D-SA2.5: VERCE platform integration: final release report of integrated services and tools with evaluation of performance and quality

30/09/2015

Project acronym:	VERCE
Project n°:	283543
Funding Scheme:	Combination of CP & CSA
Call Identifier:	FP7-INFRASTRUCTURES-2011-2
WP:	WP6/SA2, Integration and evaluation of the platform services
Filename:	D-SA2.4.pdf
Authors:	S.H. Leong (BADW-LRZ) and S. Hachinger (BADW-LRZ)
Location:	`http://www.verce.eu/Repository/Deliverables/RP4`
Type of document:	Deliverable
Status:	Final
Due date of delivery:	30/09/2015
Reviewer:	H. Schwichtenberg (SCAI) and J.-P. Vilotte(IPGP)
Keywords:	Evaluation, Integration, Testing, Tools, Services, Application codes, Software components, Seventh release, Report, Performance, Quality

History	*Author*	*Date*	*Comments*
1	Siew Hoon Leong	11/08/2015	Initial draft for comments
2	Siew Hoon Leong	14/09/2015	Update of KPIs (Tools/Service/Software releases, Security and Quality of Service) using data provided by Andre Gemuend
3	Siew Hoon Leong Stephan Hachinger	15/09/2015 - 22/09/2015	Update charts of KPIs and added the appropriate evaluation comments
4	Siew Hoon Leong Stephan Hachinger	28/09/2015	Final corrections
5	Siew Hoon Leong	29/09/2015	Improvements: Feedback from Horst
6	Siew Hoon Leong	29/10/2015	Polishing and adding comments on sustainability

Contents

List of Figures

List of Tables

Executive Summary

One of the objectives of the VERCE project is to provide a service-oriented architecture and framework that wraps the data-infrastructure resources and services with a set of distributed data-aware Grid and HPC resources provided by the European e-Infrastructure and community. To this end, the tools, services and application codes, i.e. software components, which are particularly relevant to the seismologists and the Earth Science community, are selected for integration on the VERCE platform.

The main aim of this report is to report on the seventh release of the integrated tools and services. The Plan-Do-Check-Act (PDCA) cycle is used to manage the release process. This period corresponds to the seventh completed PDCA cycle and will be the final release by the VERCE project.

The seventh release successfully completed within the schedule timeframe. In this release, one component was evaluated and was approved as shown.

Approved

- ObsPy 0.10.1

The main focus of of all work-packages in this reporting period was to continue to support the provision of a working version of the VERCE Science Gateway for user evaluation. In order to support the sustainability of the Science Gateway during the transitional period after the project ends so as to facilitate a handover to EPOS, the final set of components on the platform is carefully evaluated and selected.

The lessons learned throughout the project have resulted in now the usual smooth and efficient evaluation cycle. Members of the JRAs were as usual providing strong support to enable the on-time completion of this evaluation cycle. An evaluation of the collected data based on the previously defined key performance indicators (KPIs) is presented. Finally, a summary table on the achievements of the original objectives for this work package is also included in this final report.

1. Seventh Release Report

The seventh PDCA cycle of the project, corresponding to the last release of integrated services and tools, was completed on 31 July 2015. As in the previous releases, a release management schedule was prepared to ensure that the process was clear to each involved work package. JRA1 submitted a tool request for evaluation. One tool was approved in this release. The details of this release are described in the following subsections.

1.1. Release Management Schedule

The work performed in this six months period is depicted in Figure 1. Learning from the experience, the JRAs and SAs were once again encouraged to submit their request form as early as possible in view of the summer vacation.

