Enhancing Grid Infrastructures with
Virtualization and Cloud Technologies

Project Quarterly Report

Quarterly Report QR8 (V1.1)
14 June 2012

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	2nd
Period covered	2011-06-01 to 2012-05-31
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	15 May 2012	Initial outline.
0.2	7 June 2012	Complete draft for review.
1.0	14 June 2012	Final draft.
1.1	14 June 2012	Fix error in meeting table.

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 eighth and final quarter, the project concentrated its efforts on the finalization and release of v2.0, demonstrating the functionality via scientific and commercial applications, and working towards the implementation of the defined sustainability plan. The overall achievements of the project are contained in a large number of final deliverables covering all aspects of the project.

Adoption of the StratusLab Distribution A major activity of this final quarter has been demonstrating the wide applicability of the StratusLab cloud technologies to scientific users, commercial users, and resource providers. The diverse scientific applications taking advantage of StratusLab include astrophysics, machine learning, software engineering, high energy physics, meterology, and bioinformatics. There are a number of StratusLab deployments by partner and non-partner institutes, including several commercial deployments. EGI is also evaluating the StratusLab Marketplace as a mechanism for managing and sharing virtual appliances. The broad utility of the project's software in both academic and commercial settings bodes well for continued demand after the project.

Sustainability Plan A final plan for exploitation and sustainability has been put in place which allows for the continued development and maintenance of the StratusLab Distribution past the end of the project. The plan foresees a transition to an open, community-based consortium. Steps have already been taken in this direction to ensure continued availablility of the project's software and a smooth transition from project to consortium. For example, all of the software is now hosted in GitHub and all documents are available in an open-access repository.

StratusLab Cloud Distribution v2.0 The highlight of this last quarter is the finalization and release of the second major version of the project's cloud distribution. This release consolidates the functionality enhancements and bug fixes from the past year, integrates OpenNebula 3.2, demonstrates multi-cloud scenarios, enhanced Marketplace functionality, and support for IPv6. The release supports multiple operating systems: CentOS 6.2, Fedora 16, and OpenSuSE 12.1. As for previous releases, a concerted dissemination effort will make people in the wider scientific and technical communities aware of this release.

Reference Infrastructures The reference cloud infrastructures have followed the incremental releases of the StratusLab distribution, fine tuning the services to provide the best platforms possible for the StratusLab users and providing feedback (especially on storage services) to the developers. A second infrastructure was opened to the public at CNRS/LAL to complement the existing infrastructure at GRNET. A common authentication service was put in place between the infrastructures to allow users to have access to more resources and to allow them to test cloud federation strategies.

Both CNRS/LAL and GRNET will continue to operate their cloud infrastructures as a service after the end of the project. Providing demonstration platforms for the nascent open-source consortium and ensuring that current users do not see a disruption of service.

Multi-Cloud An important feature of release v2.0 is the support for multi-Cloud scenarios. Claudia has been enhanced with a placement module allowing resources in different clouds to be used transparently in a cloud brokering scenario. Similarly, OpenNebula has been enhanced to allow cloud bursting to public clouds and federation with other StratusLab infrastructures.

Overall, the project participants have broadly achieved the goal of the StratusLab project: creating a complete, open-source distribution for an Infrastructure-as-a-Service cloud. The relevance of the provided services has been shown via the production grid site running over a StratusLab cloud and by the wide variety of scientific and commercial applications using StratusLab. This forms a solid foundation for the transition from a project to an open-source community consortium.

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.

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 1.1: *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

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

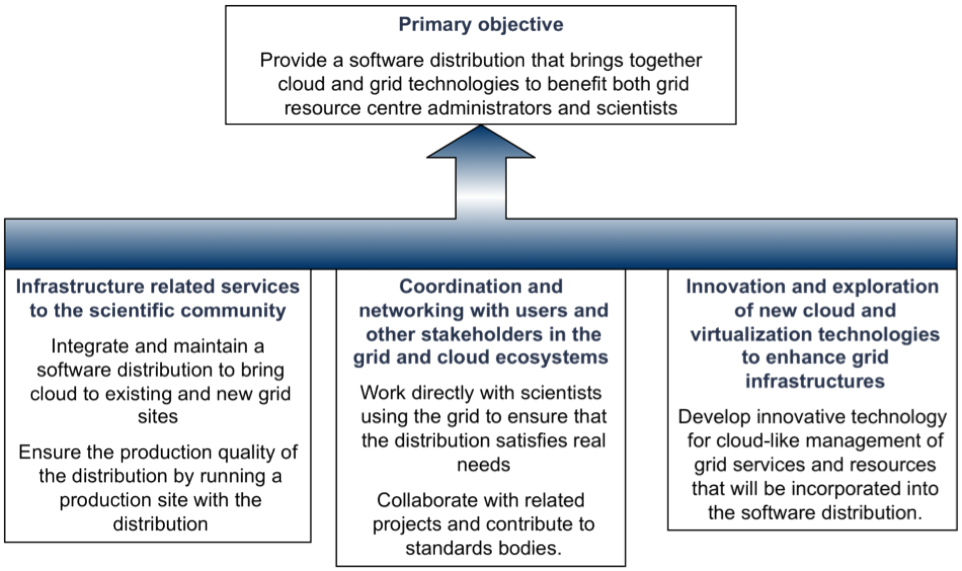

Figure 2.1: *Primary and supporting objectives*

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

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

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

2.2.4 Quarter 8

- Release of v2.0 of the StratusLab distribution for multiple operating systems
- Detailed definition and initial implementation of sustainability plan
- Demonstrate of hybrid cloud computing
- Measurement of the behavior of the system in terms of scalability and IO performance, including with alternate configurations (e.g. GPFS)
- Demonstrate a pre-production elastic Grid site
- Complete migration of gLite production Grid site to UMD-1
- Finalise sustainability plans to ensure continued development and maintenance of project outputs

2.3 Review Recommendations

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

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

   (See next response.)

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

   These delays were largely related to the difficulty of hiring new personnel in the summer months and to lengthy administrative procedures. Most of the partners compensated to some extent with effort from permanent staff. Currently, all of the partners are fully staffed and no further perturbations are expected. In general, the overall deviations from the planned effort was around 10%, although because of larger than expected travel expenses, more expensive personnel, etc., the budgets are roughly in line with the planned budgets.

