



Quarterly Report (QR2)

Charles Loomis, David O'Callaghan, Marc-Eliau Bégin, Evangelos Floros,
Henar Muñoz Frutos

► To cite this version:

Charles Loomis, David O'Callaghan, Marc-Eliau Bégin, Evangelos Floros, Henar Muñoz Frutos. Quarterly Report (QR2). 2010. hal-00687980

HAL Id: hal-00687980

<https://hal.science/hal-00687980>

Submitted on 16 Apr 2012

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

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



Enhancing Grid Infrastructures with
Virtualization and Cloud Technologies

Quarterly Report

Quarterly Report QR2 (V1.0)
15 December 2010

Abstract

In the second quarter, the project has successfully created the first public release of the StratusLab cloud distribution. To complement the release, the project participants have provided user support, deployed a reference infrastructure for outside users, and increased awareness of the release and the project. High-level features, including advanced service management features, have been defined and will be progressively added to the series of public releases leading to the v1.0 release due at the end of the first year.



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



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

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

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

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



Contributors

Name	Partner	Sections
Charles Loomis	CNRS-LAL	All
David O’Callaghan	TCD	WP3
Marc-Eliau Bégin	SIXSQ	WP4
Vangelis Floros	GRNET	WP5
Henar Muñoz	TID	WP6

Document History

Version	Date	Comment
0.1	17 Nov. 2010	Initial outline of report.
0.2	10 Dec. 2010	Nearly complete draft excepting resource descriptions and some updated activity descriptions.
1.0	15 Dec. 2010	Final draft.

Contents

List of Figures	6
List of Tables	7
1 Publishable Summary	8
1.1 Project Context and Objectives	8
1.2 Summary of Work Performed and Achievements	9
1.3 Final Results and Potential Impact and Use	10
1.3.1 Improved Interdisciplinary Scientific Collaboration	10
1.3.2 Impact on DCI Evolution.	11
1.3.3 Improved Usability of DCI Platforms	11
1.4 Contact Information	11
2 Project Objectives for the Period	13
2.1 Objectives	13
2.1.1 Quarter 1	13
2.1.2 Quarter 2	13
2.2 Review Recommendations	14
3 Progress and Achievements	15
3.1 WP2: Interaction with Targeted Communities	17
3.1.1 Progress Towards Objectives by Task	17
3.1.2 Issues and Corrective Actions	19
3.1.3 Use of Resources.	19
3.2 WP3: Dissemination.	20
3.2.1 Progress Towards Objectives by Task	20

3.2.2	Issues and Corrective Actions	23
3.2.3	Use of Resources	23
3.3	WP4: Software Integration and Distribution	24
3.3.1	Progress Towards Objectives by Task	24
3.3.2	Issues and Corrective Actions	26
3.3.3	Use of Resources	27
3.4	WP5: Infrastructure Operation	28
3.4.1	Progress Towards Objectives by Task	28
3.4.2	Issues and Corrective Actions	30
3.4.3	Use of Resources	31
3.5	WP6: Innovative Cloud-like Management of Grid Services and Resources	32
3.5.1	Progress Towards Objectives by Task	32
3.5.2	Issues and Corrective Actions	33
3.5.3	Use of Resources	33
4	Project Management	34
4.1	Consortium	34
4.2	Management Tasks	34
4.3	Issues	35
4.4	Planning	35
4.4.1	Objectives for Next Quarter	35
4.4.2	Roadmap.	36
5	Deliverables and Milestones	40

List of Figures

3.1 Visits for Q2 showing the increase in traffic around the first software release.	21
--	----

List of Tables

1.1	StratusLab Information and Support	12
1.2	StratusLab Partners	12
3.1	Talks	22
3.2	WP5 Infrastructure Services	28
4.1	Effort (in Person-Months) by Partner	34
4.2	Effort (in Person-Months) by Work Package.	35
4.3	Meetings	38
4.4	Metrics	39
5.1	Deliverables (Year 1).	41
5.2	Deliverables (Year 2).	42
5.3	Milestones	43

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 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 will create an complete, coherent, open-source private cloud distribution to allow administrators of grid resources centers to take advantage of virtualization and cloud technologies. It will also provide 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 second quarter, the project has successfully created the first public release of the StratusLab cloud distribution. To complement the release, the project participants have provided user support, deployed a reference infrastructure for outside users, and increased awareness of the release and the project. High-level features, including advanced service management features, have been defined and will be progressively added to the series of public releases leading to the v1.0 release due at the end of the first year.

Initial Public Release The project released v0.1 of the StratusLab cloud distribution. This release provides a minimal cloud distribution that allows remote access to the cloud, easy access to base images (ttylinux, Ubuntu, and CentOS) stored in the StratusLab appliance repository, contextualization of those images, and management of the full virtual machine lifecycle.

Reference Infrastructure The v0.1 release of the StratusLab cloud distribution has been installed on dedicated hardware at GRNET to provide a reference cloud infrastructure. This infrastructure has been opened to external users to allow them to easily test drive the StratusLab release and to provide feedback. In parallel with the deployment of the infrastructure, security policies and user management procedures have been put in place.

User Support To ensure that users receive the appropriate support and have the means to provide feedback, a support mailing list was created, backed by a first-line support team consisting of people from the user support, integration, and operations activities. All project participants will provide second-line support as needed. An online tutorial and video version of it help users understand the features in the release and how to use them.

Increased Visibility The project made a concerted effort to increase the visibility of the project, particularly at the EGI Technical Forum in Amsterdam. Several high-profile presentations were given and the project manned a booth showing the first results of the project. The dissemination activity has diversified the means that people can find out about the project using, for example, Twitter and YouTube. It has also improved the website with better integration with social media and bet-

ter targeted RSS feeds. It also developed a release dissemination strategy for the v0.1 release that will be used (and improved) for each of the upcoming StratusLab releases.

Improved Software Processes The successful public release of the StratusLab distribution early in the project's life is a direct result of the agile/scrum software development processes used by the project. Significant work was done to expand testing of the StratusLab software and to automate both the testing and generation of the release.

Service Manager The overall architecture of the StratusLab distribution was extended with a "service manager" that will allow ensembles of machines (complete "services") to be specified and controlled as a block. Moreover, the service manager will allow the dynamic allocation or reallocation of resources based on service load. Claudia, a development from TID, was selected as the basis for the service manager.

The successful initial public release of the StratusLab cloud distribution and associated supporting activities demonstrates the consortium's ability to create a useful open-source cloud distribution with a potentially high positive impact on European e-Infrastructures. Q3 will see this release evolve into a feature-complete beta release of the StratusLab cloud distribution.

