered in a number of Irish Education and Research-focussed websites.

Many of these media stories were also widely publicised by project members and related projects via Twitter.

A comprehensive list of media and publications is maintained on the project website.[10]

Website The redesigned website[11] went live to coincide with the release of version 1.0.

Figure 3.1 shows the number of visits to the website. The number for Q5 (5,472 visits) is up almost 20% from Q4 (4,579 visits).

The StratusLab Twitter feed now has 74 followers.

EGI Technical Forum 2011 The project booked an exhibition booth at EGI Technical Forum 2011, to be held in Lyon, France on 19–23 September 2011. The booth will feature posters, demonstrations of the software. The project is also planning dissemination material including t-shirts, brochures, and pens.

Project members are chairing the Virtualization & Cloud Computing and Agile Methodologies sessions at the Technical Forum. The project will present its status

[8] http://www.tcd.ie/Communications/news/news.php?headerID=1976&vs_date=2011-8-1

[9] http://www.techcentral.ie/article.aspx?id=17220

[10] http://www.stratuslab.eu/doku.php/press

[11] http://www.stratuslab.eu

Table 3.1: *Talks*

Title / Event	Date
Grid-Ireland Operations Centre / TCD presentation featuring StratusLab at Lero Cloud Computing Workshop Dublin, Ireland	2011-06-03
"What's the deal with Agile Contracts" Panel discussion(C. Loomis) at XP2011, Madrid, Spain	2011-06-11
"StratusLab : Les application scientifiques sur le cloud" (C. Loomis) at "Calcul intensif pour la biologie" in Lille, France	2011-06-14
"Inforum, Le Cloud en Marche (working cloud)" Stand, posters and demos at Palexpo, Geneva, Switzerland	2011-06-15
"Introduction au Cloud Computing et le projet StratusLab" (C. Loomis), seminar in Marrakech, Morocco	2011-06-17
"Building Clouds with OpenNebula 2.2 and StratusLab", tutorial at CONTRAIL Summer School 2011, Hyères-les-Palmiers, France	2011-06-27
"Challenges in Federated and Hybrid Cloud Computing", keynote at the 2011 International Conference on High Performance Computing & Simulation, Istambul, Turkey	2011-07-06
"OpenNebula and Stratuslab toolkits for virtualization", talk at CloudCamp Valencia 2011, Valencia, Spain	2011-07-13
StratusLab training day Joint European DCI Summer School, Budapest, Hungary	2011-07-15
"Challenges in Federated and Hybrid Cloud Computing", keynote at the 2011 International Conference on Parallel Computing, Ghent, Belgium	2011-08-31

and developments in the Virtualization & Cloud Computing session. StratusLab will give two 90 minute training sessions at the forum, organized by WP2.

French annual Bioinformatics Conference CNRS IBCP has presented a poster at the French annual Bioinformatics Conference JOBIM 2011, hosted by Institut Pasteur, Paris, from 28 June to 1st of July 2011. This poster, titled "Virtualisation of Bioinformatics Applications on Cloud Infrastructure", presented the virtual bioinformatics appliances that have been built by the partner CNRS and the benefits to bioinformatics scientists and engineers of using the cloud service from StratusLab. Contacts with bioinformatics researchers have been made during the discussions around the poster.

Talks Two hands-on training tutorials were organised during Q5, as well as a number of talks describing the project. A brief list is given in Table 3.1. Details and links, where available, are given on the project website[12].

[12] http://stratuslab.eu/doku.php/presentations

The project plans to submit a proposal to the EU Innovation Convention to be held on 5–6 December 2011, Brussels

3.2.3 Task 3.2: Collaboration with Standards Bodies and Related Projects

Memoranda of Understanding StratusLab has signed new Memoranda of Understanding with Cyfronet, VENUS-C.