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

   The project agrees that this is a reasonable shift in focus for the dissemination work in year 2. The project proposed a demo for the EU Innovation Convention, which unfortunately was not accepted. A general video highlighting the benefits of cloud technologies and showing concrete examples of use was prepared. Although still technical, it provides a better introduction to the cloud and its benefits then previous material.

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 made slight modifications to the metrics (mainly for the Marketplace) and updated some of the second year targets.

   In particular for dissemination, the target for people on the project announcement list was raised to 75 (with 73 currently enrolled). The goal for the discussion site was set at 100 where only 16 are enrolled, highlighting a problem we have had in jump starting a community around the StratusLab distribution.

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. As mentioned above, a general video was created that describes the project, expected benefits from cloud technologies, and gives concrete examples of real use of the StratusLab cloud. This is a start for better demonstrating the cloud benefits. Deliverable D2.4 provides further examples, which will be used as a basis for continued dissemination of the Stratuslab consortium's achievements after the project.

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

   Deliverable D4.4 that provides the reference architecture for StratusLab 2.0 was structured to provide a separate description for each service. That description includes details on where the component was developed, external

dependencies, and interactions with other components. This should provide a clearer overview of the development activities of the project. In addition, more effort has been made in the deliverables to describe clearly the project's work with respect to individual components, especially in the WP4 and WP6 deliverables.

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

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

   Although obtaining quarterly updates of the financial information turned out not to be possible, all of the partners have made better efforts in the preparation of the second year periodic report and and no delays are expected in the delivery of the final periodic report.

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.*

   Getting complete budget information from the partner's administrations quarterly has proven to be extremely difficult. Nonetheless, all of the partners were aware of their underspending and made a concerted effort to correct the problem in the second year. Preliminary figures indicate that all partners have invested more effort in the second year and have significantly closed the gap between the actual and planned budgets. See the financial section in the periodic report for details.

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. This formed the basis for the persistent disk service for the 1.x series of releases.

   Nonetheless, use of shared file systems provide a convenient alternative and provides an easy entry configuration. A shared file system option was put back into the persistent disk service for the 2.0 release. The performance of various shared file systems was compared, including GPFS, GlusterFS, and Ceph. See the D5.5 deliverable for details.

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

    In the second year of the project, a more systematic approach to testing and validation of the StratusLab software was put into place. This included a certification testbed to validate all candidate releases; tests on this infrastructure were largely automated through Hudson. Following a successful validation, a separate pre-production infrastructure was upgraded to further validate the release. Together, complete installations of new systems and upgrades of existing systems were tested. These tests extended the daily automated installation and test of the development release of the software.

    Relevant application-oriented benchmarks have existed since near the beginning of the project. Unfortunately, these were never used to collect systematic metrics on the performance of the system. Instead effort focused on performance problems noted by users and system administrators. These included the caching performance of the persistent disk services and I/O bandwidth via the network and to local disk. Both of these have been improved in the v2.0 release.

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

    In the second year, there was a stronger emphasis on pure cloud services rather than on grid services over cloud. The latest release of the StratusLab cloud distribution has support for OCCI, TCloud and Deltacloud.

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

    Our feeling is that StratusLab should remain a "neutral" carrier that allows users and administrators to implement their own mechanisms for ensuring privacy and confidentiality. Nonetheless, we see areas in which the services can improve, peoviding, for instance, better logging and auditing information that can complement user-level strategies for privacy and confidentiality. Some minor improvements were made in this area.

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 were partially implemented over the course of the second year and used to evaluate the functionality of the StratusLab cloud distribution. The deliverable D2.4 provides information on the scientific and commercial adoption of StratusLab. Many of those were used as concrete examples in the general video that describes the benefits of cloud technology. These examples form a good basis for "marketing" of the StratusLab distribution and more needs to be done to capitalize on these successes.

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

    The project agrees that having a "user champion" within the project is a good idea. However, the PMB is not the proper body as it meets infrequently and isn't in direct contact with the daily technical advances of the project. In the end, a viable solution for incorporating an outside representative has not been found. Instead we have relied on input from the scientists, engineers, and administrators with whom the project already has contact. In addition, commercial contacts through the partners have proven invaluable for validating the project's software and indicating new directions for improvement.

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.

Recently several commercial ventures have shown interest in the StratusLab software and have created demonstrators with it. This is a significant validation of the StratusLab concept for use within the commercial sector.

SixSq has done an extraordinary job in promoting StratusLab in the commercial sector. This work includes using StratusLab as a platform for SixSq's own SlipStream product, a commercial turnkey cloud solution targeted at SMEs in collaboration with IBM and Darest, a software engineering platform for automated deployment of SCOS-2000 with ESA/ESOC, and a large StratusLab deployment by Atos within the Helix Nebula initiative.

Based on the outcomes of these initiatives, StratusLab has a clear niche as an efficient, easy-to-install private cloud solution for small to medium-sized infrastructures. However, the work with Atos also shows StratusLab's potential on large infrastructures as well.

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 specific use cases and deliverable D2.4 contains concrete success stories from a wide range of scientific and engineering domains. 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.*

    We have made an effort with the second year deliverables to be more concise and provide better summaries.

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. The partners have all been very responsive concerning financial information for the second year. No delays are expected for the second periodic report or the for second year deliverables.

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

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

3 Progress and Achievements

In the eighth and final quarter, the project concentrated its efforts on the finalization and release of v2.0, demonstrating the functionality via scientific and commercial applications, and working towards the implementation of the defined sustainability plan. The overall achievements of the project are contained in a large number of final deliverables covering all aspects of the project.

Adoption of the StratusLab Distribution A major activity of this final quarter has been demonstrating the wide applicability of the StratusLab cloud technologies to scientific users, commercial users, and resource providers. The diverse scientific applications taking advantage of StratusLab include astrophysics, machine learning, software engineering, high energy physics, meterology, and bioinformatics. There are a number of StratusLab deployments by partner and non-partner institutes, including several commercial deployments. EGI is also evaluating the StratusLab Marketplace as a mechanism for managing and sharing virtual appliances. The broad utility of the project's software in both academic and commercial settings bodes well for continued demand after the project.

Sustainability Plan A final plan for exploitation and sustainability has been put in place which allows for the continued development and maintenance of the StratusLab Distribution past the end of the project. The plan foresees a transition to an open, community-based consortium. Steps have already been taken in this direction to ensure continued availablility of the project's software and a smooth transition from project to consortium. For example, all of the software is now hosted in GitHub and all documents are available in an open-access repository.