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.

StatusLab 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 Nationale 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-Eliañ BEGIN meb@sixsq.com
TID	Telefónica Investigación y Desarrollo SA	Juan CACERES caceres@tid.es
TCD	The Provost Fellows and Scholars of the College of the Holy and Undivided Trinity of Queen Elizabeth Near Dublin	Brian COGHLAN coghlan@cs.tcd.ie

2 Project Objectives for the Period

2.1 Objectives

2.1.1 Quarter 1

In this first quarter, the primary objective was to prepare the foundations for a successful project. In more detail this involved:

- Deployment of collaborative software development tools,
- Starting dialog between StratusLab and targeted communities,
- Make the project visible to targeted communities and general public,
- Put in place the software development processes and policies,
- Define the initial architecture for the StratusLab software, and
- Deploy the initial project infrastructure.

Within this quarter all of these have been obtained providing a solid basis for the first public release of the StratusLab software in Q2 with additional features appearing rapidly afterwards.

2.1.2 Quarter 2

In the second quarter, the emphasis was on making the first public release of the StratusLab cloud distribution. Detailed objectives were:

- Increase project visibility particularly at the EGI Technical Forum,
- Initial public release of StratusLab cloud distribution,
- Reference infrastructure available to outside users,
- Support provided for release, and
- Initial design of advanced management services.

All of these objectives have been met, allowing the project to build a feature-complete release during the next quarter.

2.2 Review Recommendations

Not yet applicable.

3 Progress and Achievements

In the second quarter, the project has successfully created the first public release of the StratusLab cloud distribution. To complement the release, the project participants have provided user support, deployed a reference infrastructure for outside users, and increased awareness of the release and the project. High-level features, including advanced service management features, have been defined and will be progressively added to the series of public releases leading to the v1.0 release due at the end of the first year.

Initial Public Release The project released v0.1 of the StratusLab cloud distribution. This release provides a minimal cloud distribution that allows remote access to the cloud, easy access to base images (ttylinux, Ubuntu, and CentOS) stored in the StratusLab appliance repository, contextualization of those images, and management of the full virtual machine lifecycle.

Reference Infrastructure The v0.1 release of the StratusLab cloud distribution has been installed on dedicated hardware at GRNET to provide a reference cloud infrastructure. This infrastructure has been opened to external users to allow them to easily test drive the StratusLab release and to provide feedback. In parallel with the deployment of the infrastructure, security policies and user management procedures have been put in place.

User Support To ensure that users receive the appropriate support and have the means to provide feedback, a support mailing list was created, backed by a first-line support team consisting of people from the user support, integration, and operations activities. All project participants will provide second-line support as needed. An online tutorial and video version of it help users understand the features in the release and how to use them.

Increased Visibility The project made a concerted effort to increase the visibility of the project, particularly at the EGI Technical Forum in Amsterdam. Several high-profile presentations were given and the project manned a booth showing the first results of the project. The dissemination activity has diversified the means that people can find out about the project using, for example, Twitter and YouTube. It has also improved the website with better integration with social media and better targeted RSS feeds. It also developed a release dissemination strategy for the v0.1 release that will be used (and improved) for each of the upcoming StratusLab releases.

Improved Software Processes The successful public release of the StratusLab distribution early in the project's life is a direct result of the agile/scrum software development processes used by the project. Significant work was done to expand testing of the StratusLab software and to automate both the testing and generation of the release.

Service Manager The overall architecture of the StratusLab distribution was extended with a “service manager” that will allow ensembles of machines (complete “services”) to be specified and controlled as a block. Moreover, the service manager will allow the dynamic allocation or reallocation of resources based on service load. Claudia, a development from TID, was selected as the basis for the service manager.

The successful initial public release of the StratusLab cloud distribution and associated supporting activities demonstrates the consortium's ability to create a useful open-source cloud distribution with a potentially high positive impact on European e-Infrastructures. Q3 will see this release evolve into a feature-complete beta release of the StratusLab cloud distribution.

3.1 WP2: Interaction with Targeted Communities

In parallel with the first public release of the StratusLab cloud distribution, WP2 has worked to expand the number of people (and communities) aware of the distribution, solicited feedback on its installation and use, and provided support to our first users. Strong interactions continue with the bioinformatics community and researchers from ATLAS (high-energy physics experiment) at LAL. Initial contacts with the developers of OpenMole, a task management framework, and the EDGI project promise to diversify our user community. A strong effort has been made on tutorials as a means of providing support and increasing awareness of the StratusLab distribution. WP2 has a self-guided tutorial linked from the release page, a 5-minute screencast version of it, and live tutorials at several scientific meetings. WP2 coordinates and contributes to the support team and collects feedback from users and administrators. The WP2 partners continue to evaluate the latest versions of the StratusLab distribution and to fill functionality gaps: for example, authentication of virtual machine metadata and support for grid certificate authentication.

3.1.1 Progress Towards Objectives by Task

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

ATLAS Experiment CNRS/LAL continues to work with people within the laboratory from the ATLAS experiment. We have helped them develop a virtual machine image containing commercial software for the design and testing of custom integrated circuits. This application uses license tokens and is a good example of the integration of commercial applications on a cloud infrastructure. (GRNET in WP5 has also discussed with people from MathWorks about providing a MATLAB appliance compatible with StratusLab.)

EDGI Project CNRS/LAL has discussed with the EDGI project at both the managerial and technical levels on how EDGI can use virtualization technologies to provide better quality of service on desktop grid systems. A formal Memorandum of Understanding (MoU) will likely be developed between the projects to define the collaboration. At a technical level, participants in EDGI from CNRS/LAL and INRIA have asked for support for creating machine images and deploying them on the LAL preproduction cloud service.

OpenMole CNRS/LAL has discussed with the author of the OpenMole system (a workflow/task management system) about how resources within a StratusLab cloud could be used. Although in the early stages, the authors are very interested in collaborating with the project and initial support for the use of cloud resources through OpenMole is expected in the next quarter.

Bioinformatics Community CNRS/IBCP has contributed to several events involving the French bioinformatics community and has organized meetings with the French RENABI GRISBI community. It regularly presents the StratusLab project and results at these meetings to raise awareness of the project and to promote col-

laboration. In particular, StratusLab was presented at the national GRISBI scientific school (27 September, 40 attendees, Roscoff, France) and at the regional scientific workshop (10 November, Lyon, France) of PRABI (Rhône-Alps Bioinformatics Platform). CNRS/IBCP has initiated a collaboration with the French bioinformatics platform GenOuest about the deployment and usage of StratusLab distribution for bioinformatics applications. A first workshop took place at IBCP on 18-19 October 2010. These workshops and meetings were used to collect use cases and requirements.

