



HAL
open science

Quarterly Report (QR3)

Mohammed Airaj, Marc-Eliau Bégin, Christophe Blanchet, Javier Fontan, Evangelos Floros, Eduardo Huedo, Stuart Kenny, Ignacio Llorente, Charles Loomis, Jose Lopez, et al.

► **To cite this version:**

Mohammed Airaj, Marc-Eliau Bégin, Christophe Blanchet, Javier Fontan, Evangelos Floros, et al..
Quarterly Report (QR3). 2010. hal-00687989

HAL Id: hal-00687989

<https://hal.science/hal-00687989>

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 QR3 (V1.0)
14 March 2011

Abstract

In the third quarter, the project has built on the first public release (v0.1), providing additional functionality leading up to the 1.0 production release expected at the end of Q4. An additional public, preview release (v0.2) was produced and a subsequent release (v0.3) is expected just after the close of Q3. The releases are progressing well, with v0.3 expected to be nearly feature-complete, lacking only storage functionalities. The project is well-positioned to release the StratusLab v1.0 distribution at PM12 with the complete set of expected features.



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



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

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

Copyright © 2011, Members of the StratusLab collaboration: Centre National de la Recherche Scientifique, Universidad Complutense de Madrid, Greek Research and Technology Network S.A., SixSq Sàrl, Telefónica Investigación y Desarrollo SA, and The Provost Fellows and Scholars of the College of the Holy and Undivided Trinity of Queen Elizabeth Near Dublin.

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



Contributors

Name	Partner	Sections
Mohammed Airaj	CNRS-LAL	WP2
Marc-Elian Bégin	SIXSQ	WP3, WP4
Christophe Blanchet	CNRS-IBCP	WP2, WP3
Vangelis Floros	GRNET	WP3, WP5, WP6
Javier Fontan	UCM	WP4
Eduardo Huedo	UCM	WP6
Stuart Kenny	TCD	WP5
Ignacio Llorente	UCM	WP3
Charles Loomis	CNRS-LAL	WP2, WP3, WP5, Summary, Mgt.
Jose Lopez	TID	WP2, WP3, WP4, WP6
David O'Callaghan	TCD	WP3

Document History

Version	Date	Comment
0.1	22 Feb. 2011	Initial outline of report.
0.5	11 Mar. 2011	Complete contents for internal review.
0.6	11 Mar. 2011	“Final” version for internal review.
1.0	14 Mar. 2011	Final version.

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	14
2.1 Objectives	14
2.1.1 Quarter 1	14
2.1.2 Quarter 2	14
2.1.3 Quarter 3	15
2.2 Review Recommendations	15
3 Progress and Achievements	16
3.1 WP2: Interaction with Targeted Communities	18
3.1.1 Progress Towards Objectives by Task	18
3.1.2 Issues and Corrective Actions	21
3.1.3 Use of Resources	21

3.2	WP3: Dissemination	22
3.2.1	Progress Towards Objectives by Task	22
3.2.2	Issues and Corrective Actions	24
3.2.3	Use of Resources	24
3.3	WP4: Software Integration and Distribution	26
3.3.1	Progress Towards Objectives by Task	26
3.3.2	Issues and Corrective Actions	29
3.3.3	Use of Resources	29
3.4	WP5: Infrastructure Operation	30
3.4.1	Progress Towards Objectives by Task	30
3.4.2	Issues and Corrective Actions	34
3.4.3	Use of Resources	34
3.5	WP6: Innovative Cloud-like Management of Grid Services and Resources	35
3.5.1	Progress Towards Objectives by Task	35
3.5.2	Issues and Corrective Actions	36
3.5.3	Use of Resources	36
4	Project Management	37
4.1	Consortium	37
4.2	Management Tasks	37
4.3	Issues	38
4.4	Planning	38
4.4.1	Objectives for Next Quarter	38
4.4.2	Roadmap	39
5	Deliverables and Milestones	43

List of Figures

3.1 Visits for Q3.	22
----------------------------	----

List of Tables

1.1	StratusLab Information and Support	12
1.2	StratusLab Partners	13
3.1	Talks	23
3.2	WP5 Infrastructure Services	30
4.1	Effort (in Person-Months) by Partner	37
4.2	Effort (in Person-Months) by Work Package.	38
4.3	Meetings (Q1-Q2)	40
4.4	Meetings (Q3).	41
4.5	Metrics	42
5.1	Deliverables (Year 1).	44
5.2	Deliverables (Year 2).	45
5.3	Milestones	46

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 third quarter, the project has built on the first public release (v0.1), providing additional functionality leading up to the 1.0 production release expected at the end of Q4. An additional public, preview release (v0.2) was produced and a subsequent release (v0.3) is expected just after the close of Q3. The releases are progressing well, with v0.3 expected to be nearly feature-complete, lacking only storage functionalities.

Increased Visibility Interest in the project continues to increase through the efforts of all partners. A number of researchers have contacted the project and are trying out the StratusLab software via the reference infrastructure. Two particularly interesting collaborations have been with the maintainer of the ttylinux distribution to make it more cloud-friendly and with the HEPiX Virtualization Working Group concerning the Marketplace.

Marketplace Sharing machine and disk images is critical for the adoption of cloud technologies on shared, distributed infrastructures. The project has redesigned the appliance repository concept, resulting in the StratusLab Marketplace—a repository of machine and disk image metadata—and a set of associated tools. The tools allow site policies to be applied to requested images and conveniently downloaded. Eventually, cloud-based storage will be used for storing and sharing images. A prototype version of the Marketplace is currently available (MS10).

Specialized Appliances Specialized appliances have been created for bioinformatic applications (MS3) and for grid services. The bioinformatics appliance will free researchers from the drudgery of installing by hand the large number of required applications for their research, thus streamlining their use of the cloud and research. Similarly, the grid images will make it easier for sites to create a working grid resource center, avoiding all of the usual pitfalls involved in deploying grid services by hand.

Certified Grid Site The grid appliances have been used to deploy a standard grid site. This site has been integrated with EGI and is subject to the same operations checks as for other sites. This site has passed all of the certification criteria and is now a recognized, certified site within EGI. The number of accepted VOs is

currently limited, but will expand as we build confidence in the site.

Agile Processes The Scrum philosophy adopted by the project for development and integration is now an integral part of the project, with regular cycles of planning meetings, daily standup meetings, and demonstrations taking place. A much improved continuous integration server (reconstructed after a crash) more fully tests the StratusLab components and facilitates the public releases.