StratusLab Cloud Distribution v2.0 The highlight of this last quarter is the finalization and release of the second major version of the project's cloud distribution. This release consolidates the functionality enhancements and bug fixes from the past year, integrates OpenNebula 3.2, demonstrates multi-cloud scenarios, enhanced Marketplace functionality, and support for IPv6. The release supports multiple operating systems: CentOS 6.2, Fedora 16, and OpenSuSE 12.1. As for previous releases, a concerted dissemination effort will make people in the wider scientific and technical communities aware of this release.

Reference Infrastructures The reference cloud infrastructures have followed the incremental releases of the StratusLab distribution, fine tuning the services to provide the best platforms possible for the StratusLab users and providing feed-

back (especially on storage services) to the developers. A second infrastructure was opened to the public at CNRS/LAL to complement the existing infrastructure at GRNET. A common authentication service was put in place between the infrastructures to allow users to have access to more resources and to allow them to test cloud federation strategies.

Both CNRS/LAL and GRNET will continue to operate their cloud infrastructures as a service after the end of the project. Providing demonstration platforms for the nascent open-source consortium and ensuring that current users do not see a disruption of service.

Multi-Cloud An important feature of release v2.0 is the support for multi-Cloud scenarios. Claudia has been enhanced with a placement module allowing resources in different clouds to be used transparently in a cloud brokering scenario. Similarly, OpenNebula has been enhanced to allow cloud bursting to public clouds and federation with other StratusLab infrastructures.

Overall, the project participants have broadly achieved the goal of the StratusLab project: creating a complete, open-source distribution for an Infrastructure-as-a-Service cloud. The relevance of the provided services has been shown via the production grid site running over a StratusLab cloud and by the wide variety of scientific and commercial applications using StratusLab. This forms a solid foundation for the transition from a project to an open-source community consortium.

3.1 WP2: Interaction with Targeted Communities

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

3.1.1 Summary

A major activity of this final quarter has been demonstrating the wide applicability of the StratusLab cloud technologies to scientific users, commercial users, and resource providers. The diverse scientific applications taking advantage of StratusLab include astrophysics, machine learning, software engineering, high energy physics, meterology, and bioinformatics. People from these areas were helped with direct support and training. There are a number of StratusLab deployments by partner and non-partner institutes, including several commercial deployments. Details are provided in the D2.4 deliverable. A companion deliverable (D2.5) provides an evaluation of the StratusLab distribution against defined requirements and feedback from users and administrators. The broad utility of the project's software bodes well for continued use as the partners transition to an open, community-based consortium.

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

Tutorials WP2 to update the user tutorial to reflect the most recent releases. A tutorial was given at the EGI Community Forum in Munich, Germany with an attendance of around 25 people. The dissemination section lists this tutorial and provides a link to the agenda and material. This tutorial was used to give a "private" tutorial to four people from the Astroparticle and Cosmology (APC) laboratory in Paris, France. The results of their initial use of the StratusLab cloud are described in the deliverable D2.4.

Deliverable D2.4 The Deliverable D2.4 documents the adoption of the StratusLab distribution by end-users and system administrators. This provides a summary of our work with people in these two groups over the project's lifetime. It also provides retrospective of our activities with the users and system administrators and draws some general conclusions on the successful and less successful aspects of our work.

AppStat Researchers from the Machine Learning and Applied Statistics Group have started to use the StratusLab cloud infrastructure for their research into machine learning algorithms. The project support personnel from them port an initial application to the cloud and to use the resources for an initial scientific analysis.

This activity is fully reported in the deliverable D2.4.

Updated Base Images The previous work done to automate the production of base machine images (i.e. appliances with minimal distributions of common operating systems) has allowed this images to be updated frequently with operating system patches and when problems were found with the images. These changes include increasing the disk space available in the images, ensuring that hostnames are not embedded in the images, and enabling IPv6.

Project Video A new video has been produced that gives an overview of the project, its software, and examples of use. This eight-minute video is available from YouTube and linked from the project website. This serves as a good introduction to more detailed videos that will be produced for the review concerning bioinformatics and the n-tier application prototype.

Enhanced in the E-business Application Use Case The e-business application used for demonstrating advanced StratusLab functionalities developed in WP6 (multi-tier service management, scalability and federation) has been evolved providing its database stored in the persistent disk and a GUI.

Update of the Bioinformatics Service Use case In the context of Bioinformatics Web Service use case, CNRS IBCP have added new bioinformatics software to the 'biocompute' appliance. For example, new tools about multiple alignment of sequences (gene and proteins) like Muscle and Clustal Omega have been installed and configured. This appliances are still fully contextualized to be used with the 'biodata' appliance and the persistent disk feature.

Update of the TOSCANI Use Case A new version (2.0) of the ARIA2.3 appliance following the IaaS model has been developed and released in collaboration with the group of Dr. M. Nilges at the Institut Pasteur Paris. With the usual cloud interfaces (command-line client and Web interface on IBCP's infrastructure) scientists can launch multiple virtual machines that will be used for a compute-intensive analysis. They connect to the virtual machine that will be used as the ARIA master and run the script 'aria-clouder' we have developed and installed in the ARIA2.3 appliance. This tool finishes the configuration of the virtual cluster, setting up the system services and the ARIA configuration file to adapt to the virtual infrastructure. Scientists have then simply to run the ARIA tool as usual and retrieve their results once the computations finish. This appliances is also compatible with all the instance types available on IBCP's cloud (large, bigmem, HTC) and with the StratusLab reference infrastructures at GRNET and LAL.

Online Bioinformatics Documentation CNRS IBCP has updated the online bioinformatics documentation[1]. The scientific context is described to motivate the development we are doing in the project. The two bioinformatics use cases are described in detail there. Similarly, the three bioinformatics appliances are also described, as a complement to the Marketplace descriptions.

[1] http://www.stratuslab.eu/doku.php/bioinfo

Usage of the IBCP's Cloud Infrastructure by the Bioinformatics Community The IBCP infrastructure has been opened other colleagues from the bioinformatics community. New accounts have been created and we have run experiments with real data for protein modeling. New contacts have been initiated with scientists working in the genomics area. This action will bring feedback from scientists to the project developers.