Contacts with System Administrators CNRS/LAL has discussed with system administrators at BELNET (Belgium) and at RAL (UK) about the StratusLab distribution. Both have installed the public release via Quattor and have provided feedback to the project. CNRS/IBCP has organized meetings with the French RENABI GRISBI administrators who provide resources to the French bioinformatics community.

User Tutorial CNRS/LAL has prepared a self-guided user tutorial that is available from the StratusLab web site. In addition, a 5-minute “screencast” of the tutorial has been prepared. It is also available from the release page on the StratusLab website as well as on YouTube (thanks to WP3). CNRS/IBCP has presented tutorials about how cloud technologies could meet the needs of bioinformatics users. These tutorials were based on OpenNebula and the StratusLab distribution; they included a detailed description of the StratusLab project, its goals, current releases, and functionality interesting for bioinformatics applications. These tutorials took place at the national GRISBI scientific school (1 October 2010, 22 students, Roscoff, France) and at the RENABI GRISBI steering committee meeting (19 November, 12 attendees, Toulouse, France).

3.1.1.2 Task 2.2: Intensive Evaluation of StratusLab Products

Image Metadata Sharing of virtual machine images will require a standard format for the image metadata and a method of signing that metadata to ensure its authenticity. CNRS/LAL has developed tools for signing and validating metadata information in XML files using XML Digital Signature API. Grid Certificates or DSA/RSA keys can be used to sign the metadata. When grid certificates are used, the identity of the signer is extracted and printed during the metadata validation.

Application Benchmarks The application benchmarks have been included as a part of the first public StratusLab release. These are available to both end-users and system administrators. They will eventually be used by the project to evaluate the efficiency of different deployment scenarios.

Grid Certificate Authentication Integration with EGI will require the support of the grid authentication mechanisms. CNRS/LAL had developed a proxy server that allows authentication based on grid certificates. This initial prototype proxies the OpenNebula XMLRPC interface. This proxy service will evolve as the project moves towards the OCCI interface. Full support of grid identities, groups, and

roles will require modifications to OpenNebula itself.

Quattor Configuration CNRS/IBCP has evaluated the StratusLab installation with Quattor by deploying a front-end and nodes on their local resources in Lyon, following the online “Quattor Installation Guide”.

Bioinformatics Appliance CNRS/IBCP has worked on the definition of a bioinformatics appliances consisting of a gLite Worker Node with pre-installed bioinformatics applications and NFS mounts of biological data. A first instance was integrated in the grid site IBCP-GBIO and is under evaluation.

3.1.2 Issues and Corrective Actions

No major issues related to WP2 have arisen in Q2.

3.1.3 Use of Resources

The effort involved in the activity is significantly less than planned. This has occurred mainly because of problems with hiring, not entirely related to the delayed prefinancing. To date, this underspending has not had a significant impact on the work program. Fortunately, hiring issues have been resolved. With that and the expected contributions from TID, it is expected that WP2 will be at full force for the remainder of the project.

3.2 WP3: Dissemination

StratusLab began its second quarter dissemination and collaboration activities with a strong presence at the EGI Technical Forum 2010, in Amsterdam. StratusLab prepared for its first software release—and the associated dissemination activities—in early November. The release was announced through a number of online outlets and a tutorial video produced in WP2 was published on YouTube. Prior to the software release, the website was updated to support social network sharing and to allow better analysis of visitor/download patterns.

At the EGI Technical Forum, ICT 2010, the 8th e-Infrastructure Concertation meeting, and other events, StratusLab partners had the opportunity to develop connections and collaborations with other projects and potential users. Collaborations with DCI projects are under way with the expectation that these will be formalized through Memoranda of Understanding in the coming quarter.

3.2.1 Progress Towards Objectives by Task

3.2.1.1 Task 3.1: Dissemination

Release Dissemination In preparation for the first software development release (version 0.1) on 9 November 2010, a *release dissemination plan* was created that covered the main dissemination targets for the release. For the initial release we wanted plenty of awareness of the software and the project, but with the focus on those groups who would be most interested in testing a version of the software that was not ready for full production use. With this in mind, the release was disseminated to our opt-in announcements list and relevant online media.

The release dissemination plan will be revised as necessary for future public releases. The project plans to make public, development releases every six weeks.

Media & Publications A press release was prepared for the first software development release and distributed to a number of outlets that had been interested previously in our work. This was picked up by *HPCwire*, *Sys-Con*, *CloudExpo News*, *DSA Research Blog*, and *OpenNebula Blog*. In addition, an announcement appeared in *International Science Grid This Week* on 14 November 2010.

Work continued on an article in *International Science Grid This Week* (<http://isgtw.net/>) covering the results of the StratusLab surveys carried out by WP2 in Q1 and the recent software release. This was published on 24 November 2010. StratusLab was also mentioned in the e-ScienceTalk *Grid Briefing* entitled *Mapping the e-Infrastructure Landscape* published in November 2010.

Website The project website (<http://www.stratuslab.eu>) was updated to include some social media features, such as a link to the project Twitter account (@Stratus-Lab) and social bookmarking / messaging tools, allowing visitors to post links to Twitter, Facebook, Digg and others. The website RSS feed was modified to show project news rather than wiki changes, making it more useful for outside users.

Website monitoring was improved to give the project better visibility of the number of visitors and software downloads. Google Analytics gives a clear view




Figure 3.1: Visits for Q2 showing the increase in traffic around the first software release.

of website traffic and, for instance, allows us to visualize spikes in traffic caused by media exposure (See Figure 3.1). AWStats (running locally on the web server) can provide in-depth statistics for non-web files, including software downloads.

EGI Technical Forum 2010 The project booked an exhibition booth at EGI Technical Forum 2010 and presented posters, demonstrations and other dissemination material. StratusLab had a very visible presence at the event. Project members gave a number of high-profile presentations, listed in Table 3.1.

StratusLab members participated in the session *Enabling Clouds for e-Science*, which brought together a number of international cloud computing projects.

Talks A brief list of talks describing the project delivered during Q2 is given in Table 3.1. Details and links, where available, are given on the project website at <http://stratuslab.eu/doku.php?id=presentations>.

3.2.1.2 Task 3.2: Collaboration with Standards Bodies and Related Projects

The project published its “Initial Plan for Dissemination, Collaboration and Standardization Activities” (D3.1) early in this quarter.

e-Infrastructure Charles Loomis participated in the 8th e-Infrastructure Concertation Meeting on 4–5 November 2010, CERN in Geneva, Switzerland (<http://www.e-sciencetalk.org/e-concertation/>).