Usability Improvements Numerous improvements have been added to the latest release (v0.3) based on user feedback. These include a more robust and easier to use client, quarantine of virtual machine images for forensic analysis of potentially compromised machines, new authentication modules supporting a wide range of different credentials, and automated creation of new machine images from existing ones.

Claudia Integration The build of the Claudia service manager has been integrated with the standard build procedures of the project. Packages generated via the automated integration procedures have been tested and installed and will appear in the v0.3 release which will occur just after the close of Q3.

OpenNebula Improvements During Q3 several new features has been developed in OpenNebula to address some of the requirements identified in this period, namely: integration with cluster monitoring systems (e.g. Ganglia) and new fault tolerance capabilities to recover from physical host or virtual machine failures. This features can be previewed in v0.3 and will be fully integrated in upcoming releases.

The project's cloud distribution continues to develop largely according to the foreseen plan. The only critical functionality missing from the distribution at the end of Q3 is a solution for cloud storage. This will be a focus of upcoming sprints to ensure that the 1.0 release in May 2011 will be truly complete.

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 re-

sources 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

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

information about their contributions to the project. Table 1.2 contains the list of StratusLab partners and relevant contacts.

Table 1.2: StratusLab Partners

CNRS	Centre 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-Elian BEGIN meb@sixsq.com
TID	Telefónica Investigación y Desarrollo SA	Jose LOPEZ josemll@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

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.

2.1.3 Quarter 3

In the third quarter, the primary objective was to provide a feature-complete release and demonstrate its utility for running grid services. The detailed objectives were:

- 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.

Nearly all of these objectives have been achieved with the StratusLab v0.3 release at the end of the quarter. One highlight is the certification of a production grid site running over the StratusLab distribution. The v0.3 release is nearly functionally-complete, missing only a solution for storage. This will be developed early in Q4.

2.2 Review Recommendations

Not yet applicable.

3 Progress and Achievements

In the third quarter, the project has built on the first public release (v0.1), providing additional functionality leading up to the 1.0 production release expected at the end of Q4. An additional public, preview release (v0.2) was produced and a subsequent release (v0.3) is expected just after the close of Q3. The releases are progressing well, with v0.3 expected to be nearly feature-complete, lacking only storage functionalities.

Increased Visibility Interest in the project continues to increase through the efforts of all partners. A number of researchers have contacted the project and are trying out the StratusLab software via the reference infrastructure. Two particularly interesting collaborations have been with the maintainer of the ttylinux distribution to make it more cloud-friendly and with the HEPiX Virtualization Working Group concerning the Marketplace.

Marketplace Sharing machine and disk images is critical for the adoption of cloud technologies on shared, distributed infrastructures. The project has redesigned the appliance repository concept, resulting in the StratusLab Marketplace—a repository of machine and disk image metadata—and a set of associated tools. The tools allow site policies to be applied to requested images and conveniently downloaded. Eventually, cloud-based storage will be used for storing and sharing images. A prototype version of the Marketplace is currently available (MS10).

Specialized Appliances Specialized appliances have been created for bioinformatic applications (MS3) and for grid services. The bioinformatics appliance will free researchers from the drudgery of installing by hand the large number of required applications for their research, thus streamlining their use of the cloud and research. Similarly, the grid images will make it easier for sites to create a working grid resource center, avoiding all of the usual pitfalls involved in deploying grid services by hand.

Certified Grid Site The grid appliances have been used to deploy a standard grid site. This site has been integrated with EGI and is subject to the same operations checks as for other sites. This site has passed all of the certification criteria and is now a recognized, certified site within EGI. The number of accepted VOs is currently limited, but will expand as we build confidence in the site.

Agile Processes The Scrum philosophy adopted by the project for development and integration is now an integral part of the project, with regular cycles of planning meetings, daily standup meetings, and demonstrations taking place. A much improved continuous integration server (reconstructed after a crash) more fully tests the StratusLab components and facilitates the public releases.

Usability Improvements Numerous improvements have been added to the latest release (v0.3) based on user feedback. These include a more robust and easier to use client, quarantine of virtual machine images for forensic analysis of potentially compromised machines, new authentication modules supporting a wide range of different credentials, and automated creation of new machine images from existing ones.

Claudia Integration The build of the Claudia service manager has been integrated with the standard build procedures of the project. Packages generated via the automated integration procedures have been tested and installed and will appear in the v0.3 release which will occur just after the close of Q3.

OpenNebula Improvements During Q3 several new features has been developed in OpenNebula to address some of the requirements identified in this period, namely: integration with cluster monitoring systems (e.g. Ganglia) and new fault tolerance capabilities to recover from physical host or virtual machine failures. This features can be previewed in v0.3 and will be fully integrated in upcoming releases.

The project's cloud distribution continues to develop largely according to the foreseen plan. The only critical functionality missing from the distribution at the end of Q3 is a solution for cloud storage. This will be a focus of upcoming sprints to ensure that the 1.0 release in May 2011 will be truly complete.

3.1 WP2: Interaction with Targeted Communities

The number of researchers and groups who express interest in the StratusLab distribution continues to increase through the efforts of the WP2 activity. Most of these researchers have requested access to the project's reference infrastructure and are using that platform for porting their applications to a cloud-based infrastructure. Q4 will see a formal evaluation of the StratusLab distribution based on the experiences of those using the reference infrastructure now. In addition to scientific and commercial users, there has also been a productive collaboration with the maintainer of the ttylinux distribution in order to make it more useful in a cloud setting. WP2 has contributed significantly to the design of the StratusLab Marketplace for shared machine and disk images along with developing tools to allow site policy enforcement and download of those images. The activity continues to promote collaboration with the HEPiX Virtualization Working Group to ensure that the StratusLab distribution can meet their needs and increase the chance of adoption of StratusLab at grid resource centers.

3.1.1 Progress Towards Objectives by Task

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

Virtual Spain TID is contacting the Spain-funded Virtual Spain project (CENIT research programme), which aims at enriching interactive multimedia services with geographical data (satellite images, digital terrain models, etc). TID attended the project's Technical Committee in early December. The grid-related part of the project researches the adaptation of geospatial data processing flows to massive computation platforms based on cloud computing. The project is aware of Stratuslab and the potential of the software being developed.