Bioinformatics Web Interface for StratusLab Cloud CNRS/IBCP has added new features to its bioinformatics web interface to the cloud. Scientists can now manage their persistent disks through the web interface. They can create as many disks as they need, within their defined quota. There also new features for cloud administrators to help manage the all users' disks easily.

The two bioinformatics use cases and the bioinformatics web interface for StratusLab cloud have been presented at the "Clouds: Users" session of the EGI Community Forum in Garching on March 2012.

3.1.3 Task 2.2: Intensive Evaluation of StratusLab Products

Multi-Platform Support The jobs for multi-platform builds have been further refined to provide a complete build chain for Fedora 16, CentOS 6.2, and OpenSuSE 12.1. The Fedora 14 builds are also being maintained until the full set of test jobs are migrated to CentOS 6.2.

Use Case Tests The use case tests continue to be updated with new releases of the software. They are often failing because of changes to the underlying software or because of deployment problems with the services. They remain a good indicator of "releasablility" of the current development version.

Deliverable D2.5 The Deliverable D2.5 summarizes the work done to evaluate the StratusLab software and services over the course of the project. Like the first version of this deliverable, this document compares the functionality provided to initial requirements from several sources.

StratusLab operations in a Bioinformatics Lab CNRS/IBCP has continued its adaptation of the StratusLab distribution in a bioinformatics context. The PAT (Port Address Translation) implementation has been integrated with the new releases of the StratusLab distribution. New computing resources have been added to the IBCP infrastructure, giving it a total of 172 cores, 536 GB of RAM and 25 TB of storage through the persistent disk service. A new instance devoted to "high throughput computing" with 48 CPUs and 120 GB RAM has been defined to satisfy requirements for some bioinformatics tools.

3.1.4 Issues and Corrective Actions

No particular issues have arisen in this quarter. However, a point will have to be made to disseminate the successful use of the StratusLab cloud infrastructures for a wide range of applications during the v2.0 dissemination activities.

3.2 WP3: Dissemination

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

3.2.1 Summary

Dissemination efforts have continued this quarter even as effort has turned towards sustainability and other activities in preparation for the end of the project. The release announcement for version 1.4 was widely circulated and there has been some publicity surrounding commercial exploitation efforts of the partners.

Website visits are up over the previous quarter and use of social media and mailing lists has also increased. The use of the StratusLab user forum, however, is still low.

Most formal collaborations will draw to a close at the end of the project but informal collaborations that make use of the results of the StratusLab project are expected to continue, driven by individual partners, or by the open-source StratusLab community.

A final plan for exploitation and sustainability has been put in place which allows for the continued development and maintenance of the StratusLab Distribution and other outputs of the project beyond the project lifetime.

3.2.2 Task 3.1: Dissemination

Release Dissemination StratusLab released version 1.4 during this quarter. A release dissemination plan was produced for this release. Notification of the release was disseminated via the website and the project's announcement email list. The release announcements for both releases were also sent to various media outlets and were carried by International Science Grid This Week (iSGTW) and HPC In the Cloud.

Media & Publications A press release was written about the release of the DS-Cloud Ready Pack by partner SixSq in collaboration with Swiss IT Solutions Provider Darest. This was widely disseminated and picked up by several online media outlets.

The release announcements and other media stories, events and general project news were also disseminated via Twitter, with several items being retweeted by related projects.

A major dissemination effort is expected just beyond the end of the project to coincide with the final release of StratusLab v2.0.

Website, Mailing Lists and Fora Figure 3.1 shows the number of visits to the website. The number for Q8 (5,372 visits) is up from Q7 (4,442 visits).

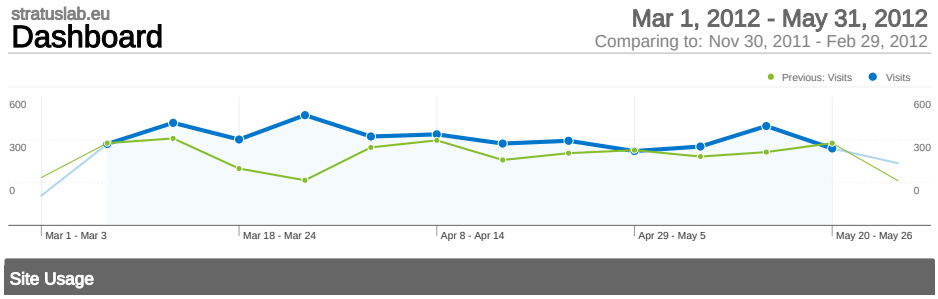

Figure 3.1: *Visits for Q7.*

The StratusLab Twitter feed now has 101 followers.

The announcements mailing list, which allows interested members of the public to keep up-to-date with news from the StratusLab project such as new releases, has 73 members.

The user announcements mailing list for users of the StratusLab reference infrastructure has 47 members, all of whom are users of the reference infrastructure.

The public user forum has 16 members, although there has again been relatively little use of this forum, despite attempts to disseminate the url more widely.

Booth at EGI Community Forum The project prepared a booth for the EGI Community Forum event in Munich, Germany. Four posters were prepared covering the StratusLab service, applications using StratusLab, the trusted appliance "ecosystem" and our use of agile software development methodologies. A large plasma screen provided a continuous loop of videos from the project. In addition, an updated project brochure, a take-home flyer of the posters, and copies of the StratusLab overview paper were provided. The entire run of 100 StratusLab t-shirts were given away during the event.

Talks A brief list of talks featuring StratusLab given by project members is given in Table 3.1.

3.2.3 Task 3.2: Collaboration with Standards Bodies and Related Projects

ERINA+ StratusLab has attempted complete the ERINA+ Socio-Economic Impact Self-Assessment of e-infrastructure Projects using the online self-assessment tool provided by ERINA+. Initial results of the assessment are relatively favourable, but updates to the tool and some technical issues have impacted our ability to complete the assessment process.

StratusLab also participated in the ERINA+ session at the EGI Community Forum in March.