SIENA At the 1st SIENA Roadmap Event, 27 October 2010, in conjunction with OGF30, Brussels, Juan Caceres represented StratusLab at the Roadmap Editorial Board: a group of national and international distributed computing initiative members, active contributors to standards, and collaborating e-Infrastructure users serve on an editorial board to consolidate the Roadmap’s wide input.

Venus-C Project partners participated in several meetings with the Venus-C project about collaboration between both projects.

IGE Project partners met with the IGE project about collaboration between both projects, to deploy Globus middleware and services on StratusLab sites.

Table 3.1: Talks

Title / Event	Date
EuroPar 2010	31 August – 3 September 2010
Integrating Cloud Monitoring and Accounting with Grid Operational Tools	14 September 2010
Operational Considerations From Running Grid Services on Cloud Resources	14 September 2010
StratusLab Accounting Requirements, EGI Technical Forum 2010	15 September 2010
Cloud Challenges, EGI Technical Forum 2010	16 September 2010
ICT 2010, Brussels, Belgium	27–29 September 2010
GRISBI Bioinformatics School, Roscoff, France	1 October 2010
OGF30/Grid2010	25–28 October 2010
CloudComp 2010	26–28 October 2010
ISC Cloud 2010	28–29 October 2010
HP Labs Seminar	November 2010
7th International Cloud Computing Expo	1 November 2010
SlipStream and StratusLab, Cloud & ICT 2.0 Summit, Geneva	18 November 2010
OW2 Annual Conference	24–25 November 2010

EGI-InSPIRE StratusLab will contribute to EGI-InSPIRE deliverable D2.6 on *Integration of Clouds and Virtualization into the European Grid Infrastructure*.

FP7 Proposals StratusLab partners were approached by a number of project consortia interested in future collaborations. The project issued Letters of Support for four projects.

3.2.1.3 Task 3.3: Development of Exploitation and sustainability Plan

This task will begin in Q4.

3.2.2 Issues and Corrective Actions

No major issues related to WP3 have arisen in Q2.

3.2.3 Use of Resources

Effort increased in Q2 with respect to Q1 as all partners are now actively participating in project dissemination and collaboration. The initial effort in Q1 was lower as hiring was not complete at that time, notably at the WP3 lead institution. There was not a significant impact of on the work programme as some of the work was performed by unfunded resources.

While effort is expected to increase in the next quarter, it will see a greater increase in Q4 when Task 3.3 begins and the deliverable reports D3.2 and D3.3 will be prepared.

3.3 WP4: Software Integration and Distribution

During Q2, WP4 focused on the first release of StratusLab distribution: v0.1. This was also the opportunity to assert that our architecture defined during Q1 was sound and able to support the key features planned in StratusLab v0.1. This major effort included the integration, test and certification of a number of tools to ensure that StratusLab would be consistent, easy to install, as well as easy to use.

The agile/scrum process put in place during Q1 was consolidated and worked well in providing feedback from short iterations (i.e. three weeks) and ensuring that all partners are working on a coherent set of objectives and that fluid and constant communication takes place between all.

A significant effort was also deployed in continuing automation efforts (noticeably by CNRS/LAL and SixSq) that had started during QR1 of our build, test and release process.

3.3.1 Progress Towards Objectives by Task

3.3.1.1 Task 4.1: Definition of Reference Architecture

This task was not active during Q2. Having said that, the architecture of the system is regularly reviewed to ensure that all features required of the StratusLab distribution can be provided effectively. To date, no dramatic changes are required or foreseen in the architecture captured in D4.1 and extended in D6.1.

As part of this ongoing work, the StratusLab roadmap was discussed during the Lyon F2F meeting, with resulting new interfaces identified for future versions of the system. See the planning section for details.

3.3.1.2 Task 4.2: Integration of Open-source Distribution

To support the test, certification and release process, WP4 invested significant efforts in creating an automated and continuous integration system (noticeably CNRS/LAL and SixSq). Using a series of open source tools (e.g. Hudson, Maven, Yum repository). We now have several machines, with the support from GRNET, able to test several possible deployments, including both CentOS and Ubuntu operating systems. This gives us the ability to verify that StratusLab can successfully be configured in different ways to match system administrators requirements.

This automation strategy was also adopted by the OpenNebula team at UCM where integration tests were developed to check the functionality of OpenNebula. A set of functionality tests were also created that manage real virtual machines. All these tests are integrated into a hudson server that runs them every day. Similarly to what SixSq and GRNET created for the StratusLab's integration facilities, UCM set up a private infrastructure to develop StratusLab components with multiple configurations in terms of hypervisors (Xen, KVM and VMware), storage systems (Shared FS, ssh) and installation modes (system wide, installation and runtime directories). This setup was effective at finding new bugs, such as: new remote action management issues, improvements in the OpenNebula packages, sup-

port for Xen4 and specific disk drivers and new architecture attribute for OS section of VM templates.

A series of tools and services were integrated, developed and tested in order to deliver the functionality required for StratusLab v0.1. Here are the main features included in the first version of the distribution:

OpenNebula virtual machine management

Web Monitor administrator dashboard for cloud

User CLI Python-based command line tools for remote access to cloud

System Administrator CLI Command line tools to facilitate manual installation and management of the cloud

Automated (Quattor) Installation Templates for automatic installation using the Quattor fabric and configuration management tool

All of these features were selected by the project technical group, defined and tracked using JIRA/GreenHopper.

The Appliance Repository developed during Q1 by TCD was also included in the release, but not officially, since not considered critical. A reference Appliance Repository was deployed by TCD, and a mirror instance was also provided by GRNET for backup.

StratusLab distribution v0.1 was released under the Apache 2.0 license, with copyright owned by the major contributors of each respective tool and/or service.

All technical issues identified, in preparation for each sprint, as well as during, were tracked using the project tracking tool (JIRA/GreenHopper).

In preparation for future features, we have also explored requirements on OpenNebula to better support image creation. For example, qcow could be used to automatically create and use sparse images was identified and could be integrated. Further, in order to better manage user authentication, we should be able to support ldap and (grid) certificates. In OpenNebula, an analysis was performed by UCM and work started to provide such support.

Preparation work also took place for the integration of Claudia, the Service Manager component developed by TID and scheduled for a future release of StratusLab.