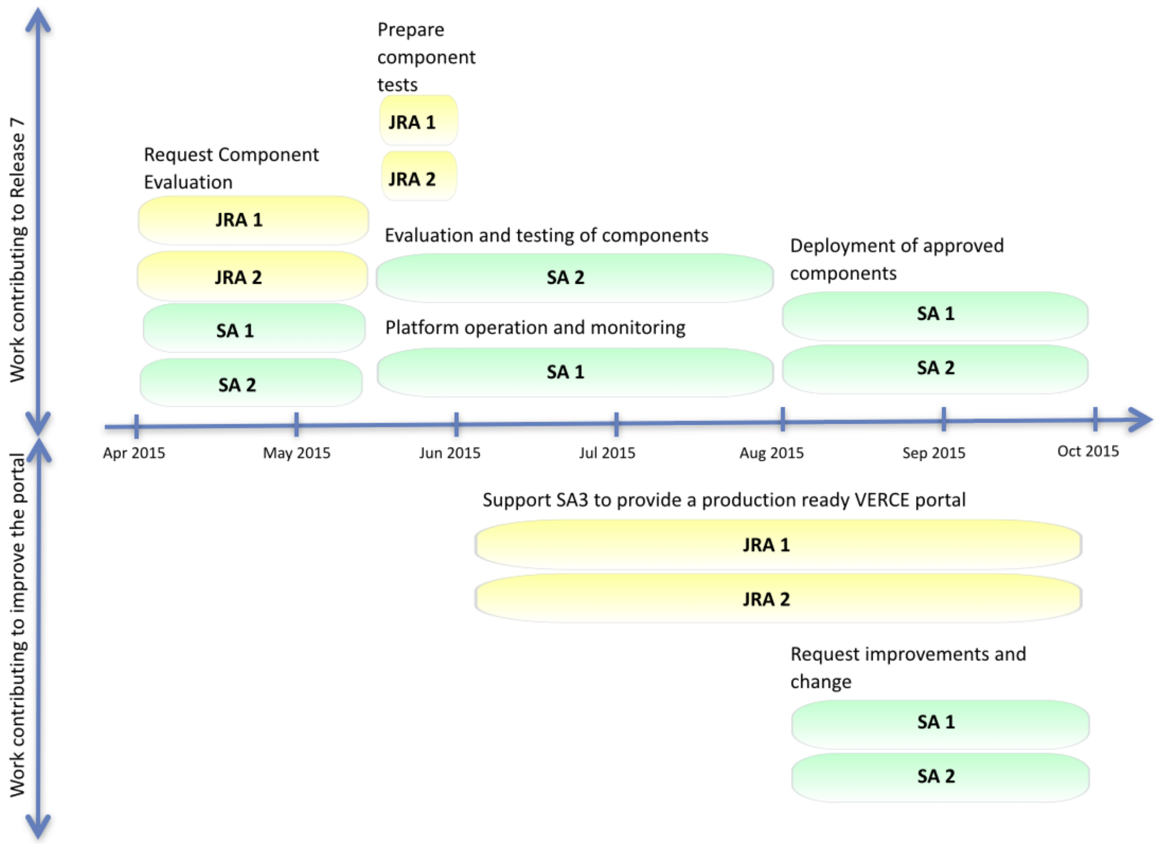

Figure 1 – Release Management Schedule (Apr 2015 - Sep 2015)

1.2. Requested components

One tool request was received from JRA1 as shown in the Table 1. Detailed information about each component is available internally on Redmine.

1.3. Assignment of Evaluators/Testers and Resources

The assigned resources and testers for each component are shown in Table 2

Component	Version	Type	Purpose	Submitted by
ObsPy[1]	0.10.1		Python framework for processing seismological data	JRA1

Table 1 – Requested Component for Evaluation and Testing

Component	Assigned Resource (SA1 Definition)	Assigned Tester
ObsPy	SuperMUC (HPC-LRZ-01)	LRZ: Siew Hoon Leong
	EGI Cluster (GRI-LRZ-02)	LRZ: Siew Hoon Leong
	Departmental Resource (DEP-SCAI- 01)	SCAI: Andre Gemünd
	Departmental Resource Cluster (DEP-UEDIN-01)	UEDIN: Amy Krause/Iraklis Klampanos

Table 2 – Assigned Resources and Testers of each component

The assigned resources are based on the profile of the components. Components that can potentially utilise HPC, GRID and Departmental resources are assigned to at least one of such resources for evaluation.

1.4. Evaluation and Testing

The evaluation and testing phase commenced in early May 2015 and was completed by the end of July 2015. In anticipation of the summer vacation, the team was encouraged to begin the evaluation process as early as possible. The detail of each specific component test is described in the following sub-section.

1.4.1. Component Specific Tests

ObsPy

ObsPy was recommended for an upgrade in this evaluation cycle. Significant number of new features were implemented in this new version to support the misfit analysis with dispel4py. dispel4py is a python based module with a strong dependency on ObsPy, which enables the description of misfit analysis as workflows. This new version of ObsPy is compatible with any Python 2.7.* version and thus no update of python is required. All required python modules, e.g. SciPy, NumPy and MatPlotlib, and ObsPy test suites were run to verify the correctness of the installation and its integration with the existing VERCE platform. It was noted that the older Matplotlib, version 1.2.1, on the platform is causing a memory leak issue on some systems. An update of Matplotlib to 1.4.3 was thus evaluated and finally also recommended.