VENUS-C Support to Develop OVF4ONE This software (http://grids15.eng.it/svn/main/ovf4one) is a Java implementation of OCCI according to GFD.184

Table 3.1: *Talks*

Title / Event	Date
"Cloud, Ready for Bioinformatics?" at the EGI Community Forum, Munich, Germany	2012-03-28
"Supporting grid-enabled GPU workloads using rCUDA and StratusLab" at the EGI Community Forum, Munich, Germany	2012-03-28
"StratusLab: Use cases, features and sustainability" at the EGI Community Forum, Munich, Germany	2012-03-29
"StratusLab Training" at the EGI Community Forum, Munich, Germany	2012-03-29
"Operating a Public Cloud with StratusLab" at the EGI Community Forum, Munich, Germany	2012-03-29
"Public vs Private Cloud Usage Costs: The StratusLab Case". 2nd International Workshop on Cloud Computing Platforms (CloudCP) - In conjunction with EuroSys 2012, Bern, Switzerland	2012-04-10
"The StratusLab distribution: Use-cases and support for scientific applications" at Earth & Space Science Information session 2.10, EGU General Assembly 2012, Vienna, Austria (Invited Talk)	2012-04-26
"An open source cloud solution based on KVM" Pierre Perdeams (IBM) / Marc-Elian Begin (SixSq) at the IBM Technical University, Birmingham, UK	2012-05-10

using OVF messages and OpenNebula as backend. ovf4one translates RESTful calls to the OCCI methods into OCA RESTful calls, and the OVF XML message is translated into OpenNebula VM templates. ovf4one is part of the OpenNebula software ecosystem[2].

VENUS-C Support, Feedback and Experience Sharing Members of VENUS-C have received support about OpenNebula (contextualization process, use of public Cloud APIs, Cloud federation..). VENUS-C has provided feedback to StratusLab about OpenNebula, including feature requests and bugs, and the StratusLab IaaS model. Both projects have shared experiences installing and deploying OpenNebula.

Support and Integration for Spanish NGI During this period StratusLab has assisted the Spanish-Portuguese NGI to set up several cloud sites to build a federated cloud infrastructure. Along with StratusLab components, such as OpenNebula and the Marketplace, other IaaS platforms are used (e.g. OpenStack). The UCM team also helped with the integration of the authentication system with the LDAP services of the sites by developing specific multi-site authentication drivers.

EGI Collaboration with EGI has continued, including delivery of a StratusLab tutorial at the EGI Community Forum 2012 and participation in the EGI Federated Cloud Task Force, where the StratusLab Marketplace is in use.

The project is also providing input to EGI for a proposed session on business models at the EGI Technical Forum which will take place in September 2012 in Prague, Czech Republic. StratusLab will provide examples of potential and actual business models for organisations using the StratusLab Distribution and other outputs of the project. These examples will draw upon the real commercial products and use cases which have been developed by the StratusLab project partners.

SIENA The SIENA project StratusLab submitted a response to the SIENA *call for actions* outlining how the project has and will continue to contribute to a range of initiatives, including:

- collection of deployoment scenarios
- collaboration and standards
- contributing to the European Digital Market
- cloud federation
- developing and identifying business models
- sustainability and reuse of project outputs beyond the project
- inventory of grid and cloud standards for research

[2] http://www.opennebula.org/software:ecosystem:ovf4one

3.2.4 Task 3.3: Development of Exploitation and Sustainability Plan

Initial Sustainability Actions As part of the sustainability plan currently running services have to be moved to locations that provide the possibility of long-term support. Similarly documentation and materials need to be archived in public locations. Some of these actions are already underway.

- The code repositories have been moved from a dedicated server to GitHub, allowing current members to access the repositories after the end of the project and other people in the community to contribute.
- All project deliverables and milestones have been unloaded to an open access repository (HAL). This will ensure that these documents are visible and available for the long-term.

Other actions and long-term development of a community around StratusLab will occur after the project's close.

The project has written its "Exploitation and Sustainability Final Plan" (D3.5) which outlines in more detail the exploitation outcomes and plans for sustainability beyond the project life time. D3.5 is due concurrently with this report.

3.2.5 Issues and Corrective Actions

It was hoped that the ERINA+ Socio-Economic Impact Self-Assessment could be fully completed before the end of the project, however, as the self-assessment tool is still in development, the results of the assessment are not yet available to be included in the StratusLab final deliverables.

Usage of the StratusLab User Forum is still low despite efforts to increase the visibility of the forum. As the StratusLab Distribution will become an open-source project going forward the collaboration will continue to publicise the user forum and encourage its use by the user and developer community.

It is worth mentioning that a final dissemination effort is expected with the imminent release of StratusLab v2.0, although the results of this are not reported here.

3.3 WP4: Software Integration and Distribution

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

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

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

3.3.1 Summary

During Q8, the integration and testing effort continued in order to deliver a solid v2.0 release of the entire StratusLab distribution. Several services were upgraded. We also completed the integration of OpenNebula 3.2.

The build and test system was also refactored to simplify it, in order to add support for release on CentOS 6.2, Fedora 16 and OpenSuSE.

Four sprints were completed during this period, resulting in the release of StratusLab v1.4 and v2.0 beta and v2.0.

3.3.2 Task 4.1: Definition of Reference Architecture

The architecture of StratusLab is stable and no significant modifications were required.

3.3.3 Task 4.2: Integration of Open-source Distribution

Placement Decision Module The placement decision module, developed in WP6, which is required for federation, has been included as part of the Claudia distribution in StratusLab.

The Persistent Storage Service was upgraded with a new backend architecture built in Python. This provides a cleaner separation of concerns between the Java Restlet service implementation, responsible for disk state, from the logic manipulating the backend storage system. The result of this separation of concerns also includes a cleaner integration with OpenNebula drivers. The Persistent Storage Service now supports iSCSI TGTD (as before) and iSCSI NetApp. Further, a File based backed is also available, and used at GRNET, making the Persistent Storage Service able to work with distributed file systems, as well as iSCSI solutions. We also have prototyped an iSCSI and Fiber Channel IBM V7000 implementation, although this prototype is not production ready.

OpenNebula 3.2 is now integrated with StratusLab. Several backward incompatible changes were detected during the integration, but only after the release of 3.2, which forced us to adapt and spend more effort than expected.

Version 2.0 of StratusLab is now being released under CentOS 6.2 as the default platform. This will help enterprise adoption of the StratusLab distribution, since

compatible with Redhat Enterprise Linux 6.2 (RHEL 6.2).

3.3.4 Task 4.3: Contextualization of Grid Services

An initial discussion took place around Cloud-init, proposed by Ubuntu, and endorsed by Amazon, which proposes a first convention for contextualisation. We have not had the chance to converge on this interesting development, but this discussion will continue, with a potential implementation, after the project's conclusion.