3.3.1.3 Task 4.3: Contextualization of Grid Services

The generic mechanism devised during Q1 for contextualization was simplified and implemented for production use during this quarter and released as part of StratusLab v0.1. This included the decision to rely on DHCP for public IP address / MAC address assignment, which significantly simplifies the configuration work for the system administrators. Further simplifications are also possible on this topic, for example for private IP assignment.

StratusLab v0.1 was released with three fully contextualized reference images for the Ubuntu, CentOS, and ttylinux operating systems.

The StratusLab contextualization strategy was implemented as a coherent set, using OpenNebula idiomatic setup, configuration and end-user command-line tools.

3.3.1.4 Task 4.4: Technical Support

WP4 has provided support for the StratusLab tools and OpenNebula to the whole project. This support has been provided via the daily standup meetings, phone calls, Skype, and email.

WP4 has continued to support and manage our software development procedure based on scrum, an agile software development process. WP4 manages this process by running the daily standup meetings, the sprint demo meetings and the sprint preparation meetings. During Q2, our Scrum process was consolidated, with the addition of effort estimates such that the project can be better managed and scheduling made more reliable.

As the project identifies bugs and issues with each new version of StratusLab, including OpenNebula, both SixSq and UCM have been active at fixing these bugs and addressing the identified issues. When bugs were identified in core components of StratusLab, when possible, patches were created and integrated. In parallel, these patches were sent to the owner of the component for future fix. Once these patches were integrated and released by the provider, the patches were then removed from the StratusLab code-base. This ensured that we were not blocked moving forward. As a result of bug fixes, including ones found by StratusLab partners, UCM will release a maintenance version of OpenNebula (2.0.1) this quarter.

The technical support is now extended to a number of wiki pages on our website, as well as a FAQ page to which WP4 will contribute as recurrent questions and issues are raised by members and users. The previously created wiki pages for internal use have been cleaned-up and updated in order to be made public.

3.3.2 Issues and Corrective Actions

Following from our first release, during the Lyon face-to-face meeting, a retrospective was conducted. The retrospective is an important event that all agile methods mandate. It is a mechanism that encourages the team, and all its stakeholders, to reflect on its performance. This is an important tenet in the concept of ‘continuous improvement’ that lies at the heart of agile. The two items on which the retrospective focused were:

1. Increasing clarity and priorities on our technical program of work prior, to improve the planning meeting
2. Contrasting views on the effectiveness of the daily meetings

For the first item, we agreed that having the technical group meetings would provide the right platform to discuss ‘medium term strategic’ issues. The technical group had not met as often as it could have. This was felt by all participants as the right forum to address this issue. We also agreed to keep the three weeks sprint as a meeting frequency.

For the second item regarding the daily meetings, since we had different views on the effectiveness of this meeting, we discussed the issue and tried to better understand everybody's viewpoints. The meetings are managed such that they never exceed 15 minutes (with only one exception since the beginning of the project, with 16 minutes), and normally are over in 10 minutes. The daily meeting take place every day at 10:30 (sharp!), Paris time. For people participating regularly to the daily meetings, they are efficient and useful. Others felt that the stand-ups were too frequent. Alternatives were discussed such as traditional sit-down weekly meetings, but these would reduce the flow of communication, as well as our responsiveness, and would force us to start taking and distributing minutes.

The compromise in the end was that the daily stand-ups will continue with some people participating less frequently, but regularly.

3.3.3 Use of Resources

The effort involved in the activity dropped during this quarter. This was partly due to the delayed pre-financing. New staff are being hired and will be in place during quarter three. The program of work was not impacted during this quarter with the successful release of StratusLab v0.1. WP4 will be at full capacity, from quarter three, for the remainder of the project.

Table 3.2: WP5 Infrastructure Services

Production Cloud Service	GRNET	http://cloud-grnet.stratuslab.eu:2633
Preproduction Cloud Service	GRNET	http://node006.one.ypepth.grnet.gr:2633
Project Tools (Hudson Server)	GRNET	http://hudson.stratuslab.eu:8080
Appliance Repository	TCD	http://appliances.stratuslab.eu
App. Repository Mirror	GRNET	http://appmirror-grnet.stratuslab.eu/images
Test Infrastructure	LAL	https://onehost-2.lal.in2p3.fr:2643/RPC2
Test Infrastructure	GRNET	http://node003.one.ypepth.grnet.gr:2633

3.4 WP5: Infrastructure Operation

WP5 is responsible for the provision of the computing infrastructure required by the various activities within the project. During the second quarter of the project, work towards many of the initial goals of the work package progressed satisfactorily delivering important results. Among the highlights of the past quarter is the opening for public access of a reference cloud service in parallel with the release of v0.1 of StratusLab distribution. Along with the appliance repository that contains a set of basic VM images, also prepared by WP5, this service plays the role of the technological preview of the results delivered by the project and provides a set of elementary IaaS services to third parties.

Overall, WP5 has deployed a significant number of services, either as production service for public access or private services for testing or support for the development activities of the project. These services utilize the physical infrastructure offered by the project partners. Table 3.2 summarizes the services that currently offer web-based access in the context of WP5.

In the following paragraphs we provide more details about the various achievements of WP5 during Q2, grouped by subtask.

3.4.1 Progress Towards Objectives by Task

3.4.1.1 Task 5.1: Deployment and Operation of Virtualized Grid Sites

Production Cloud Service With the release of StratusLab 0.1 distribution, a reference cloud service was deployed and made available to the public. The StratusLab distribution has been installed in the physical infrastructure allocated by GRNET. Coupled with the Appliance Repository maintained by TCD, this infrastructure will play the role of the production cloud service deployed by the project, allowing people outside the project to test-drive a reference installation of the 0.1 release and the IaaS cloud developed by the project. A total of 11 nodes have been allocated for this purpose offering 160 CPU cores and 528 GB of memory. Depending on the demand more nodes will be added to the reference cloud service.

For what concerns the offered QoS, the StratusLab cloud services are provided on a best-effort basis, with no guarantees about the availability and stability of the service. As the distribution matures the infrastructure is expected to become more stable and reliable. Our goal is to offer, in the coming months, a production cloud service with a high-quality Service Level Agreement.

Detailed instructions for accessing and using the service are available from the project's wiki site at <http://www.stratuslab.org/doku.php?id=referenceservices>.

Pre-production Services LAL has been discussing this quarter with the Security Officers of IN2P3 to come to a workable solution for the firewall around LAL's preproduction cloud service. An agreement was reached at the end of the quarter that will allow all grid service ports as well as ssh, http(s), and ldap(s) to be accessible from the WAN for running virtual machines. Access to physical machines will have more severe restrictions with ssh access only allowed from within the LAL site. The purchase order for hardware for the preproduction cloud service has been prepared. The hardware should arrive in December 2010 with rapid deployment of the service.