1.5. Results and Recommendations

Approved components

The evaluated component is approved for release after the evaluation and testing phase. The approved component in the seventh release is:

1) ObsPy

ObsPy was evaluated on one Grid, one HPC and two departmental resources. Installation procedures have greatly been simplified due to our earlier feedback. Additionally the most complex step with regards to the installation of ObsPy was the installation of its dependency python libraries, in particular SciPy and NumPy. However, this dependency library update was not required in the release and thus the update on most systems went smoothly. The new departmental resource at SCAI required an installation from scratch. Correspondingly, some issues were faced due to version conflicts. Using the detail documentations from earlier releases, the issues were quickly ironed out. All tests were successful completed. The ObsPy upgrade is thus accepted and recommended for all resource types.

1.6. Documentation

All issues faced and solutions that the SA2 team encountered during this release cycle were documented in the SA2 wiki and in the request forms for the component. These documents contribute to the technical documentation that SA1 will use to coordinate the deployment of the approved components on the VERCE testbed. The documentation can be found on our internal Redmine under the heading "RP5 Results Issues and Documentations".

2. Evaluation of Performance and Quality

The evaluation of performance and quality is performed with the support from SA1 and SA3. Data collected by both WPs is shared with SA2 and the defined metrics are computed. The equations used to quantify these metrics can be found in D-SA1.5.

2.1. Availability of Services

Availability of VERCE services can be measured using the metrics defined under this section. This set of KPIs will provide an insight into how stable and reliable the VERCE platform is.

The availability of service data is collected from two sources, Inca and the logs from the services. The offered services are grouped into three categories as show below:

- Data Transfer (GridFTP)
- System Modules
- Other core services
  - Job Submission
  - iRODS
  - Scientific Gateway
  - Grid Middleware

The above services were used in production for both beta user evaluation and the training.

2.1.1. Service Availability

Figure 2, 3 and 4 show the service availability from July 2014 to July 2015. To ensure that our beta-users have a good overall experience, most components, i.e. services and tools, manage to improve their performance in 2015. Core client components, the system modules, science gateway and grid middleware, have availability of 90-100%. This implies a "Very Good" or "Excellent" QoS classification. More volatile services, e.g. GridFTP and iRODs, have mostly a "Good" to "Excellent" classification. One reason for the sudden drop in availability (both GridFTP and Job Submission) in June 2015 was found to be due to an internal issue with credential and VOMS in Inca. Consequently, these tests failed, i.e. the tests could not be carried out, across multiple resources that use this version of the credential. iRODS also experienced a reduction in availability due to an increased load with the introduction of beta-testers. The lack of good error logging within iRODS delayed the resolution of the problem.

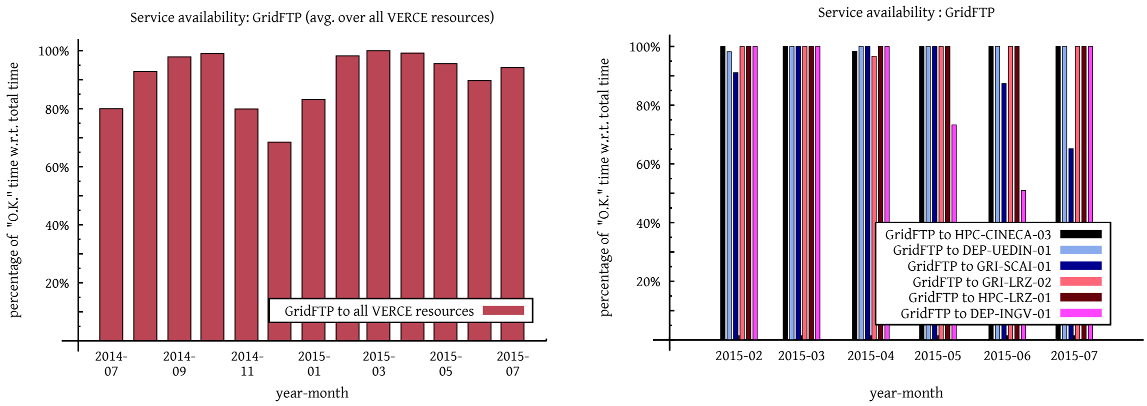

(a) Average Service Availability of Data Transfer (GridFTP) across all Resources

(b) Service Availability of Data Transfer (GridFTP) of each Resource

Figure 2 – Service Availability of Data Transfer (GridFTP)