LAPP The Laboratoire d'Annecy-le-Vieux de Physique des Particules (LAPP) is a CNRS particle physics laboratory located in Annecy, France. They had proposed a cloud-based infrastructure for the laboratory and the surrounding university campus, which was unfortunately not directly funded. However, they view cloud technologies as a natural complement of their existing grid infrastructures and wish to continue to work towards such an infrastructure on a best-effort basis. LAL has discussed in detail how they could take advantage of the StratusLab distribution to simplify their grid administration and to provide cloud-based resources to their users. LAPP already has extensive experience with VMware and is willing to test the StratusLab distribution over VMware. This would help the project in providing direct feedback on another underlying virtualization technology.

ASSYST Meeting on Cloud Computing Following a presentation of StratusLab at the ASSYST Meeting on Cloud Computing in Paris, we have been contacted by a Portuguese researcher who is interested in porting his master/worker calculation framework, Conillion, to the cloud. WP2 will help him to do this porting using the reference infrastructure.

ttylinux The ttylinux distributions are small, but fully functional, linux distributions that are typically used for embedded systems. Images based on these distributions are extremely useful for testing the StratusLab cloud and for training. Scientists may also be interested in these images because of their extremely fast start-up times. LAL had a strong collaboration with the maintainer of the ttylinux distributions in order to make it more appropriate for use in a cloud. Essentially all of the customization that StratusLab had been doing for making cloud ttylinux images have been adopted by the maintainer and now appear in the standard ttylinux distributions. This makes it easier for the project (and others) to keep pace with the new ttylinux releases. StratusLab will prepare new images based on the updated ttylinux distributions in Q4.

HEPiX Virtualization Working Group Significant work in Q3 was related to the collaboration with the HEPiX Virtualization Working Group concerning the Marketplace and trusting of machine images through signed metadata. There are ongoing discussions with this group to allow for interoperable metadata descriptions between various services. There is an interest on their part to use the StratusLab Marketplace if it can fulfill their requirements.

IIT Vietnam An engineer from the Institute of Information Technology in Vietnam visited LAL for a week to learn about the StratusLab distribution. During this week, this engineer learned how to use the StratusLab reference infrastructure, how to create new machine images, how to upload them to the image repository, how to deploy grid services within the cloud, and how to install a StratusLab cloud. In addition, the engineer also started porting a drug-discovery scientific application based on the WISDOM platform. This collaboration is continuing after his return to Vietnam; other scientists from his institute will come for training in Q4.

NIIF The National Information Infrastructure Development (NIIF) Program is the framework for the national research network in Hungary. They have already developed a working cloud prototype for their users. In particular, they have based their image repository infrastructure around iSCSI technologies. LAL discussed various points for collaboration, in particular around using these storage technologies in the StratusLab distribution. Providing storage services will be a major focus for Q4, so it is expected that technical collaboration will begin then.

Virtual Spain TID is contacting the Virtual Spain project (CENIT research program), which aims to enrich interactive multimedia services with geographical data (satellite images, digital terrain models, etc.). A part of the project researches the adaptation of geospatial data processing flows to massive computation platforms based on cloud computing. The project is aware of StratusLab and the potential of the software being developed.

NUBA TID is contacting the Spanish NUBA project (Avanza research program), which aims to develop a multi-vendor federated IaaS platform for corporations. The project team takes into consideration the use of the Stratuslab tools to deploy their testbed.

TID Private Cloud TID is discussing with the in-house Private Cloud project on how the Service Manager can benefit from the tasks that are being carried out in StratusLab. Although in the early stages, the development team is very interested in collaborating and expressed their interest in areas such as monitoring and scalability applied to the use cases of grid services.

French Bioinformatics RENABI IBCP has created two customized machine images for the bioinformatics community: “biological databases repository” and “bioinformatics compute node”. The “biodata repo” VM aims to provide users with access from any cloud node to international reference databases recording biological resources such as protein or gene sequences and associated data, protein structures, or complete genomes. This appliance acts as a proxy between the internet where all the reference databases are published and the cloud internal virtual nodes that will compute the bioinformatics analyses. The “biocompute node” VM has pre-installed bioinformatics software such as ClustalW, BLAST, FastA and SSearch. Because these methods require access to reference data for processing, this appliance is linked via an NFS mount to the “biodata repo” appliance.

ELIXIR IBCP has participated to an ELIXIR workshop about Bioinformatics Infrastructures. StratusLab was presented as a possible solution for bioinformatics distributed infrastructures, especially for the web services and portal interfaces deployed by IBCP and for the French distributed infrastructure set-up by RENABI GRISBI. Attendees are very interested in StratusLab. Contact with institutes Germany (Rost Lab, Munich), Denmark (CBS, Copenhagen) and Netherlands (CMBI, Nijmegen) has been made. One member of the Rost Lab who is in charge of cloud evaluation has already requested for an account on the StratusLab Reference Infrastructure.

3.1.1.2 Task 2.2: Intensive Evaluation of StratusLab Products

Documentation The activity lead the effort to restructure on the online documentation of StratusLab and will continue to work to keep it up to date with the project’s preview releases.

Semi-Production Use of StratusLab LAL has deployed a StratusLab cloud for use by the computer services section of the laboratory. This is used to test and to deploy laboratory services to gauge how well the cloud paradigm works in a production setting. Several bugs have been found and corrected (e.g. incorrect reporting of the number of virtual CPUs). Also several feature requests have come out of this work: need for group machine management, “tags” of virtual machines to easily identify them, and finer control over the resource allocations for a particular machine instance.

Evaluation of StratusLab Manual Installation IBCP is deploying a StratusLab cloud for use by the bioinformatics section of the laboratory. The manual installation procedure is used in order to be strongly evaluated. In case of deployment of StratusLab clouds in other bioinformatics Labs such as the French RENABI ones,

there will be no expertise about Quattor usage. Then the manual installation will be very probably the most used one in the Bioinformatics community, at least for the first times. Several bugs have been found and most were fixed, e.g. the daemons httpd and oned were not configured to start at boot time (chkconfig on), mis-configuration of the passwords between the different files (jetty-7 login.properties, oned one_auth) or problem in the configuration file ('app_repo.use_ldap' in stratuslab.cfg).

Marketplace As part of the collaboration with the HEPiX Virtualization Working Group, LAL has contributed extensively to the design and development of the StratusLab Marketplace. The design has been captured in the Marketplace technical note. LAL has also created scripts to allow for policy enforcement and downloading of images based on image metadata. These will be critical for the integration of the Marketplace into the distribution and allowing cloud administrators better control over the images running on their infrastructure.