CYFRONET StratusLab has signed a new MoU with the Academic Computer Centre CYFRONET of the University of Science and Technology in Krakow, Poland. Under the terms of the MoU, the two parties will share knowledge related to cloud and virtual technologies and StratusLab will provide a software platform and support to CYFRONET for the Polish Grid infrastructure, while CYFRONET will provide requirements to the project.

VENUS-C StratusLab has signed a new MoU with the VENUS-C project. VENUS-C will deploy OpenNebula and provide feedback and requirements to StratusLab based on their experiences. In addition to this the two projects will aim to make their software inter-operable on the interface level and to share experiences in the area of accounting. StratusLab will also investigate the use of the VENUS-C storage solutions.

EGI Under its MoU with EGI, StratusLab has contributed to the EGI prioritised list of standards by identifying the standards which the project implements or intends to implement. StratusLab is also continuing to provide input on the EGI UMD roadmap and is represented in the EGI Security Coordination Group, and the EGI training workgroup.

EMI Discussions have continued with EMI where StratusLab is in close collaboration with the EMI Virtualization group, providing comments on the virtualization-enabled architecture proposed by EMI. Other areas of collaboration include possible integration on the security level, such as adoption of Argus authorization framework, support for SAML, etc. It is planned to discuss these areas further with EMI project members at the EGI Technical Forum in Lyon in September.

Contrail Telefnica I+D StratusLab has been in contact with Contrail and hope to work together to share code and modules and push for standards for authorisation and accounting.

IGE A draft MoU has been drawn up between StratusLab and IGE which 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 In July StratusLab provided an update on the project after one year in development and operation. StratusLab also continues to provide SIENNA with

input, for example, on the recent SIENNA gap analysis document. Project partner Vangelis Floros is participating in SIENNA phone calls as member of the REB, and in the SIENNA session at the EGI Technical Forum in September.

European Cloud Computing Infrastructure Project representatives attended the European Cloud Computing Strategic Plan workshop at ESA/ESRIN from 28–29 June. This workshop aimed to develop an initial vision and direction for a European Industrial Strategy for a Scientific Cloud Computing Infrastructure to be implemented by 2020, and included research and academic stakeholders as well as commercial cloud providers. StratusLab will continue to engage with this group and will participate in a follow-up Supply-Side (cloud provider) meeting in September.

3.2.4 Task 3.3: Development of Exploitation and sustainability Plan

The initial exploitation and sustainability plan was delivered in PM10. Since then work in this area has begun following the plan set out in that document.

During Q5 SixSq Srl released the first version of its SlipStream product for automated creation virtualised system testbeds with fully integrated support for StratusLab version 1.0. This provides StratusLab with a new use case and potential new user community. Partner Telefnica I+D has also continued to develop the Claudia service manager within the StratusLab project, as part of the Telefnica I+D Global Cloud Initiative.

StratusLab partners have made contact with local business consultancy organisations as part of an effort to develop a more comprehensive business plan.

Initial contacts have also been made with some commercial companies interested in the StratusLab Distribution including Axceleon, a cloud computing consultancy firm.

3.2.5 Issues and Corrective Actions

The project review conducted in July 2011 highlighted some issues and included a number of recommendations related to this Work Package. These, along with the corrective actions to be taken are described in the Review Recommendations section of this document. WP3 will attempt to become more marketing-driven in its approach to dissemination and to focus on specific user-communities, and develop demos and dissemination materials aimed at these communities.

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

QR5 was an opportunity, over the summer, to consolidate several features released in v1.0 and start development and integration work of several new exciting features. These range from persistence storage and better understanding of Amazon EC2 interfaces in preparation for hybrid deployments, but also better integration of Claudia client tools and general configuration.

The architecture for v2.0 was also reviewed and updated and documented in D4.4.

Work also took place on consolidating our continuous integration strategy with several new tests extending the test coverage to most services.

3.3.2 Task 4.1: Definition of Reference Architecture

The architecture of StratusLab v1.0 was updated to fulfill the needs of StratusLab v2.0. The main new focus of v2.0 is support for hybrid cloud deployment. This important task is ongoing and includes inputs from all work packages and partners.