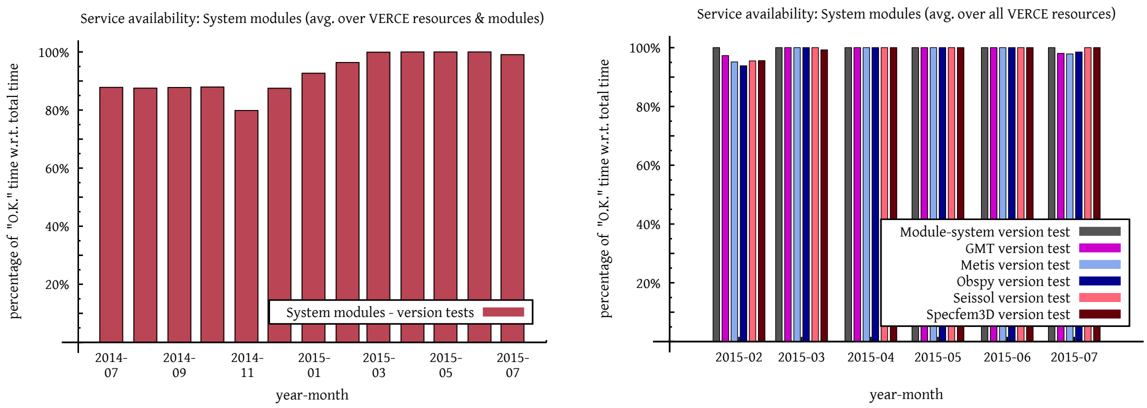

(a) Average Service Availability across all System Modules

(b) Average Service Availability of each System Module

Figure 3 – Service Availability of System Modules

2.1.2. Number of Service Interruptions

Figure 5, 6 and 7 show the number of service interruptions from July 2014 to July 2015. In February 2015, the integration of addition Inca reporters introduced a number of errors that contribute to the number of service interruptions across services monitored by Inca. Similar to service availability, the increase in service interruptions in June 2015 was also due to the VOMS credential issue in Inca. In July 2015, the "Job Submission" service interruption increment was due to a few unscheduled interruptions, i.e. cooling issues, of the departmental resource at INGV (not used by the beta users). In most cases. there are less than 5 per month in 2015. Crucial resources that are integrated with the Science Gateway generally have zero or a low number of service interruptions.

2.1.3. Duration of Service Interruptions

Figure 8, 9 and 10 show the duration of service interruptions from July 2014 to July 2015. It can be seen from the statistics that strong effort was made from the beginning of this year to reduce the duration of service interruptions upon the introduction of beta-users. Core resources that are available to the beta users, e.g. HPC-LRZ-01 and GRI-LRZ-01, are particularly ensured to have low duration of unscheduled interruptions. The spike in February 2015 in 9(b) was a result of the reported Inca reporter integration as