3.1.2 Issues and Corrective Actions

No major issues related to the WP2 work plan have arisen in Q3.

3.1.3 Use of Resources

WP2 continues to use less effort than foreseen in the project's Technical Annex, even though this has not impacted the activity's program of work. Upcoming evaluations and an expected increase in users and administrators are expected to increase the effort that is invested in this activity and close the gap with respect to the Technical Annex.




Figure 3.1: Visits for Q3.

3.2 WP3: Dissemination

The project continued to increase its visibility among European projects. The dissemination effort has been primarily concerned with announcing development releases of the software, and with preparing for the EGI User Forum 2011. Project partners continue to give talks to promote and explain StratusLab.

Collaboration activity has increased through involvement with SIENA, EGI-InSPIRE and other DCI projects, and Memoranda of Understanding are in the preparation or exploration phases in some cases.

3.2.1 Progress Towards Objectives by Task

3.2.1.1 Task 3.1: Dissemination

Release Dissemination The release dissemination plan devised for the first software development release in the previous quarter was updated for version 0.2 and 0.3. Preparation has begun for the release of version 0.3, due early in Q4.

Media & Publications The StratusLab grant is acknowledged in a book chapter “Architectures for Enhancing Grid Infrastructures with Cloud Computing” by Eduardo Huedo, Rafael Moreno-Vozmediano, Ruben S. Montero and Ignacio M. Llorente (of StratusLab participant UCM) in *Grids, Clouds and Virtualization* (Springer, 2011). An initial, internal draft of an academic position paper on StratusLab was prepared.

Website The project website (<http://www.stratuslab.eu>) was updated to allow visitors to provide comments. This will be used initially to solicit feedback on the project roadmap. The website content management system was also upgraded to get the latest features and security updates. Some design work has begun to improve the appearance and usability of the project website for visitors.

Figure 3.1 shows the number of visits to the website. The number has increased from 2922 in Q2 to 4623 in Q3 (+58%).

Online, the project’s Twitter account has been used to announce development progress such as planning meetings, demos and releases.

Table 3.1: Talks

Title / Event	Date
“StratusLab: The European Initiative to Bring Cloud to Grid Infrastructures” (R. S. Montero, UCM) at Spanish NGI meeting, Barcelona, Spain	2010-12-02
“StratusLab : Le projet et sa distribution cloud” (C. Loomis, LAL) at France-Grille Cloud Event in Lyon, France	2010-12-13
UCM talk at France-Grille Cloud Event in Lyon, France	2010-12-13
“Cloud Computing – Anatomy and Practice” (M.-E. Bgin, SixSq) at <i>Groupe romand des utilisateurs/trices de GNU/Linux et de Logiciels libres</i> , Morge, Switzerland	2011-01-11
“StratusLab Cloud Distribution” (C. Loomis, LAL) “Opening the Cloud” 2011 ASSYST Meeting on Cloud Computing, Paris, France	2011-01-31
UCM talk at Cloud Expo Europe 2011, London, United Kingdom	2011-02-03
“Bioinformatics distributed infrastructure, services and cloud computing” (C. Blanchet, CNRS IBCP) at <i>ELIXIR Workshop - Bioinformatics Infrastructures</i> , Amsterdam, Netherlands	2011-02-22

EGI User Forum 2011 The project has booked an exhibition booth at EGI User Forum 2011 – to be held in Vilnius, Lithuania – and WP3 is coordinating the preparation of dissemination materials for the event including t-shirts, flyers, and demonstrations.

GRNET has submitted an abstract to the conference describing the release infrastructure provided for interested users. LAL will provide training on StratusLab as part of the ‘Heavy User Communities’ training sessions at the event.

Project partners will participate in the virtualization strand of the conference program.

Talks A brief list of talks describing the project delivered during Q3 is given in Table 3.1. Details and links, where available, are given on the project website at <http://stratuslab.eu/doku.php/presentations>.

3.2.1.2 Task 3.2: Collaboration with Standards Bodies and Related Projects

SIENA Project partners prepared for the upcoming *Cloudscape III* meeting (March 15–16) organized by SIENA. Vangelis Floros will represent StratusLab.

TID wrote several sections on the SIENA Initiative (<http://www.sienainitiative.eu/>) wiki, which has been set up specifically for the writing of the *SIENA European Roadmap on Grid and Cloud Standards for e-Science and Beyond*. The sections aim at disseminating the activities carried out in Stratuslab, and consisted of an introduction to the project, a description of challenges and a description of engagement. In addition to that, deliverable D6.1 was uploaded as a reference document

for StratusLab's *Cloud-like management of grid sites design*.

Memoranda of Understanding StratusLab has taken initial steps towards establishing a Memorandum of Understanding with EGI. The project has begun exploring the possibility of MoUs with other projects: namely EGI, IGE, EMI, EDGI, ERINA+.

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

StratusLab was represented at an EGI-InSPIRE Security Policy Group meeting in Amsterdam in January 2011. StratusLab input will be included in a number of security policies which touch on virtualization and cloud computing. StratusLab may be invited to participate formally in the Security Policy Group once the MoU is in place.

EMI Since Q2 Vangelis Floros (GRNET) has engaged an informal interaction with EMI project. This was initiated by two members of EMI, Shahbaz Memon and Bjorn Hagemeyer (both from the Julich Supercomputing Center) mainly with the purpose for EMI to keep up to date with the developments of StratusLab integration activities. During the past quarter there have been various communications between the two parties answering questions and clarifying specific technical details about the StratusLab cloud solution. In addition a user account has been created for Shabaz Memon in GRNET's reference cloud infrastructure providing access to the test installation and tools of StratusLab. This informal collaboration is expected to be formally defined in the coming months in the context of the MoU with EMI which is currently under preparation.

OpenNebula Discussion has begun on promoting new StratusLab features and services (such as enhanced security, and the under-development appliance marketplace) to the OpenNebula community through the mutual project partners.

Venus-C UCM participated in a Venus-C project meeting at Aachen Microsoft Innovation to study integration between both projects.

FP7 Proposals TID was approached by Fraunhofer Fokus (Germany) expressing interested in future collaborations in cloud-related projects. Fraunhofer and TID participated in former project consortia.

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 Q3.

3.2.3 Use of Resources

Effort decreased slightly in Q3 with respect to previous quarters. This is probably due to the holiday period during this quarter. Effort is expected to increase in

Q4 when Task 3.3 will begin and the deliverable reports D3.2 and D3.3 will be prepared.