The promising element of cloud-init is that it can be implemented on top of the current StratusLab contextualisation strategy, confirming that our initial assumptions and strategy were right.

3.3.5 Task 4.4: Technical Support

Support for Integrating OpenNebula 3.2 A new set of configuration files has been created for StratusLab installer. Code has been adapted to make OpenNebula 3.2 behave more similarly to 3.0. Patches previously done for 3.0 have been translated to work with 3.2. Finally, the CLI has been also changed to comply with the new XML-RPC API.

New Hooks for User Management New hook functionality has been added to execute scripts on user creation or modification. This was needed for the stratuslab-one-proxy and the SunStone web interface integration.

Integration of the SunStone Web Interface The SunStone web interface for OpenNebula has been integrated in StratusLab, to graphically show the deployment of VMs in demonstrations.

Preparation of a Scenario for Cloud Bursting An scenario has been prepared where OpenNebula provides access to Amazon EC2, using its adapter. In OpenNebula, a public Cloud service like this is managed as any other host, but it may provide "infinite" capacity for the execution of VMs. Uploading VM images on-the-fly to the external cloud is still not feasible (because the Amazon S3 datastore is not ready yet), therefore it is assumed that a valid VM image has been already uploaded to EC2. And, since VMs deployed on different clouds have to communicate through the Internet, OpenVPN is used to establish a secure channel between them.

Preparation of a Scenario for Cloud Federation Another scenario has been prepared where OpenNebula provides access to other StratusLab sites running a different OpenNebula instance. Other StratusLab sites are accessed using the ONE2ONE adapter. For storage, the sites just use the same Markeplace service. For networking, the solution based on OpenVPN is applied, as in the Cloud bursting scenario.

Support was also provided for a number of deployments, including at GRNET and LAL, but also of evaluation infrastructures in Vietnam and for the Helix Nebula collaboration.

3.3.6 Issues and Corrective Actions

No serious issue encountered during this quarter.

Backward incompatible changes were detected with OpenNebula v3.4. The changes relate to the fact that image registration with the OpenNebula image repository is now made mandatory. This means that StratusLab is not currently compatible with the latest OpenNebula version. A possible way forward would be to build a rich facade to the OpenNebula Daemon, hiding some of these incompatible changes. However, this was not possible within the effort and time constraints left in the project. This is an issue that will have to be carefully handled as we transition towards the community driven open source project.

3.4 WP5: Infrastructure Operation

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

3.4.1 Summary

The focus of work for WP5 during Q8 evolved naturally around the fine tuning of services and the resolution of outstanding issues in order to deliver the final version of services and the overall activity results. Among the highlights of this last quarter is the introduction of common authentication service between LAL and GRNET sites and the testing and deployment of file-based persistent disk service.

3.4.2 Task 5.1: Deployment and Operation of Virtualized Grid Sites

GRNET Reference Cloud The site continued to operate during Q8 using the iSCSI-based pDisk service running on the frontend node. After the upgrade to StratusLab 1.4 the behavior of the service is much more stable and the only real issue is the lack of storage space in order to deploy a large number of VMs. With 1 TB local storage the maximum capacity of the site is around 30-40 VMs depending on the volume of cached images. This restrictions have motivated the development of the file-based pDisk version which is planned to be included in StratusLab 2.0. As of the writing of this report, we are currently testing this version and we plan to deploy it in the production site once certified.

LAL Reference Cloud The cloud infrastructure at LAL has been opened up to external users. This serves as an alternate site to the GRNET infrastructure and allows for testing of different storage solutions. (LAL has both storage that can be accessed via an iSCSI target like tgtd and a commercial NetApp solution.)

Unified Authentication The two reference cloud services now use a single LDAP server for authentication. This allows users to access both sites using the same credentials. This LDAP server is also connected to a simple registration service that allows users to register for the StratusLab cloud and manage their own password and certificate information.

Hosted grid site HG-07-StratusLab continued to run on an an auxiliary cloud site in order to mitigate the restrictions of iSCSI-based pDisk. The site will also be moved to the production cloud service in GRNET once the upgrade to v2.0 is completed and the site starts exploiting the new file-based pDisk service.

3.4.3 Task 5.2: Testing of the StratusLab Toolkit

NetApp Support Integration of existing resources at a site is important for the adoption of the StratusLab cloud distribution. This is especially important for storage as sites are very likely to have their own storage infrastructure already in place. LAL has worked to integrate its NetApp server into the persistent disk service. This has resulted in a refactoring of the persistent disk "backends" to make the integration of new services more flexible. This work was also necessary to allow shared file systems to be used with the persistent disk service again.

IPv6 Support With the exhaustion of the IPv4 address space, support for IPv6 is becoming more important. The StratusLab services have been tested in a mixed IPv4 and IPv6 environment. All of the java-based services (the majority) worked without any changes in this environment; those services were accessible to both IPv4 and IPv6 clients. OpenNebula cannot at the moment provide both IPv4 and IPv6 addresses to clients. This is not really a limitation in the StratusLab configuration as IPv4 can be used in OpenNebula to assign addresses and the site DHCP server can be configured to provide both IPv4 and IPv6 addresses to the associated MAC address. A slight change was required in the StratusLab client to provide the users with domain names rather than raw IPv4 addresses. It was also required to update the standard base images to configure both IPv4 and IPv6 networking.

File-based pDisk service As mentioned above, the restrictions imposed by the LVM/iSCSI based pDisk service has led to the development of a file-based approach which can exploit any local or network file service for hosting and sharing VM instance images and storage volumes. The "file-based" term refers to the fact that VM images are simple files and not logical volume devices as in the iSCSI implementation. This approach not only enables the utilization of a large range of file systems (NFS, GPFS, Ceph, etc) but maintains the original file format for VM instances which in turn is very useful for features like instance live-migration, VM backup etc. This version of pDisk is currently under certification and will be released officially as part of StratusLab v2.0 by the end of the project.

3.4.4 Task 5.3: Virtual Appliances Creation and Maintenance

Marketplace During this quarter development of the Marketplace has continued. The focus has been on bug fixing and interface improvements.

Marketplace availability is critical to the StratusLab infrastructure, and as such a replication mechanism is required to provide fault-tolerance in case of failure. A simple scheme has been implemented that will allow for a replica Marketplace to be quickly and easily deployed using standard tools (rsync, and cron). A replica Marketplace was deployed at TCD to demonstrate this.