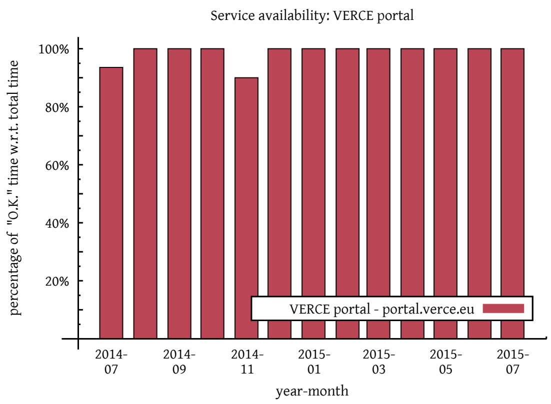

(a) Service Availability of Science Gateway

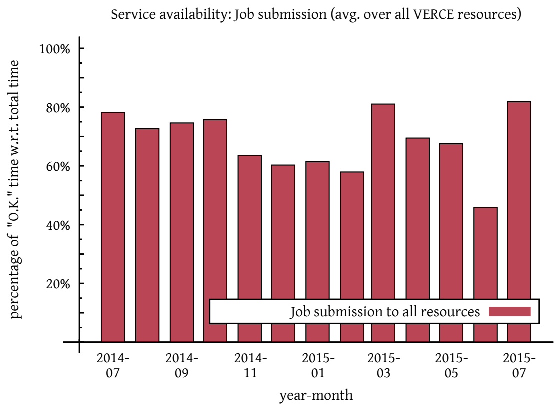

(b) Service Availability of Job Submissions

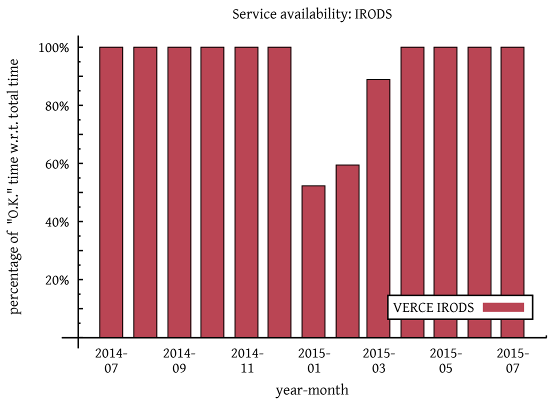

(c) Service Availability of iRODS

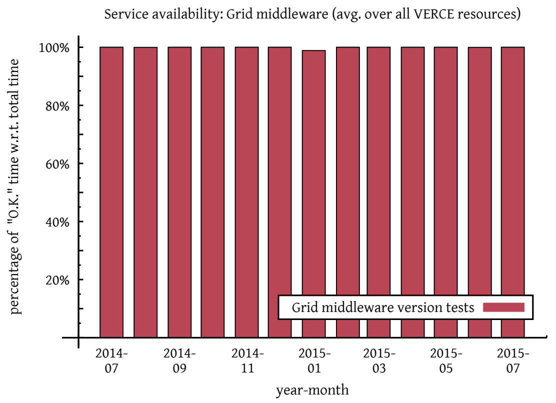

(d) Average Service Availability of Grid Middleware

Figure 4 – Service Availability of Core Services

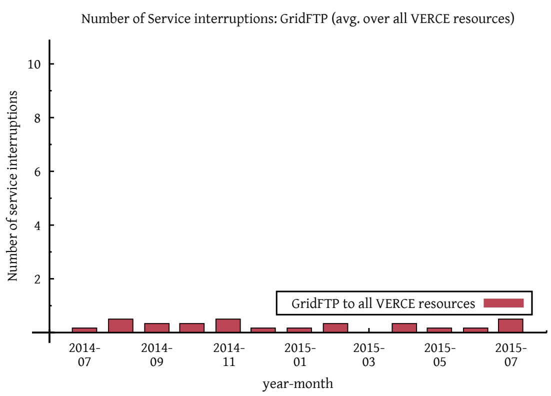

(a) Average Number of Service Interruptions of Data Transfer (GridFTP) Per Resource

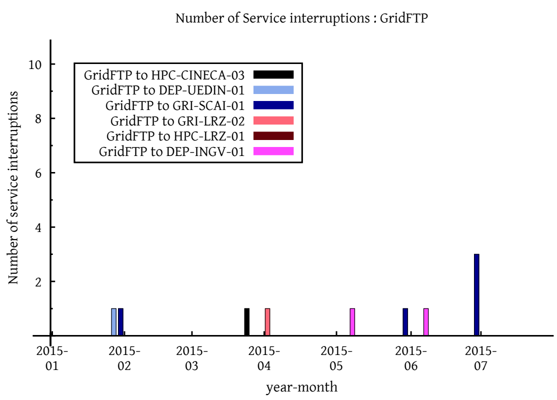

(b) Number of Service Interruption of Data Transfer (GridFTP) for each Resource

Figure 5 – Number of Service Interruptions of Data Transfer (GridFTP)

mentioned in the previous subsections.