Support Infrastructure GRNET went through a re-organization process of the pre-production infrastructure in order to optimize the resource utilization and plan for future workload demands. Currently this infrastructure is used for hosting project services: the hudson continuous integration service, the appliance repository mirror, and two testing sites installed with the StratusLab distribution.

3.4.1.2 Task 5.2: Testing of the StratusLab Toolkit

Testing Infrastructure LAL continues to maintain the Quattor configuration for the StratusLab toolkit. The configuration and associated deployment are frequently updated according to feedback from administrators who have installed the StratusLab release.

Verification of StratusLab Release 0.1 During the period that preceded the release of StratusLab v0.1, the distribution went through an intensive testing process. For this purpose a number of nodes were allocated from GRNET's support and pre-production infrastructure. During this testing process a significant number of bugs were identified and fixed, and the overall stability of the distribution was enhanced. This process took place in close collaboration with WP4 developers and focused on mainly two aspects of the StratusLab distribution: the manual installation of the service and remote client access via the command line tools.

In both cases, the online tutorials were used as a guides, namely the "Manual installation tutorial" and the "User tutorial". As a result, this testing and verification process improved the online support material as well as exposed problems with the software.

3.4.1.3 Task 5.3: Virtual Appliances Creation and Maintenance

Appliance Repository The virtual appliance repository is available as a service hosted by TCD. This first version the appliance repository is a standard Apache

web server, accessed using WebDAV (Web-based Distributed Authoring and Versioning) with authentication via the StratusLab LDAP server. The focus during this quarter has been on providing a reliable, stable service to support the first release. The main activities have been the following:

- The set of appliances available from the repository has been rationalized. Only images that have been tested, and are known to work with the StratusLab release are included. Currently three reference images are provided for Ubuntu, CentOS, and ttylinux.
- The front-end of the repository has been modified to match the look and feel of the main StratusLab web page. Users can now more easily find the available images, and also obtain information about them. This information is loaded from the metadata files stored with the images.
- A statistics package has been installed on the repository to provide a way of tracking the downloads of images. This is particularly useful to track the impact of the first release.
- The WP4 installation tools have been extended to allow for the installation of a local appliance repository. The tools were used by GRNET to deploy an appliance repository mirror which is intended to serve as a backup in case of a failure of the TCD repository.

Appliance Repository Mirror A mirror of the appliance repository has been installed in GRNET using the automated installation process supported by the project. The service is deployed on a VM running on GRNET's support infrastructure. A total storage of 1 TB has been allocated from the local storage server in order to accommodate VM images. Currently the repository is mirrored from TCD once every day.

Creation of Standard Base Images Three standard base images were created for the first public StratusLab release. These were images for ttylinux 9.5, CentOS 5.5, and Ubuntu 10.04. The procedure for generating these images has also been documented on the StratusLab website as part of the user tutorial. These images are currently available from the appliance repository and can be used to instantiate VMs on the StratusLab reference cloud service.

3.4.2 Issues and Corrective Actions

The provisioning of physical infrastructure was impacted during the reported period by a number of unscheduled and particularly long downtimes at the GRNET site. The downtimes were required to perform various maintenance activities of the physical infrastructure. These activities included the recabling of the physical nodes interconnection and the upgrade of the networking infrastructure (routers and switches) in order to deal with some serious firmware bugs affecting the network equipment. These downtimes affected the testing process of the StratusLab distribution and slightly delayed the date of the first official release of the software.

Apparently, these circumstances are rather rare and the corrective actions performed by the datacenter administrators are expected to fix the various problems in the physical setup of the infrastructure. Nevertheless, in order to avoid similar situations in the future, the StratusLab team negotiated with the datacenter the process to be followed in the future. In particular it was agreed that:

- StratusLab will be identified as one of the datacenter's official users and will be notified well beforehand about downtimes. The datacenter support help desk has been notified for this and the GRNET operations team has been included in the relevant mailing lists.
- In case these downtimes can be postponed and if the project's current status demands such a delay, the StratusLab team will prevent the downtimes till the issues imposed by the project are resolved (e.g. planned release of a new version).
- A specific window of maintenance has been set in which the datacenter administration can schedule a downtime. Specifically, there are two windows on Monday and Wednesday morning from 07:00-09:00 CET.

So far the above actions seem to have brought the expected results and the number and length of downtimes have decreased. Obviously these are major issues that will be followed closely in the coming period in order to ensure the provision of envisioned Quality of Service from the project's operations activity.

3.4.3 Use of Resources

During the first two quarters of the project, WP5 experienced a significant under-spending compared with the planned usage of resources. One of the reasons for this has been the delay from all the partners of WP5 for what concerns the assignment of new staff and the overall allocation of human resources. In other cases work was performed without consuming financial resources from the project as a consequence of the late arrival of the prefinancing. This is one reason why no deviation from planned work was experienced during the first months of the project. On the contrary there have been significant achievements that have established a solid basis for future work. Already in the second quarter the rate of spending has increased significantly and has reached to large percentage the expected levels of consumption. The situation is expected to normalize during the Q3 as most of the personnel is now in place and the prefinancing of the project has been received.

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

WP6 investigates and develops services for the innovative automatic deployment and dynamic provision of grid services as well as scalable cloud-like management of grid site resources. The main result in the second quarter has been the deliverable D6.1 “Cloud-like Management of Grid Sites 1.0 Design Report”, where an extension of the initial StratusLab architecture including these innovative services has been defined. This document has established the starting points for development: cloud-like APIs, service definition language and contextualization, scalable cloud frameworks and monitoring and accounting solutions. For scalability, a service manager, Claudia, will be included in the StratusLab distribution; the code has been moved into a StratusLab repository. A set of specific initial uses cases have been identified to test the StratusLab distribution, for example, Torque, Sun Grid Engine (SGE), and a dynamically-managed grid computing element.

3.5.1 Progress Towards Objectives by Task

3.5.1.1 T6.1: Dynamic Provision of Grid Services

Service Manager For the dynamic provision of grid services, WP6 has introduced a layer on top of current IaaS clouds that allows users to manage a service (ensemble of machines) as a single entity. This service manager, Claudia, analyzed in D6.1, has been moved into the StratusLab repository and will appear in upcoming releases. This solution provides a wider range of scalability mechanisms and a broader set of actions that can be undertaken (addition, removal, reconfiguration, federation, etc.) than on a simple IaaS cloud. Claudia can work on top of several different cloud infrastructure providers.