A number of fixes were made for the email verification, which is now in use at IBCP for their Bio-Informatics Marketplace. Also, some improvements have been made to the search interface to make it easier for users to filter entries. Search performance has been improved through changes to the internal SPARQL queries.

3.4.5 Issues and Corrective Actions

Security incident in SlipStream nodes One of the physical nodes in GRNET that hosts a production instance of SlipStream was hacked and used to launch further attacks to VNC servers running in other hosts. The incident was brought to the attention of the WP5 operations team by GRNET's Network Operations Center. A forensic analysis in the node revealed a rootkit installed under /tmp/.ICE-unix/.run containing all the necessary scripts to take control of the local system and initiate the attacks to VNC ports of external systems. The security break was tracked back to host with IP 85.183.72.18 located in Germany. After the forensics completed, the node was re-imaged and returned to SixSq in order to re-install SlipStream. GRNET's NOC was informed regarding the above findings for further actions.

Grid site elasticity The ability to deploy elastic grid sites that autonomously change their size using Claudia's KPI capabilities has been demonstrated in previous quarters. Nevertheless, this features was never moved to production usage. This is because we finally decided to give priority to other more important issues like the testing and feedback of the file-based pdisk services as well as the optimization of the existing static grid site and the migration to UMD-1.

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

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

3.5.1 Summary

The work done in Q8 has involved working on Inter-Cloud Connectors for different scenarios. On one side a decision placement module has been created and integrated in Claudia for deciding in which site to deploy the service, providing Cloud brokering. On the other hand, Amazon EC2 and ONE2ONE drivers have been improved for having Cloud bursting and federation, respectively. In addition, Claudia has been evolved to include the persistent disk interaction. Finally, OpenNebula has been improved with support for multiple Datastores, new transfer drivers, or Cloud partitioning.

3.5.2 T6.1: Dynamic Provision of Grid Services

Placement decision module During this quarter a placement module has been created. This module is in charge of deciding the site where the service is going to be deployed. Thus, according a set of requirements specified in the service definition in the OVF, the placement module returns the site and its management URL. It is used for Cloud brokering.

Cloud Brokering with Claudia Cloud service brokering is a form of cloud service intermediation, in which an entity adds value to one or more cloud services on behalf of one or more consumers of that service. In StratusLab, Claudia is taking care of this role and is performing this functionality by using the placement decision module.

Integration of Claudia with OpenNebula for Hybrid Cloud Computing OpenNebula has been enhanced to provide the hybrid Cloud functionality. Claudia has been adapted to this functionality changing the way it obtains the IP where the virtual machine is deployed in a remote site.

Enhanced Claudia Works with Persistent Disk Service Claudia code has been enhanced to include the persistent disk functionality. References in the OVF it are translated to the appropriate fields in the OpenNebula machine template.

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

Support for Multiple Datastores The storage capabilities of OpenNebula have been improved in StratusLab with the addition of the Datastore abstraction. A datastore, previously known as an Image Repository, is an abstraction of any storage medium for VM disk images. Datastores are distributed to the hosts with specific transfer drivers. This allows a single host to include multiple datastores of different types. OpenNebula 3.4 includes four datastore types (system, file-system, iSCSI/LVM and VMware). The system has been architected to be highly modular, so these base types can be easily adapted to any specific deployment.

New Transfer Drivers Hosts are not tied to a single transfer mechanism (transfer driver) and now can access images from different datastores in different ways. Even, a VM can have its disks in different datastores. Also the transfers associated with persistent or `save_as` images have been simplified. There are also new drivers to use in combination with the datastores: qcow2, iSCSI and an improved version of vmware that uses the vmdkfs tools, which add to the shared and ssh drivers in OpenNebula 3.4.

Improved Amazon EC2 Adapter for Cloud Bursting OpenNebula 3.4 includes an improved driver to create hybrid clouds with Amazon EC2, to support most of the EC2 features like tags, security groups or VPC (Virtual Private Cluster).

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

OCCI updates The OCCI server has been improved to include user/group information in resources and extended information of resources.

New CloudAuth Driver This new driver delegates the authentication to the OpenNebula core. Therefore any OpenNebula auth driver can be used to authenticate cloud users or the Sunstone web UI.

Cloud Partitioning Cloud requests can be routed to an specific cluster with its own storage and network resources to better isolate public cloud users. Usually medium to large cloud sites structure their resources on multiple cluster each one with their own storage and networking systems. Cloud users are assigned to an specific cluster to prevent image trashing across large datacenters. This feature extends the previous cluster concept available in OpenNebula 3.0.

Improved Logging A new framework has been included to add logging information to the servers. In particular the logging facilities of the OpenNebula's Cloud API have been extended to ease the maintenance and deployment of several services like the Sunstone graphical interface or the OCCI API server (also used by the DeltaCloud API drivers).

ONE2ONE Driver for Cloud Federation OpenNebula manages external cloud as if they were local hosts. In this way the remote cloud is abstracted as a single but powerful resource. In order to interact with this special host a specific driver is needed that interacts with the remote cloud API in order to monitor and manage VM instances. In this period we have developed an driver that interacts with OpenNebula based clouds. It also integrates resources from other StratusLab services, specially the Marketplace.

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; however, SixSq is now recognized as an SME by the European Commission. The representatives for TCD and TID have changed because of retirements and internal reorganization of activities.

4.2 Management Tasks

Meetings Tables 4.1–4.8 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 4.9 contains the metrics for the project. The table groups related metrics together.

Deliverables and Milestones Tables 5.1, 5.2, and 5.3 list all of the documents. In addition, these are available from the project website. The deliverables and other project documents (excluding milestones) have been uploaded to the HAL[1] open access repository to ensure their longterm availability.

4.3 Issues

No particular issues have arisen in Q8.

4.4 Planning

See the deliverable describing the sustainability and exploitation plan (D3.5) for information about the planning after the project.