3.3 WP4: Software Integration and Distribution

During Q3, WP4 continued integrating features resulting in the production of StratusLab distributions: v0.2 and v0.3 (although the actual release of v0.3 officially took place early March). During this, we improved the build, test and release procedures. This also included devising an upgrade procedure, such that an infrastructure running a previous versions of StratusLab could be easily upgraded. This work has been performed in tight collaboration with WP5.

The Scrum process adopted in Q1 and improved in Q2 is now fully integrated, as can be seen from the regular sprints, demos and planning meetings that are now taking place, as well as the releases that are being produced on average every two sprints (approximately every six weeks). All the Scrum events are now part of fabric of the project and most partners contribute actively.

The Hudson continuous integration server was considerably improved following a hardware failure that forced us to rebuild it. It now includes machines from GRNET and LAL, including machines being re-imaged at LAL daily, providing clean machines on which to perform automatic tests.

As defined during the Scrum planning meetings, WP4 focused its integration effort on the following services:

- Marketplace
- Authentication proxy
- Claudia
- Automatic image creation
- Policy enforcement (e.g. quota, Marketplace images)
- Quarantine

WP4 also simplified the configuration file required for the manual installation as well as completing the ability for Quattor to manage an entire StratusLab installation.

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 Q3. However, the architecture was revisited to include the proxy service which provides a single point for performing authentication. The Marketplace was also integrated as planned in the architecture. 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 is now published on our website to collect comments from our community.

3.3.1.2 Task 4.2: Integration of Open-source Distribution

As mentioned in the introduction, SixSq, GRNET and LAL have continued to improve the Hudson continuous integration server. The Hudson infrastructure is now composed of over five machines (three at GRNET and two at LAL). The LAL machines are managed by Quattor, which allows us to re-image the machines every day (at 2:00 in the morning), from which an automatic installation, configuration and end-to-end test is conducted. Additionally, the testing infrastructure for the OpenNebula component at UCM has been improved with two new slave nodes and specific test jobs for the development branches. This contributes significantly in building confidence that StratusLab remains functional with each code commit.

Several services and components, such as the command-line client, the proxy and the Marketplace, now include unit tests, which are automatically executed by Hudson on each code commit. The end-to-end smoke tests have also been extended to cover more features and key functionality.

From v0.2, StratusLab includes a new authentication service, referred to as proxy server, which provides a single authentication point for the entire StratusLab distribution. A new OpenNebula driver was developed to disable the OpenNebula authentication feature, removing duplication of configuration information, without compromising security.

A quarantine feature was also developed to quarantine each virtual machine, such that post-mortem forensic can be performed on suspicious instances. This is an important feature to investigate security issues.

Significant work took place in improvements in the monitoring system of the distribution. This included a partial integration of Ganglia with OpenNebula and the installation system. In particular, the host CPU load now takes into account the hypervisor overhead. The VM network metrics (received and transmitted bytes) are now obtained for the KVM hypervisor. Additionally a new system probe to gather the disk free space of each host is now in place and ready for integration and testing.

UCM has setup a git repository that tracks changes in the mainstream source tree of OpenNebula and also includes specific patches and developments for StratusLab. This allows StratusLab to benefit more quickly from improvements and bug fixes to OpenNebula, instead of having to wait for public releases.

Several bugs in OpenNebula were fixed, including: shared file system issues, special OpenNebula configurations and wrong handling of remote action scripts, partial URL encoding and decoding of usernames when using certificates for authentication. In preparation for an elastic-IP like feature, OpenNebula now supports the ability to dynamically add and remove leases from a virtual network. These features could be integrated and leveraged in future versions of the StratusLab distribution.

During this quarter, Claudia (the Service Manager component developed by TID) was integrated in the project's continuous integration system, with the help of CNRS/LAL and SixSq. Claudia's source code was organized to comply with

the policies shared by the rest of the components (e.g. using the same version of libraries across components). New jobs were created in Hudson so that compilation, deployment and documentation are automated sharing a common strategy with the rest of the components. Python scripts were written to install, configure and start Claudia services and integrated with the Stratuslab command-line tools.

3.3.1.3 Task 4.3: Contextualization of Grid Services

The generic contextualization mechanism devised during Q1 and improved during Q2 was further enhanced during this quarter. From v0.2 onwards, StratusLab supports three network levels: public, local and private. The local network can now be NATed using the StratusLab front-end machines using the manual installation configuration file, making it possible for private system administrators to deploy StratusLab on private resources with no public IPs.

Following the agreed new metadata for describing machine images for the Marketplace, the StratusLab base images were re-created and simplified. These base images can also be used as base images for creating new custom images.

More grid images were created, such that an entire grid site can run on a StratusLab infrastructure. In support to VMs running grid services, instances can be started requesting a specific IP address, such that server certificates remain valid across VM restarts.

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. Further, as planned in the program of work, WP5 is taking over more of the first line support, as WP5 gains in operational knowledge in the distribution.

WP4 has continued to support and manage our software development procedure based on Scrum. All main Scrum events are now routinely taking place (daily stand-up, demo and planning meetings), to which the large majority of partners are taking part regularly.

The user-stories and tasks selected during the planning meetings include a balance of bug fixes and new features, such that current versions are maintained, while new versions are being prepared. Several of the patches fixing OpenNebula bugs have been integrated in new versions of OpenNebula, such that these are not required in the StratusLab code-base.

The technical support is now extended to a number of wiki pages on our website, as well as a FAQ page to which WP4 contributes as recurrent questions and issues are raised by members and users. The wiki pages have been maintained and updated prior to each release.

TID provided the project's website (<http://stratuslab.eu/doku.php/claudia>) with technical support on Claudia (the Service Manager). This information covers the following subjects:

- Claudia Architecture

- How to compile Claudia and to install Claudia Platform
- Install Claudia from deb and rpm packages
- Claudia User Manual

3.3.2 Issues and Corrective Actions

Following a major hard-disk failure, part of the Hudson server was lost for several days. Since this is a critical part of our release procedure, for building and testing each release, we also decided to insert a small two weeks sprint focused on finalizing and testing release 0.3. During that time, the server was rebuilt, which was also an opportunity to considerably improve the system, now including Quattor controlled machines. The new system now includes a backup process.

While our implementation of Scrum already works very well, it would perform even better if all the partners were to take part more regularly to the different events.

3.3.3 Use of Resources

The effort involved in the activity is nominal for this quarter and the work is progressing according to the plan.