The creation of D4.4 - Reference Architecture for StratusLab Toolkit 2.0, was also an opportunity to create service specific description, which will be used on our online documentation to provide better visibility and detail on our cloud distribution.

3.3.3 Task 4.2: Integration of Open-source Distribution

Persistent Storage A new important StratusLab feature is the support for persistent storage. While the foundations for such feature were developed and integrated during year 1, this quarter say this feature used and exposed to a number of end-user feature, notably the ability to create new persistent disk, which can be attached to virtual machines via standard commands.

Integrated Improvement from Operations This work also included integration of code developed by WP5 to improve the performance of the reference infrastructure, back into the mainstream StratusLab distribution, including testing and installation support.

Development of StratusLab client tools for Claudia The development of StratusLab client tools for Claudia has started 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, obtain information about its status, and so on.

Code for RESTful Tests SixSq donated to the project a REST test framework, developed for SlipStream, such that StratusLab RESTful services can be better unit tested. Integration of this framework with services, such as the Marketplace has started, which should result in a more robust service with high test code coverage.

OpenNebula OpenNebula v3 is being actively developed, which includes StratusLab required features such as group support.

3.3.4 Task 4.3: Contextualization of Grid Services

Contextualization was improved via richer metadata creation using the StratusLab command line tools.

A deprecation feature was also defined, in collaboration with WP5 and WP3, including new commands as part of the existing command line client.

3.3.5 Task 4.4: Technical Support

Claudia Technical Support An important effort has been done during this quarter for providing technical support for Claudia. In a way, several bugs in the Claudia code have been solved and some improvement has been carried out in the generation of packages (poms, hudson jobs etc.) for version 1.0. In addition, technical support is being provided to StratusLab end-users for installation, configuration and usage of Claudia.

Technical support was provided, tightly coordinated with WP5, on a wide range of issues, such as troubleshooting virtual image contextualization and command line tools semantic.

3.3.6 Issues and Corrective Actions

A higher number of build and test job failures were recorded during the second part of the quarter. This is in part due to the holiday period were fewer people were available to quickly fix problems as they occur. This is also due to the high rate of code commits performed around the new features for persistence storage and Claudia tools. This had no impact on production. Specific tasks were created in Sprint 17 to address this issue, to return our continuous integration strategy to a nominal state.

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

WP5 started the Y2 activities having two clear goals; continue the successful operation of public cloud services by constantly integrating and improving the StatusLab distribution and improve on the areas identified during the first year project review as being weak. The recommendations of the review have been taken promptly into consideration and in the past quarter and have been the driving force for two main activities:

- Expand the scope of target use cases demonstrating the capabilities of the software to satisfy a multitude of applications.
- Improve the performance of the offered public services, mainly in the area of storage management.

3.4.2 Task 5.1: Deployment and Operation of Virtualized Grid Sites

Production Cloud Service The reference cloud service was upgraded to StratusLab v1.0 immediately after the official release of the software. The interference of the upgrade was limited since the service was already running on v0.4 of the distribution which incorporated most of the new features introduced to the software in the last months (including the switch of the hosting OS from CentOS 5.5 to Fedora 14).

Operation of a Production Grid Service The grid site was reconfigured in order to support more EGI VOs. Thus by the end of Q5 HG-07-StratusLab supports the following scientific VOs:*alice, atlas, biomed, cms, compchem, env.see-grid-sci.eu, esr, gridcc, magic, meteo.see-grid-sci.eu, nwchem.vo.hellasgrid.gr, planck, scier, seismo.see-grid-sci.eu, vo.complex-systems.ey* and *vo.dorii.eu*. The selection of the VOs to support was imposed by the policies and procedures defined by NGI-GRNET which is the hosting NGI were the grid site is located. By supporting these VOs, StratusLab's production grid site provides computing resources to scientific domains like HPC, Bioinformatics, Biomedicine, Computational Chemistry, Seismology, Meteorology, Complex Systems, Computer Science and Astrophysics.