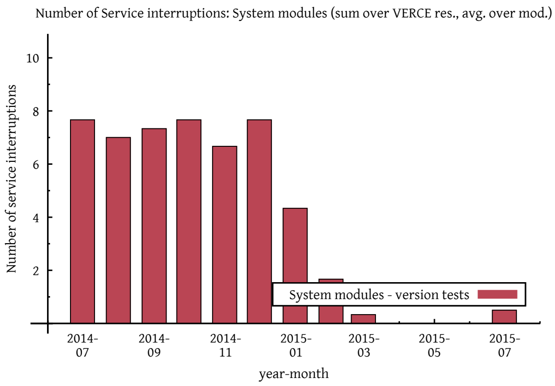

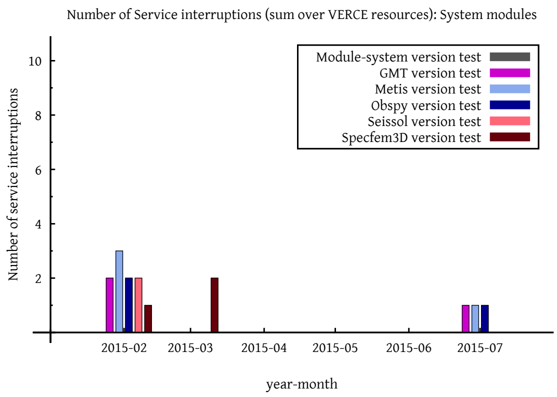

(a) Average Number of Service Interruptions per System Module across all Resources

(b) Number of Service Interruptions of each System Module across all Resources

Figure 6 – Number of Service Interruptions of System Modules

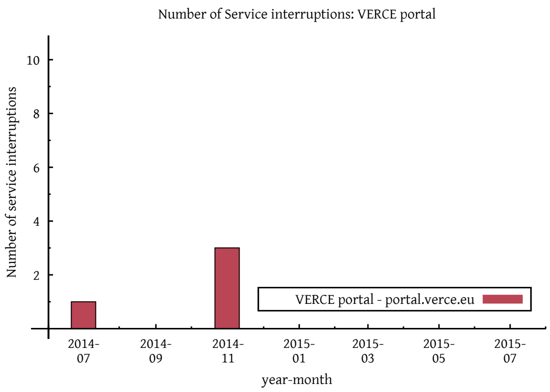

(a) Number of Service Interruptions of Science Gateway

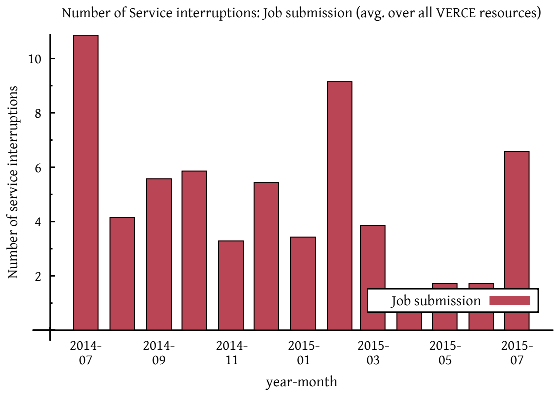

(b) Number of Service Interruptions of Job Submissions

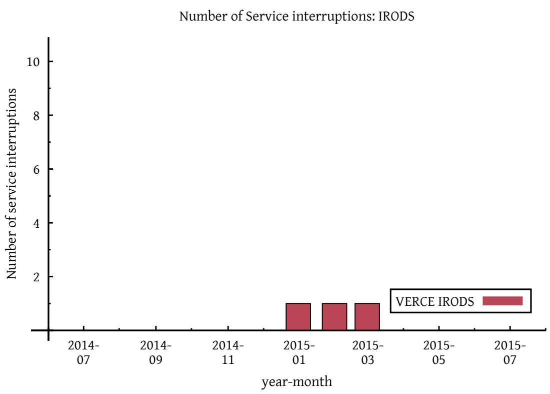

(c) Number of Service Interruptions of iRODS

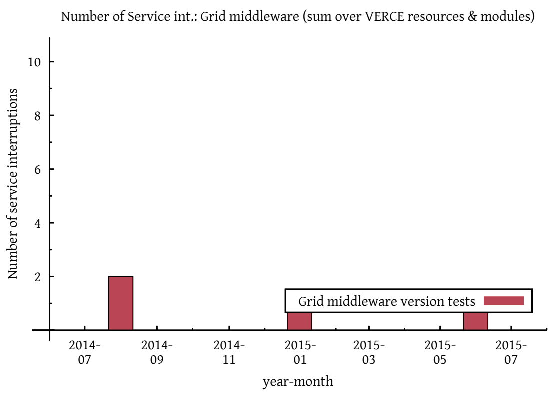

(d) Number of Service Interruptions of Grid Middleware

Figure 7 – Number of Service Interruptions of Core Services

2.1.4. Availability Monitoring

Figure 11 shows the availability monitoring information from July 2014 to July 2015. The decrease in the percentage of components monitored in November and December 2014 was due to a manpower issue. The support personnel of Inca left the partner institution and consequently, there was a gap in the support and maintenance of the Inca system. In March 2015, the introduction of a new resource at INGV