Table 3.2: WP5 Infrastructure Services

Reference Cloud Service	GRNET	https://cloud-grnet.stratuslab.eu:2634/RPC2
Pre-production Infrastructure	GRNET	https://62.217.120.158:2634/RPC2
Pre-production Infrastructure	LAL	https://onehost-4.lal.in2p3.fr:2643/RPC2
Project Tools (Hudson Server)	GRNET	http://hudson.stratuslab.eu:8080
Appliance Repository	TCD	http://appliances.stratuslab.eu
App. Repository Mirror	GRNET	http://appmirror-grnet.stratus.eu/images
Prototype Marketplace implementation	TCD	http://appliances.stratuslab.eu/marketplace/
Test Infrastructure	LAL	https://onehost-2.lal.in2p3.fr:2643/RPC2

3.4 WP5: Infrastructure Operation

WP5 is responsible for the provision of the computing infrastructure required by the various activities of the project. During the third quarter of the project, we achieved noticeable progress with all of the tasks planned by the activity. Among the highlights of the past quarter is the certification of the first virtualized production grid site, the expansion of the cloud infrastructure, the migration to StratusLab release 0.2, the preliminary work on the Marketplace and the provision of first grid and bioinformatic-specific appliances. Table 3.2 summarizes the services that currently offer web-based access in the context of WP5.

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 The production cloud service in GRNET evolved and expanded during the reporting quarter attracting at the same time new external users. The service was upgraded to StratusLab 0.2 a few days after this was released. Due to significant changes that this revision introduced, the service had to be re-installed from ground up. In normal cases this would cause a major disruption in the hosting services but since this was still in a very early stage of the project no production services, hosted in running VMs, were impacted. As the project progresses though, and the infrastructure is used for hosting production-level applications and grid sites, it will be imperative to streamline the upgrade progress as much as possible making it transparent to the end users. This requirement has already been conveyed to WP4 and was taken under consideration during the production of release 0.3 expected at the beginning of Q4.

The physical infrastructure hosting the project's reference cloud service, was also expanded significantly during the past quarter, with the addition of 6 more nodes. Thus the total capacity offered is 17 physical nodes (1 frontend and 16

hosting nodes) providing a total of 256 CPU cores and 768 GBytes of total main memory (48 GBytes per node).

For what concerns external users, 10 more user accounts were created in response to requests from people outside the project who got interested in StratusLab and wanted to test-drive the tools integrated by the project. For the time being the reference cloud service still depends on a username/password authentication scheme, but with the advent of 0.3 we plan to adopt the more secure certificate-based authentication mechanism implemented in this latest release.

The project decided to adopt the EGI Acceptable Usage Policy (<https://documents.egi.eu/public/ShowDocument?docid=74>) which defines the rights and obligations of all StratusLab users related to the usage of the offered cloud services. The web link to the above document is communicated to every new user upon the creation of their account informing them that the usage of the service implies the acceptance of the EGI AUP conditions.

Certification of a production Grid Service During the first six months we experimented extensively with the installation and operation of grid sites on top of cloud services using pre-configured VMs. In this quarter we decided to take this effort one step further and formalize the installation of a virtualized grid site by certifying it within the EGI infrastructure. In parallel we prepared a number of VM appliances for the basic node roles of a gLite-based grid site and namely: the Computing Element, the Storage Element, the Worker Node, the User Interface and the APEL service. All of the above-mentioned images are available from the appliance repository. The appliances follow the evolution of gLite middleware; with every new release a new image snapshot is created and is uploaded on the repository.

The certified grid site is named HG-07-StratusLab. The site was certified within the GRNET NGI (the Greek National Grid Initiative) and has joined the national grid infrastructure (HellasGrid). The site offers a CE and 8 dual-core WNs thus providing a total capacity of 16 cores for job submission. The site also supports submission of MPICH-2 and OpenMPI parallel jobs. Communication among the nodes for MPI execution is supported through ssh host-based authentication. Each WN is configured with 2GB of main memory. The site also comprises a SE that offers a total storage space of 2TB. It should be noted that this storage is configured directly as an NFS mount-point from the local storage server and is not yet virtualized (e.g. it cannot be managed as a EBS service from the StratusLab command line tools).

The site currently supports the StratusLab VO (vo.stratuslab.eu) as well as the required EGI-wide and local ops (operations) VOs. Obviously the job processing capacity of the site is currently rather limited. In this first phase the site primarily serves as a testbed for grid-cloud interoperability tests and for evaluating the implications of operating grid sites on public cloud services. Already a number of issues have been identified which will be classified and reported in the coming quarter. In the future depending on the workload and potential requests to support additional VOs it should be rather trivial to expand the workload execution capacity of the

site (i.e. number of available cores and/or WNs).

The GStat page with all the details of the site as they are reported from the Site-BDII are available at <http://gstat-prod.cern.ch/gstat/site/HG-07-StratusLab/>.

Pre-production Services LAL has deployed an initial pre-production cloud service which is available to the laboratory's users and system administrators. This service will be opened to the wider community once registration procedures are in place and the StratusLab release provides a mechanism for enforcing site machine image policies.

GRNET has also deployed a pre-production site dedicated for beta-testing and validation of StratusLab releases. The site is comprised of 3 nodes (1 frontend and 2 hosting nodes). Access to the site is granted only to a few system administrators and developers within the project.

Support Infrastructure The Hudson continuous integration service had to be relocated to a new node, within the datacenter, in order to perform required maintenance tasks to the previous hosting node (see "Issues and Corrective Actions" section for more details). Three servers have been allocated to WP4 for development testing. Additionally, two servers have been allocated to WP6 for the Claudia integration tasks.

3.4.1.2 Task 5.2: Testing of the StratusLab Toolkit

Testing Infrastructure Work was done to more completely automate the installation of the StratusLab installation with Quattor. This includes a new configuration module for the authentication proxy and an updated module for OpenNebula. With the deployment of the pre-production cloud service at LAL, the test infrastructure will be regularly reinstalled from scratch with the latest software to ensure that the distribution installs and functions correctly.

Validation of StratusLab Releases A validation procedure for new releases of the StratusLab distribution has been defined and reported in project deliverable D5.2 "Infrastructure Tools and Policy Specifications", released in M7. According to this procedure each new candidate release is first tested in the pre-production infrastructure and then, if accepted, is deployed in the reference cloud service. This procedure enforces the close interaction between WP4 and WP5 in order to promptly resolve any issues concerning the stability and validity of new releases. The procedure was applied during the release of StratusLab 0.3 which is due in Q4.