Service Language and Contextualization For the service definition language, Open Virtualization Format (OVF) has been chosen, since it provides a standard way for describing service and virtual machines as well as the networks involved in the service. Furthermore, as virtual machines need contextualization information, (that is, configuration information passed at boot time) some mechanisms from OpenNebula and some OVF recommendations (for instance ISO images) have been introduced.

Cloud-like APIs The usage of standard APIs has been identified as an important point for StratusLab. Thus, TCloud and OCCI are the APIs that will provide access to the Service Manager and Virtual Machine Manager, respectively.

Monitoring and Accounting Regarding monitoring and accounting, the extension of OpenNebula for both monitoring (integration with Ganglia) and accounting has been identified as a better solution.

Identification of Use Cases D6.1 has identified a set of use cases for the WP6 development based on end-user’s requirements. Concretely, these are Sun Grid

Engine (a service that has already been tested with the service manager and all the required artifacts exist), Torque, and gLite services.

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

Service-level open-source elasticity frameworks Grid applications deployed over cloud technologies should benefit from scalability at the service level, which conceals low level details from the user. WP6, in Q2, has selected Claudia as a service manager because it fulfills the requirements and to take advantage of TID's experience with it. Claudia is an advanced service management toolkit that allows service providers to dynamically control the service provisioning and scalability in an IaaS Cloud. Claudia manages services as a whole, controlling the configuration of multiple VM components, virtual networks and storage support by optimizing the use of them and by dynamically scaling up/down services applying elasticity rules, SLAs and business rules.

3.5.2 Issues and Corrective Actions

No major issues related to WP6 have arisen in Q2.

3.5.3 Use of Resources

WP6 started in M4, so no effort was expended in Q1. Regarding the effort spent in Q2, WP6 has significantly underspent comparing with the planned usage of resources. This is somewhat expected as this activity is in the startup phase. Most of the work has revolved around defining and use cases as well as finding agreement on the final architecture of the StratusLab distribution. It is expected that the effort will increase significantly as development begins.

4 Project Management

4.1 Consortium

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

The effort consumed by partner and by work package are shown in Tables 4.1 and 4.2, respectively. See the “Issues” section for a discussion of the lower expended effort than expected.

4.2 Management Tasks

Meetings Table 4.3 contains a list of the meetings 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 4.4 contains the metrics for the project. Because of continued discussion within the project and changes in the development priorities, these have been slightly modified and expanded to provide a better picture of the project’s progress. 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. Deliverables and milestones D3.1, D6.1, D5.2, MS2, and MS7 have been produced in this quarter.

Table 4.1: *Effort (in Person-Months) by Partner*

Partner	Q1	Q2	Total	Expected	Diff. (%)
CNRS	6.84	6.61	13.45	20.25	-34
UCM	4.40	5.50	9.90	14.54	-32
GRNET	1.75	6.68	8.43	14.21	-41
SIXSQ	5.10	3.29	8.39	12.75	-34
TID	0.10	2.70	2.80	10.29	-73
TCD	1.50	3.50	5.00	6.00	-17

Table 4.2: Effort (in Person-Months) by Work Package

WP	Q1	Q2	Total	Expected	Diff. (%)
WP1	0.98	0.93	1.91	3.00	-36
WP2	5.38	4.46	9.84	12.00	-18
WP3	0.70	3.92	4.62	8.75	-47
WP4	9.00	5.29	14.29	24.00	-40
WP5	3.63	8.43	12.06	21.00	-43
WP6	0.00	5.25	5.25	9.29	-43

Letters of Support The project has provided letters of support for four EC project proposals aiming to provide domain-specific computing infrastructures. All expressed an interest in cloud technologies in general and StratusLab products in particular. Should they be approved, the project will work with those projects in order to broaden our user communities and increase our impact. A letter of support for a commercial tender potentially using StratusLab software was not provided. The Project Management Board clarified the policy for letters of support in the future to make the process as clear and as efficient as possible.

Memoranda of Understanding It is expected that we will want to conclude Memoranda of Understanding (MoUs) with other projects to clarify how we can collaborate on common interests and tasks. The Project Management Board approved a process which involves PMB approval for opening and concluding an MoU negotiation. The negotiation itself will be managed by WP3 with input from all of the activities.

4.3 Issues

Delayed Hiring As identified in Q1 there has been a delay in hiring personnel for the project because of the late distribution of the pre-financing. To date, this has not significantly affected the work plan of the project; however for the next quarter (creating a feature-complete release) it will be imperative that the project is at full strength. Fortunately, hiring has taken place at many of the partners and new people will start 1 January 2011.

4.4 Planning

4.4.1 Objectives for Next Quarter

- Continued dialog with and support of targeted communities,
- Increasing visibility of project by targeted communities,
- Regular public releases concluding with functionally complete beta,

- Production grid site running over a stable StratusLab cloud, and
- Integration of the service manager into the distribution.

4.4.2 Roadmap

At the Face-to-Face meeting in Lyon on 15-16 November 2010, the collaboration discussed the high-level features to be included in the 1.0 release of the StratusLab cloud distribution (due in M12). Tasks related to these features will be included in the integration sprints. A public release of the StratusLab cloud distribution will be made every other sprint (around every 6 weeks) as a compromise between release overheads (documentation, release notes, changelogs, etc.) and development.

The varied infrastructures provided by the project will each have its own upgrade cycle. Test infrastructures will be updated most frequently following developments within a sprint. Preproduction infrastructures will follow the most recent public release. Production infrastructures will be updated when benefits outweigh inconveniences. The project will support the latest development release and the release installed on the production infrastructure.

The following paragraphs provide more details on the high-level functional features that are planned for the 1.0 release of the StratusLab cloud distribution. Other outcomes of the Face-to-Face meeting are discussed in the summaries of the individual activities.

Interfaces/APIs For the management of virtual machines, the project will evolve from OpenNebula's XMLRPC interface to OCCI. The service manager will additionally provide a TCloud interface that allow the specification and management of ensembles of machines. The service manager will eventually use the OCCI interface to the underlying IaaS cloud. In parallel, interfaces will be developed for the image registry and repository. Storage interfaces are discussed below.

Service Manager Claudia will be integrated into the StratusLab cloud distribution to allow the specification and management of ensembles of machines. This will enable automated and dynamic control of resources for the specified service. Initial use cases will be for deployments of Sun Grid Engine and Torque. These will lead into specification of grid computing elements and eventually the other grid services.