In parallel, the total number of computing resources were doubled in size in order to better accommodate the expected increasing demand from the new VOs. The current total amount of CPU cores offered by the site are 32 hosted in 16 virtualized WNs running on the StratusLab reference cloud service.

Adaption of MapReduce platform In order to demonstrate the applicability of StratusLab to satisfy different use cases, we adapted the well-known Hadoop platform and demonstrated the ability to run MapReduce applications. MapReduce (http://en.wikipedia.org/wiki/MapReduce) is a patented software framework introduced by Google in 2004 to support distributed computing on large data sets on clusters of computers. Apache Hadoop (http://hadoop.apache.org/) is an open source implementation of MapReduce. In order to demonstrate this capability we produced a Virtual Appliance based on CentOS 5.5 that comes pre-installed and configured with *Hadoop 0.20.203* and *JDK 1.6*. This appliance is available from StratusLab's Marketplace (image ID: *Gz3B6D3z1F1m4k6RXaQA_gc0nop*). The instantiation of the Hadoop cluster takes advantage of the *stratus-run-cluster* command as documented in the project site: http://stratuslab.eu/doku.php/tutorial:mapreduce.

Following this effort we continued the tests with Hadoop by exploiting the SlipStream platform as a tool to instantiate, configure and run a sample MapReduce application. For this purpose one additional virtual appliance was produced, based on Fedora 14. This was enforced due to the requirements imposed by SlipStream on the virtual images, that have to support Python 2.6 or later. The appliance has been configured similarly to the CentOS one and is available from the marketplace with the ID: *FcQdN7cOjfueaFmb-aJ80_TqHxJ*.

Currently, we plan to demonstrate both of the above activities during the EGI Technical Forum in Lyon (Sept 2011).

Improved Contextualization The contextualization framework was improved making it more flexible (especially for automated installations) and packaging it to make it easier to install in RedHat-based distributions.

3.4.3 Task 5.2: Testing of the StratusLab Toolkit

Improving the cloud service features and performance Two activities were carried in the pre-production service; the development of an NFS-based volume management service similar to EBS (Elastic Block Storage) and the introduction of a caching mechanism in order to speed up the VM instantiation process.

The NFS-based volume management service exploits the storage space provided by NFS in order to create virtual machine disk images and attach them to running instances. The functionality has been installed on the pre-production cloud service in GRNET and was demonstrated during Sprint 16 and is currently in the process of being integrated to the release code.

The image caching mechanism has been introduced in order to tackle the latency issue that we experience during VM instantiation. This latency is introduced by the fact the each time a VM is instantiated the respective image has to be trans-

ferred from a remote repository (typical from the site in TCD). The caching mechanism utilizes the capabilities offered by QEMU and uses a combined configuration of SSH and NFS installation on the cloud site that permits copies of VM images to be stored locally, shared among all hosting nodes and used whenever a new VM is instantiated. The mechanism has been developed and demonstrated also on the pre-production site in GRNET and we plan to integrate it in the stable version of StratusLab distribution at the beginning of Q6.

3.4.4 Task 5.3: Virtual Appliances Creation and Maintenance

Marketplace In this quarter the focus of the task has been on the continued design and implementation of the Marketplace. The Marketplace has been made available for use at http://marketplace.stratuslab.eu. The Marketplace has now replaced the appliance repository as the central location for sharing of image metadata. The appliance repository is maintained to provide storage resources to the project for the actual image files. The Marketplace interface has been improved to provide easier browsing and searching of metadata entries. Work has also started on implementing a deprecation mechanism for metadata entries that is compatible with the EGI draft policy on the endorsement and operations of virtual machine images.

Existing Hudson jobs have been extended to further test the deployment, configuration, and functionality of the Marketplace. These run regularly and flag any errors that have been introduced in the code or the deployment procedures.