Investigation of Storage Solutions Cloud storage services will be added to the StratusLab distribution shortly. LAL has investigated using iSCSI technologies as a component of those storage services. An iSCSI target (server) was deployed. Disks residing on this server could then be mounted by hosts in the cloud and used as input devices for virtual machines running on the host. No obstacles were encountered to using this as part of the StratusLab cloud storage services.

GRNET also investigated the usage of NFS as a shared storage solution among the frontend and the hosting nodes. The migration to a shared storage solution

seems inevitable at the moment since it will allow us to take complete advantage of the physical nodes processing capabilities and also will give us the ability to perform live migration of VM instances among the hosting nodes. The current reference could service resides on SSH transfers for sharing images among the hosting nodes. The nodes themselves are limited to a total storage space of 80 GB. Currently all VM images are copied locally in every host before instantiation, this space quickly gets filled, especially when large images are used (e.g. a 10GB CentOS base image currently offered from the appliance repository). The storage server in GRNET infrastructure currently support only NFS so the usage of this protocol is the only solution even though it does not scale very well. Nevertheless, for this size of infrastructure (17 nodes) this approach should be sufficient. For larger infrastructures technologies like iSCSI, currently tested by LAL, should be considered more appropriate.

3.4.1.3 Task 5.3: Virtual Appliances Creation and Maintenance

Appliance Repository During this quarter the initial version of the appliance repository, using a WebDAV enabled Apache web server, has been maintained. The initial set of reference images has been updated as required, and a new set of images has been made available. These include the first *appliances*. Five grid appliances have been provided containing gLite grid middleware: a Computing Element, a Storage Element, a Worker Node, an APEL node and a User Interface. Additionally, in collaboration with WP2 activity a Bioinformatics specific appliance has been created based on a CentOS 5.5 base image.

The operation of the first version of the repository was marked by the delivery of milestone MS10 - "Initial Virtual Appliance Repository", which was prepared during month PM9 as planned.

In this quarter the focus of the task has been on the design, and initial implementation of the next version of the repository, the 'Marketplace'. The Marketplace will serve as a registry for images that can be shared. Rather than providing a centralized storage location, the storage of the actual images will be handled by the owners of the images, allowing them to control access to the images if desired. The Marketplace itself will contain cryptographically-signed metadata about machine and disk images allowing users to find existing images and allowing system administrators to define policies on trusting those images.

A reference implementation of the Marketplace will be made available for testing at the beginning of the next quarter.

Creation of Standard Base Images New base images were made available during this period. After user requests a 10 GB CentOS 5.5 image was generated providing larger storage space on the root file system. Also in response to security issues (see below) related to ttylinyx 9.5 images, a new image based on version 9.7 of the distribution was created fixing the security holes of its predecessor.

3.4.2 Issues and Corrective Actions

Reaction to Security Incident During the past quarter we have experienced two security incidents related with vulnerable VM images.

A machine running on the LAL preproduction cloud service was hijacked and used for a password scanning attack on another site. The forensic analysis showed that the method of entry was use of an image with a commonly-known root password. This analysis also highlighted deficiencies in the available logging information. Changes to OpenNebula and the authentication proxy were made to ensure that all necessary tracing information is conveniently available to system administrators. In addition, a quarantine mechanism was implemented that keeps terminated instances for 48 hours (by default) so that detailed forensic analysis can be done afterwards on any image.

A similar incident occurred also in a VM running the in the reference cloud service. Again an instance was hijacked using a commonly known username/password combination in the ttylinux 9.5 distribution. The instance was brought off-line immediately and the faulty image was removed from the appliance repository.

Both incidents demonstrated the problems of public cloud infrastructures related with the vulnerabilities of VM images. Apart from reacting promptly to such events we believe that the introduction of the Marketplace and the security mechanisms implemented within will improve the overall security of the service. Nevertheless, security will remain one of the critical challenges for the project operations activity and will in our focus for improvement in the coming months.

Disruption of support services The Hudson continuous integration service was disrupted just before the release of StratusLab 0.3 due to a failed hard disk in the hosting node. This delayed the release process since the server had to be re-installed and configured from scratch. The server employed a RAID-1 mirroring configuration which means that the hosting data were not lost, but due to lack of an on-line back-up procedure, recovering of the information took longer than expected. For this reason it was decided to put in place a backup mechanism in the hudson server which will store daily snapshots of the system data. This way the time to bring the service back on line will be much shorter should a similar incident occur again.

3.4.3 Use of Resources

After the significant underspending experienced in the first months of the project, resource utilization has reached to the expected level during QR3. Actually the current pace of spending effort is slightly above the initially projected average quarterly level, thus we expecte that by the end of the project all resources will be consumed according to the DoW.

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 third quarter has been the development of new functionalities in the Service Manager (Claudia) and the Virtual Infrastructure Manager (OpenNebula). Claudia has been provided with an extended lifecycle engine and a more flexible rule language. The new rule definition language will allow for improved scalability in the management of grid services. OpenNebula was integrated with a monitoring tool and fault tolerance mechanisms were developed to reduce service downtime.

The work done in Q3 led to the achievement of milestone MS14 *Release of Cloud-like Management of Grid Services and Resources 1.0 Beta*.

3.5.1 Progress Towards Objectives by Task

3.5.1.1 T6.1: Dynamic Provision of Grid Services

Lifecycle engine Claudia, the Service Manager developed by TID, has now an improved lifecycle engine that supports a broader set of actions over the Virtual Machines.

Rule language Regarding the service definition language, Open Virtualization Format (OVF) was extended with Rule Interchange Format (RIF), a W3C standard, to support complex scalability rules.

Requirements and use case definition Since the project has a grid site officially running and maintained by GRNET, the team is starting a phase for establishing requirements and use cases to be covered in the next quarters.

Integration of the Service Manager Claudia was added to the project's continuous integration system and is ready to be part of the Stratuslab distribution/release. Installation and configuration mechanisms were prepared during this quarter, and more effort on that will be done in the next quarter in order to further automate Claudia's installation and configuration.

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

Improved service-level elasticity Service-level elasticity will be brought to grid applications by Claudia, the Service Manager. During this quarter, Claudia went through a series of changes to increase the flexibility of the rule engine and accept complex scalability rules. Rule Interchange Format (RIF) was chosen as the preferred language for rule description. A RIF parser was implemented so that Claudia supports RIF rules embedded in the OVF file. Services that scale up demanding many Virtual Machines benefit from load balancers. For that reason Claudia was integrated with a load balancer to be ready for high load situations.