Security All exterior interfaces will be encrypted and secured via protocols like https and ssh. Internal interfaces will be blocked from exterior access via file-based sockets or firewall rules. The authentication and authorization mechanism will allow use of the grid security model with access control via subjects (DNs), groups, and roles. Appropriate security policies will be developed in collaboration with EGI.

Accounting/Monitoring There are three levels of monitoring: physical infrastructure, virtual machines, and grid services. The grid services will be monitored with the standard grid tools. The other two will use Ganglia as a backbone with

custom Ganglia probes for virtual machine information. Accounting information will be collected through OpenNebula (possibly via Ganglia). Resources are more rich in a cloud environment and will require extension of accounting databases used on EGI.

Appliance Repository Evolution The current appliance repository will be split into two services: an appliance marketplace (or registry) and an appliance store. The marketplace will allow people to register metadata concerning shared machine (and disk) images; it will allow that information to be searched as well. An appliance store contains the actual machine or disk image and may be, for example, a simple web server. An appliance store will be maintained for images generated by StratusLab. These services will be integrated with the machine deployment mechanisms on the StratusLab cloud.

Storage Services The project will not provide a specific service for file-based storage. It will provide services for the use of read-only disk images (e.g. distribution of standard databases) and persistent (read-write) disk images (e.g. storing of service state and logs). The OCCI and CDMI interfaces will be studied to see which is most appropriate.

Networking The project will initially implement a mechanism for allocating three types of network addresses to virtual machines: public, local, and private. Private addresses are NATed to the outside and can only be contacted from the physical host, local addresses are NATed to the outside but can be contacted by virtual machines running on the site, and public addresses that can be contacted from the WAN. In the future, the configuration and allocation of virtual networks will be investigated.

Table 4.3: Meetings

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
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 4.4: Metrics

Metric	Q2	Q3	Q4	Y1 Target	Q5	Q6	Q7	Q8	Y2 Target
No. of people on StratusLab announcement list	67			25					75
Registered users on StratusLab discussion site	N/A			50					100
No. of views of website	2922			–					–
No. of completed sprints	5			–					–
No. of releases	1			–					–
No. of open user stories	38			–					–
No. of implemented user stories	69			–					–
No. of open bugs	6			–					–
No. of fixed bugs	7			–					–
No. of prod. sites running StratusLab dist.	1			5					10
No. of sites exposing the cloud API	1			0					5
Availability of sites	N/A			80%					95%
Reliability of sites	N/A			80%					95%
No. of VOs served via StratusLab sites	0			10					30
No. of sci. disciplines served via StratusLab sites	0			3					7
Delivered CPU	N/A			–					–
Delivered CPU through cloud API	N/A			–					–
Storage used	N/A			–					–
Storage used through cloud API	N/A			–					–
No. of sites providing scale-out	0			–					–
Fraction of resources by scale-out of a site	0			–					–
No. base machine images	5			5					10
No. of base machine image downloads	783			–					–
No. appliances	0			5					15
No. of appliance downloads	0			–					–

5 Deliverables and Milestones

Tables 5.1, 5.2, and 5.3 list the deliverables for Y1, deliverables for Y2, and milestones, respectively. The deliverables and milestones for this reporting period were D3.1, D6.1, D5.2, MS2, and MS7.

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		WP6	TID	P	PU	PM11				
D2.2	Report on Evaluation of StratusLab Products		WP2	CNRS	R	PU	PM12				
D3.2	Report on Dissemination, Collaboration and Standardization Activities		WP3	TCD	R	PU	PM12				
D3.3	Exploitation and Sustainability First Plan		WP3	TCD	R	PU	PM12				
D4.2	StratusLab Toolkit 1.0		WP4	SIXSQ	P	PU	PM12				
D4.3	First Year Software Integration Report		WP4	SIXSQ	R	PU	PM12				
D5.3	First Year Infrastructure Operations Report		WP5	GRNET	R	PU	PM12				
D6.3	First Year Cloud-like Management of Grid Sites Research Report		WP6	TID	R	PU	PM12				

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				
D4.4	Reference Architecture for StratusLab Toolkit 2.0		WP4	SIXSQ	R	PU	PM15				
D6.4	Cloud-like Management of Grid Sites 2.0 Design Report		WP6	TID	R	PU	PM17				
D5.4	Economic Analysis of Infrastructure Operations		WP5	GRNET	R	PU	PM18				
D6.5	Cloud-like Management of Grid Sites 2.0 Software		WP6	TID	P	PU	PM23				
D2.4	Final Report on StratusLab Adoption		WP2	CNRS	R	PU	PM24				
D2.5	Report on Evaluation of StratusLab Products		WP2	CNRS	R	PU	PM24				
D3.4	Final Review of Dissemination, Collaboration and Standardization Activities		WP3	TCD	R	PU	PM24				
D3.5	Exploitation and Sustainability Final Plan		WP3	TCD	R	PU	PM24				
D4.5	StratusLab Toolkit 2.0		WP4	SIXSQ	P	PU	PM24				
D4.6	Software Integration Final Report		WP4	SIXSQ	R	PU	PM24				
D5.5	Infrastructure Operations Final Report		WP5	GRNET	R	PU	PM24				
D6.6	Cloud-like Management of Grid Sites Research Final Report		WP6	TID	R	PU	PM24				

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			
MS10	Initial virtual appliance repository	WP5	GRNET	PM9			
MS14	Release of Cloud-like Management of Grid Services and Resources 1.0 Beta	WP6	TID	PM9			
MS8	Release of StratusLab 1.0 Beta	WP4	SIXSQ	PM10			
MS11	Operation of Site Running StratusLab toolkit v1.0	WP5	GRNET	PM10			
MS4	Adoption of StratusLab Software by External Grid Sites	WP2	CNRS	PM14			
MS12	Delivery of Virtual Appliance Repository	WP5	GRNET	PM18			
MS5	Opening of Virtual Appliances Repository to External Application Communities	WP2	CNRS	PM20			
MS15	Release of Cloud-like Management of Grid Services and Resources 2.0 Beta	WP6	TID	PM21			
MS9	Release of StratusLab 2.0 Beta	WP4	SIXSQ	PM22			
MS13	Operation of Site Running StratusLab Toolkit v2.0	WP5	GRNET	PM22			

Glossary

Appliance	Virtual machine containing preconfigured software or services
CDMI	Cloud Data Management Interface (from SNIA)
DCI	Distributed Computing Infrastructure
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
NFS	Network File System
NGI	National Grid Initiative
OCCT	Open Cloud Computing Interface
Public Cloud	Cloud infrastructure accessible to people outside of the provider's organization
Private Cloud	Cloud infrastructure accessible only to the provider's users
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