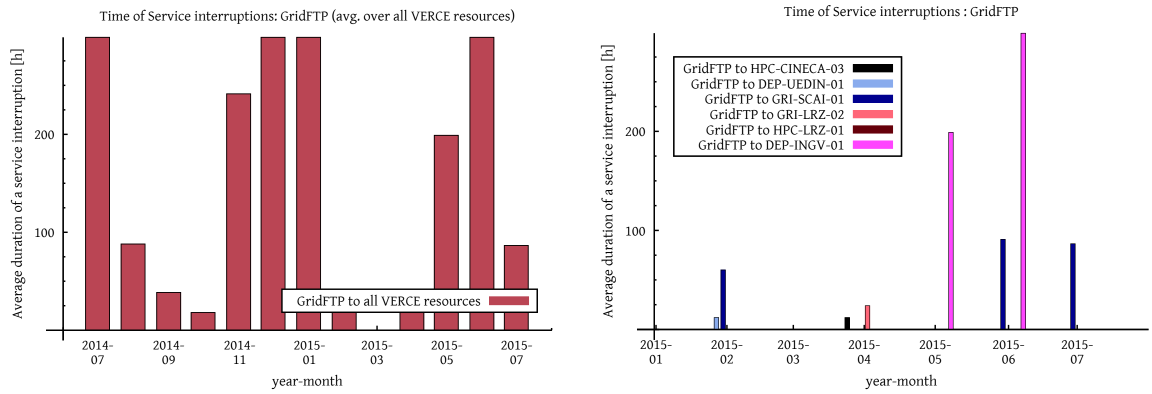

(a) Average Duration of Service Interruption of Data Transfer (GridFTP) Per Resource

(b) Duration of Service Interruption of Data Transfer (GridFTP) for each Resource

Figure 8 – Duration of Service Interruption of Data Transfer (GridFTP)

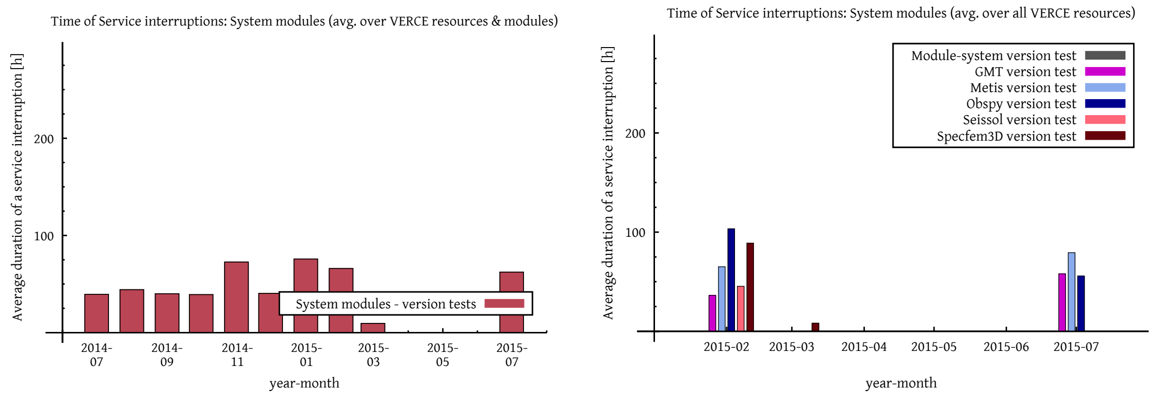

(a) Average Duration of Service Interruption per System Module across all Resources

(b) Average Duration of Service Interruptions for each System Module across all Resources

Figure 9 – Duration of Service Interruptions of System Modules

within Inca, triggered numerous bugs that affected the monitoring system. Thus, we can conclude that more than 80% of the deployed components are monitored.

2.1.5. Number of reported service failures

Table 3 shows the number of reported (and announced) service failures from January to July 2015. For scheduled maintenance, the failure results are not included in the previous section since they are not considered as downtime. However, reported failures, which include unscheduled maintenance, are included in the availability charts in Section 2.1. One unscheduled maintenance occurred on each of LRZ's resources this year. The unscheduled maintenances, a failure in the cooling system and a failure after maintenance, are reported and included in the data collected in Section 2.1. Similarly, there was an unscheduled maintenance due to a cooling issue at INGV for 2 days in June 2015.