3.4.5 Issues and Corrective Actions

Availability and Reliability of production Grid site The average reliability and availability metrics for HG-07-StratusLab in Q5 exhibit slightly lower numbers than the foreseen target for Y2. In particular the average availability was 92% and the reliability was 91% (both targets for both values are ¿95%). The main reason for these degraded values was the major upgrade that we performed during June that required the complete re-installation of the reference cloud service physical hosts. The operation system of the hosts had to be switched from CentOS 5.5 to Fedora 14. Obviously such a radical upgraded required a few days of downtime in order to backup the images, re-format the nodes, install StratusLab and re-instantiate the VMs in the new setup. This is reflected also in the monthly statistics of June in which reliability and availability were 82% and 79% respectively. During July and August the numbers returned back to the regular high level thus increasing the overall average.

The above figures first of all underline an obvious fact that the performance and quality of a virtualized service depends heavily on the underlying cloud platform. Moreover it's a clear indication that one of the things we should improve in the operations part is the process of core software update. This process should be as smooth as possible limiting the time that the service is brought of line. For this goal we will need the combined effort from WP4 activity which has to ensure that the core components of the cloud service (VM management, storage, network etc.)

are able to migrate to new versions of the software fast and with limited manual intervention from the cloud administrators.

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

The Joint Research Activity (JRA), carried out in WP6, develops advanced technology/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 QR5 mainly has involved the resolutions of bugs for the first review demo and the continuation of some missing tasks for that. In addition, some new functionalities have been implemented as the inclusion of new policies for scaling down and the management of the physical infrastructure just in the virtual machine layer. A new networking model have been included also in OpenNebula. Finally, some work has been done in authentication for the service manager API, the support for groups and roles, the development of a new VM Template Repository and the inclusion of new drivers for accessing to more Cloud providers.

3.5.2 T6.1: Dynamic Provision of Grid Services

Scale down strategy policies Due to grid service requirements, it has been created more scale down strategy policies. Now the node to be removed should be decided by the Compute Element, which knows the status of each worker node. It has implied the implementation of this policy in the service manager and the communication of service manager - CE in order to obtain this information.

Skip the IP management from Claudia In previous versions, Claudia had the management of networks and assigned IPs to virtual machines. In a way this is not a normal working mode, since the infrastructure should be managed by the virtual machine manager. Thus, Claudia has skipped this infrastructure managements and deploy the VMs in the networks created by StratusLab (public, private, local) where the DHCP assigns the right IP.

Service Manager - CE integration During the tests with site elasticity (i.e. the dynamic addition and removal of WNs) we've realized that the Local Resource Manager (Torque) could not cope with the continuous changes of the underlying infrastructure. The problem resided in the way torque communicates with the worker nodes and in particular with the setup of the ssh (Secure Shell) service. A new unix demon had to be introduced (wnMonitor) that continuously monitors the yaim configuration file (wn-list.conf file) and whenever a new Worker Node is added in the

cluster it updates the ssh host keys in the CE. This way it is ensured that torque server will be able to communicate without problem with the workers over the ssh protocol.

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

New Networking Model OpenNebula now provides an easily adaptable and customizable network subsystem in order to better integrate with the specific network requirements of existing data centers. The default configuration connects the virtual network interface to a bridge specified in the network template. Although this is flexible enough to fit into every datacenter it lacks two essential features: network isolation and firewall management. These features are provided as hooks which need to be activated and configured in OpenNebula's main configuration file.

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

Drivers for accessing to Virtual Machine Managers Claudia accesses to the different Virtual Machine Managers (e.g. OpenNebula) by using the TCloud API. The implementation of this API has been done to work as an aggregated API, so that, it can be used to invoke different Cloud providers. In Y1, a driver for OpenNebula was created. In QR5, a driver to access to Flexiscale public Cloud has been incorporated in order to work towards an brokered federation.