[1] http://hal.archives-ouvertes.fr/

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: *Meetings (Q6)*

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

Table 4.7: *Meetings (Q7)*

Title	Date	Venue	Comments
Sprint 20 Demo	19/12/2011	Phone/EVO	Sprint 20 demonstration meeting. http://indico2.lal.in2p3.fr/indico/categoryDisplay.py?categId=131
Technical Meeting (TSCG)	31/01/2012	Phone/EVO	F2F agenda; priorities for upcoming sprints. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1709
Sprint 21 Demo	30/01/2012	Phone/EVO	Sprint 21 demonstration meeting. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1705
Sprint 22 Demo	28/02/2012	Phone/EVO	Sprint 22 demonstration meeting. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1749
Technical Meeting (TSCG)	01/03/2012	Phone/EVO	Priorities for upcoming sprints. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1709

Table 4.8: *Meetings (Q8)*

Title	Date	Venue	Comments
Sprint 23 Demo	22/03/2012	Phone/EVO	Sprint 23 demonstration meeting. http://indico2.lal.in2p3.fr/indico/categoryDisplay.py?categId=1779
Technical Meeting (TSCG)	18/04/2012	Athens, Greece	Planning for EU Review. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1770
Project Management Board (PMB)	18/04/2012	Athens, Greece	Planning for EU Review; financial review. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1770
Face to Face Meeting	18-20/04/2012	Athens, Greece	StratusLab Face to Face Meeting http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1770
Sprint 24 Demo	27/04/2012	Phone/EVO	Sprint 24 demonstration meeting. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1828
Sprint 25 Demo	??/05/2012	Phone/EVO	21 stories demonstrated; see JIRA for details.
Sprint 26 Demo	29/05/2012	Phone/EVO	Sprint 26 demonstration meeting. http://indico2.lal.in2p3.fr/indico/conferenceDisplay.py?confId=1845

Table 4.9: *Metrics*

Metric	Q2	Q3	Q4	Y1 Target	Q5	Q6	Q7	Q8	Y2 Target
No. of people trained on StratusLab software	N/A	N/A	~ 25	–	~36	~89	0	25	–
No. of people on StratusLab announcement list	67	67	67	25	70	72	72	73	75
Registered users on StratusLab discussion site	N/A	N/A	N/A	50	N/A	12	13	16	100
No. of views of website	2922	4623	4579	–	5472	4985	4442	5372	–
No. of completed sprints	5	5	4	–	3	3	3	4	–
No. of releases	1	1	1	–	2	1	2	2	–
No. of open user stories	38	72	101	–	118	107	121	102	–
No. of implemented user stories	69	40	67	–	50	48	57	61	–
No. of open bugs	6	15	22	–	28	51	51	66	–
No. of fixed bugs	7	11	27	–	14	20	17	14	–
No. of prod. sites running StratusLab dist.	1	1	1	5	1	3	5		10
Availability of hosted grid sites	N/A	N/A	100%	80%	91%	74%	93%	98%	95%
Reliability of hosted grid sites	N/A	N/A	100%	80%	92%	78%	93%	98%	95%
No. of VOs served via StratusLab hosted grid sites	0	1	1	10	21	18	18	18	30
No. of sci. disciplines served via StratusLab hosted grid sites	0	0	0	3	11	9	9	9	15
Delivered computing resources through hosted grid services	N/A	16 cores	16 cores	–	32 cores	32 cores	32 cores	32 cores	–
Delivered computing resources through hosted cloud services	N/A	256 cores	256 cores	–	256 cores	288 cores	256 cores	256 core	–
Storage provided through cloud service	N/A	N/A	N/A	–	0	0	0	3 TB	3 TB
No. of jobs run in hosted grid site	N/A	N/A	N/A	–	13,960	16,916	28,345	12,930	–
Norm. CPU time consumed in the hosted grid site (hrs)	N/A	N/A	N/A	–	26,202	14,231	87,671	51,895	–
No. base machine images	5	7	8	5	8	13	–	–	10
No. of base machine image downloads	783	2628	7072	–	7225	6657	–	–	–
No. appliances	0	6	7	5	7	7	–	–	15
No. of appliance downloads	0	252	687	–	1010	426	–	–	–
No. of Marketplace metadata entries	–	–	–	–	–	–	111	114	–
No. of Marketplace endorsers	–	–	–	–	–	–	24	35	–
No. of Marketplace base images	–	–	–	–	–	–	86	71	–
No. of Marketplace appliances	–	–	–	–	–	–	25	43	–

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	03/10/2011	Done	Yes	Delayed PM15-16
D6.4	Cloud-like Management of Grid Sites 2.0 Design Report		WP6	TID	R	PU	PM17	15/12/2011	Done	Yes	Delayed PM17-18
D5.4	Economic Analysis of Infrastructure Operations		WP5	GRNET	R	PU	PM18	19/03/2012	Done	Yes	Delayed to PM21
D6.5	Cloud-like Management of Grid Sites 2.0 Software		WP6	TID	P	PU	PM23	11/05/2012	Done	Yes	
D2.4	Final Report on StratusLab Adoption		WP2	CNRS	R	PU	PM24	06/06/2012	Done	Yes	
D2.5	Report on Evaluation of StratusLab Products		WP2	CNRS	R	PU	PM24	14/06/2012	Done	Yes	
D3.4	Final Review of Dissemination, Collaboration and Standardization Activities		WP3	TCD	R	PU	PM24	29/05/2012	Done	Yes	
D3.5	Exploitation and Sustainability Final Plan		WP3	TCD	R	PU	PM24	22/05/2012	Done	Yes	
D4.5	StratusLab Toolkit 2.0		WP4	SIXSQ	P	PU	PM24	28/05/2012	Done	Yes	
D4.6	Software Integration Final Report		WP4	SIXSQ	R	PU	PM24	28/05/2012	Done	Yes	
D5.5	Infrastructure Operations Final Report		WP5	GRNET	R	PU	PM24	05/06/2012	Done	Yes	
D6.6	Cloud-like Management of Grid Sites Research Final Report		WP6	TID	R	PU	PM24	04/06/2012	Done	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	Yes	19/03/2012	
MS12	Delivery of Virtual Appliance Repository	WP5	GRNET	PM18	Yes	15/12/2011	
MS5	Opening of Virtual Appliances Repository to External Application Communities	WP2	CNRS	PM20	Yes	19/03/2012	
MS15	Release of Cloud-like Management of Grid Services and Resources 2.0 Beta	WP6	TID	PM21	Yes	19/03/2012	
MS9	Release of StratusLab 2.0 Beta	WP4	SIXSQ	PM22	Yes	02/05/2012	
MS13	Operation of Site Running StratusLab Toolkit v2.0	WP5	GRNET	PM22	Yes	02/05/2012	

6 Use of Resources

See the “Use of Resources” section in the second periodic report.

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)