Monitoring with Ganglia OpenNebula is now able to obtain monitoring information from Ganglia. GRNET and UCM wrote custom Ganglia probes that were installed on each host in order to expose hypervisor information to the Ganglia monitoring tool. Large deployments may benefit from the scalability of Ganglia and use it as the monitoring source both for physical hosts and VMs. Moreover, the administrator is now able to have a real time graphic representation of the resource consumption using the Ganglia Web Frontend.

Fault tolerance Fault tolerance has been improved to automatically trigger recovery actions when a physical host or VM fails. When OpenNebula detects that a host is down, a hook can be triggered to deal with the situation. This can be very useful to limit the downtime of a service due to a hardware failure, since it can redeploy the VMs on another host.

3.5.2 Issues and Corrective Actions

No major issues related to WP6 have arisen in Q3.

3.5.3 Use of Resources

The effort spent in Q3 is close to the planned usage of resources. New developments got off to a good start and effort increased significantly. It is expected that this effort reach a steady state in the next quarters. These are the spent effort figures in person-months by the partners involved in WP6: UCM 3.99, GRNET 0.52, TID 4.29. This adds up to a total 8.80 person-months.

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 Tables 4.3 and 4.4 contain 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 4.5 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. Milestones MS3, MS10, and MS14 have been produced in this quarter.

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

Partner	Q1		Q2		Q3		TOTAL		
	Actual	Exp.	Actual	Exp.	Actual	Exp.	Actual	Exp.	Diff. (%)
CNRS	6.84	10.13	6.61	10.13	7.49	10.13	20.94	30.38	-31
UCM	4.40	5.13	5.50	9.41	6.93	9.41	16.83	23.95	-30
GRNET	1.75	6.75	6.68	7.46	7.40	7.46	15.83	21.68	-27
SIXSQ	5.10	6.38	3.29	6.38	4.14	6.38	12.53	19.13	-34
TID	0.10	3.00	2.70	7.29	7.30	7.29	10.10	17.57	-43
TCD	1.50	3.00	3.50	3.00	3.30	3.00	8.30	9.00	-8

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

WP	Q1		Q2		Q3		TOTAL		
	Actual	Exp.	Actual	Exp.	Actual	Exp.	Actual	Exp.	Diff. (%)
WP1	0.98	1.50	0.93	1.50	0.55	1.50	2.46	4.50	-45
WP2	5.38	6.00	4.46	6.00	4.96	6.00	14.80	18.00	-18
WP3	0.70	4.38	3.92	4.38	2.94	4.38	7.56	13.13	-42
WP4	9.00	12.00	5.29	12.00	8.03	12.00	22.32	36.00	-38
WP5	3.63	10.50	8.43	10.50	11.28	10.50	23.34	31.50	-26
WP6	0.00	0.00	5.25	9.29	8.80	9.29	14.05	18.57	-24

Letters of Support The project has provided letters of support for a proposed Slovenian e-Infrastructure project. This project would include cloud computing capabilities and, if funded, would be an excellent candidate for using the Stratus-Lab cloud distribution.

Memoranda of Understanding The PMB has defined the procedure for negotiating Memoranda of Understanding and given its approval to start with EGI, EMI, IGE, EDGI, and ERINA+.

4.3 Issues

Underused Effort The effort numbers for Q3 are much closer to the expected expenditures than in previous quarters. Notable exceptions are for WP1 and WP2. In both cases, this gap is likely to be closed in Q4 as much more management effort and evaluation of the upcoming 1.0 release will be needed.

4.4 Planning

4.4.1 Objectives for Next Quarter

- In depth evaluation of the StratusLab v1.0 distribution,
- Increasing visibility of project by targeted communities and evaluation of our dissemination strategy,
- Continued public, preview releases culminating with the StratusLab v1.0 distribution,
- Continued operation of reference infrastructure and production grid site, and
- Complete integration and use of the service manager in the v1.0 release.

4.4.2 Roadmap

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

1. The tasks regarding having a public (user-visible) cloud and an associated cloud API have been moved from Y2 to Y1, largely because of interest from scientific communities and resource centers wanting to provide public clouds.
2. The tasks about hybrid clouds will be expanded to include also cloud federation models. This will be moved to Y2 to balance the change above. Also having a solid release will make these investigations easier.
3. As foreseen in the TA, the appliance repository consists of a single service that contains appliance metadata, appliance storage, and services for changing appliance formats. This has been split into different services. The Marketplace will handle appliance metadata. Storage will take place with normal cloud storage or outside of the cloud. Instead of providing a service for appliance format changes, client tools will be provided instead.

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

Table 4.3: Meetings (Q1-Q2)

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: 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 4.5: Metrics

Metric	Q2	Q3	Q4	Y1 Target	Q5	Q6	Q7	Q8	Y2 Target
No. of people on StratusLab announcement list	67	67		25					75
Registered users on StratusLab discussion site	N/A	N/A		50					100
No. of views of website	2922	4623		–					–
No. of completed sprints	5	5		–					–
No. of releases	1	1		–					–
No. of open user stories	38	72		–					–
No. of implemented user stories	69	40		–					–
No. of open bugs	6	15		–					–
No. of fixed bugs	7	11		–					–
No. of prod. sites running StratusLab dist.	1	1		5					10
No. of sites exposing the cloud API	1	1		0					5
Availability of sites	N/A	N/A		80%					95%
Reliability of sites	N/A	N/A		80%					95%
No. of VOs served via StratusLab sites	0	1		10					30
No. of sci. disciplines served via StratusLab sites	0	0		3					7
Delivered CPU	N/A	16 cores		–					–
Delivered CPU through cloud API	N/A	256 cores		–					–
Storage used	N/A	3 TB		–					–
Storage used through cloud API	N/A	N/A		–					–
No. of sites providing scale-out	N/A	N/A		–					–
Fraction of resources by scale-out of a site	N/A	N/A		–					–
No. base machine images	5	7		5					10
No. of base machine image downloads	783	2628		–					–
No. appliances	0	6		5					15
No. of appliance downloads	0	252		–					–

5 Deliverables and Milestones

Tables 5.1, 5.2, and 5.3 list the deliverables for Y1, deliverables for Y2, and milestones, respectively. The milestones for this reporting period were MS3, MS10, and MS14.

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	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			
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
OCCI	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