Authentication in the Service Manager API In order to make possible the usage of Claudia in a public infrastructure, authentication mechanisms are being integrated in Claudia. A authentication proxy is being developed in WP2, and has been integrated in the claudia client used to access to Claudia.

Authorization Using Groups and Roles One of the main characteristics of an IaaS cloud is its multi-tenancy nature. In order to efficiently implement a multi-tenant system it is needed a flexible user system that allows the definition of user groups and access control lists to define specific access rights to each virtual resource. Therefore, the user system of OpenNebula has been extended to support groups. A user now is part of an user group, by default users in the same group can list and share (if labeled as public) any resource type (network, virtual machine, disk images etc.). The access control to each resource has been also improved with the addition of ACLs. An ACL express the user (or set of users) that may perform a given operation (e.g. create, delete or deploy) on a given virtual resource or set of them (e.g. virtual machines, networks, hosts or images).

VM Template Repository Usually IaaS clouds offer a predefined set of virtual machines (instance types) that users may instantiate. This leads to a simplification of the provisioning interface for final users that are only allowed to instantiate a predefined error-free set of virtual machines. The predefined instances may include

different OS types, packed with multiple software stacks (i.e. virtual appliances). This mechanism is even more robust in combination with the new group and user access control lists. In this way, access to a given instance type can be granted only to a specific set of users. Therefore, OpenNebula has been extended to include a new Template Repository that allows OpenNebula administrators and users to register virtual machine definitions in the system, to be instantiated later. These templates can be instantiated several times, and also shared with other users.

3.5.5 Issues and Corrective Actions

None.

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 4.1–4.5 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 4.6 contains the metrics for the project. The table groups related metrics together. The first group aimed towards dissemination show steady interest in the project; an open question is how to encourage and manage discussion with the community. The second group concerns the integration processes; all of the metrics show good progress, which is reflected in the regular releases of the distribution. The third group concerns the operations and deployments. The metrics show that the quality of the software is good. However, more effort needs to be made by the project in having external users deploy and use the distribution in production. The release of the StratusLab v1.0 release will help improve these metrics. In the fourth group, the resources provided by StratusLab are steady. Storage services have just been added in the v1.0 release, so related metrics can now be collected. The last group shows that the maintained appliances and the Marketplace are well used. Further growth in these metrics are expected in Y2.

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. Deliverables D2.3 and D4.4 have been produced in this quarter. Milestone MS4 has been delayed; see the WP2 Issues for details.

Memoranda of Understanding The project has signed Memoranda of Understanding (MoU) with the EGI and EDGI projects; an MoU with VENUS-C has been concluded but not yet signed. The project has decided not to pursue an MoU with the ERINA+ project because of the large additional effort required by Stra-

tusLab and incompatibilities with the timelines. MoUs with EMI and IGE are still under negotiation.

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. A detailed budget including both effort and spending will be provided at the end of September 2011.

Review Recommendations The results of the first periodic review were known informally at the end of the review itself. Based on that feedback, the work package managers have already adjusted the focus and work to respond to those recommendations. However, the formal feedback was not received until mid-August, leaving no time for detailed, formal responses before the end of the quarter. Initial responses have been supplied in this report (see Review Recommendations section), but those responses will be updated as the project works to analyse those recommendations and to implement the solutions.

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.
- 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.
- Finalize experiments with Grid service elasticity and move results to production
- 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

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. 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 Kick-Off Meeting	14-15/06/2010	Orsay, FR	Kick-off of project. Detailed planning for accomplishing objectives. 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: *Metrics*

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

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 website[1].

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			Yes	Delayed until PM16
D6.4	Cloud-like Management of Grid Sites 2.0 Design Report		WP6	TID	R	PU	PM17			Yes	
D5.4	Economic Analysis of Infrastructure Operations		WP5	GRNET	R	PU	PM18			Yes	
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			
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. An updated budget will be available at the end of September that will include both the new planned values for year 2 and the values for QR5.

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
OCCI	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)