Table 3 – No. of Reported Service Failures (January 2015 - July 2015)

	HPC-LRZ-01	GRI-LRZ-01	DEP-SCAI-01	GRI-SCAI-01	DEP-INGV-01
Scheduled Maintenance	4	2	3	5	0
Reported Failures	1	1	0	0	2

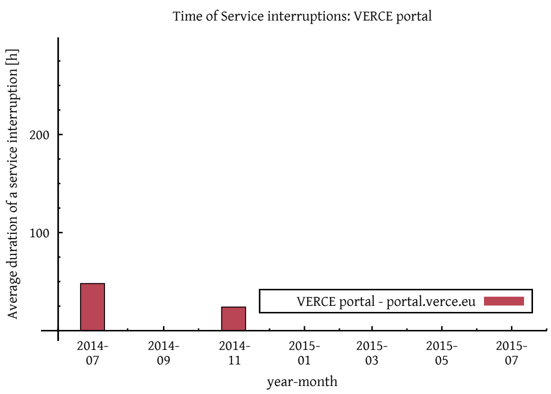

(a) Duration of Service Interruption of Science Gateway

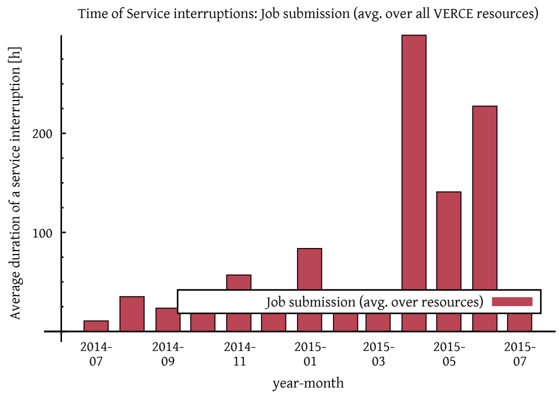

(b) Duration of Service Interruption of Job Submissions

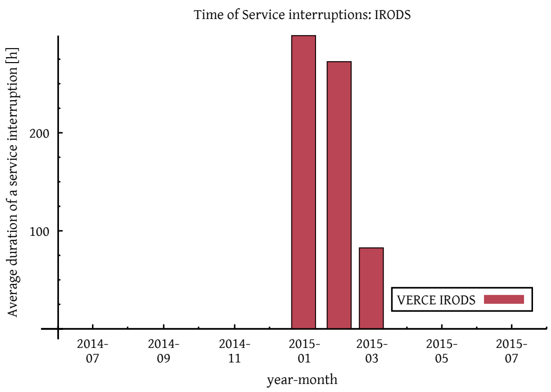

(c) Duration of Service Interruption of iRODS

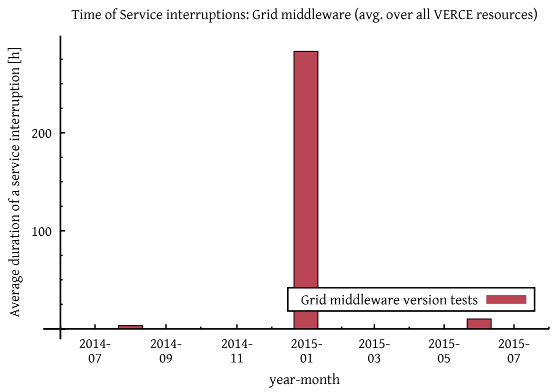

(d) Average Duration of Service Interruption of Grid Middleware

Figure 10 – Duration of Service Interruption of Core Services

2.2. Tools/Service/Software releases

Figure 12(a) shows the number of components that are evaluated and accepted with each platform release throughout the VERCE's project lifetime. In the last four releases, a 100% acceptance rate of evaluated components is achieved. Figure 12(b) illustrates 100% of completed evaluations and deployments within the scheduled timeframe per platform release over VERCE's entire project lifetime.

2.3. Security

Table 4 shares the security measurements in this reporting period. There was a major security incident where the Science Gateway was hacked. Quick action was taken by UEDIN to mitigate the risk and reduce the impact on users. A total of four standard security patches, e.g. heartbleed and SSL, were added to our operation services, e.g. the Science Gateway and Inca.

Table 4 – Security Measurements

Number of major security incidents	1
Number of security patches	4

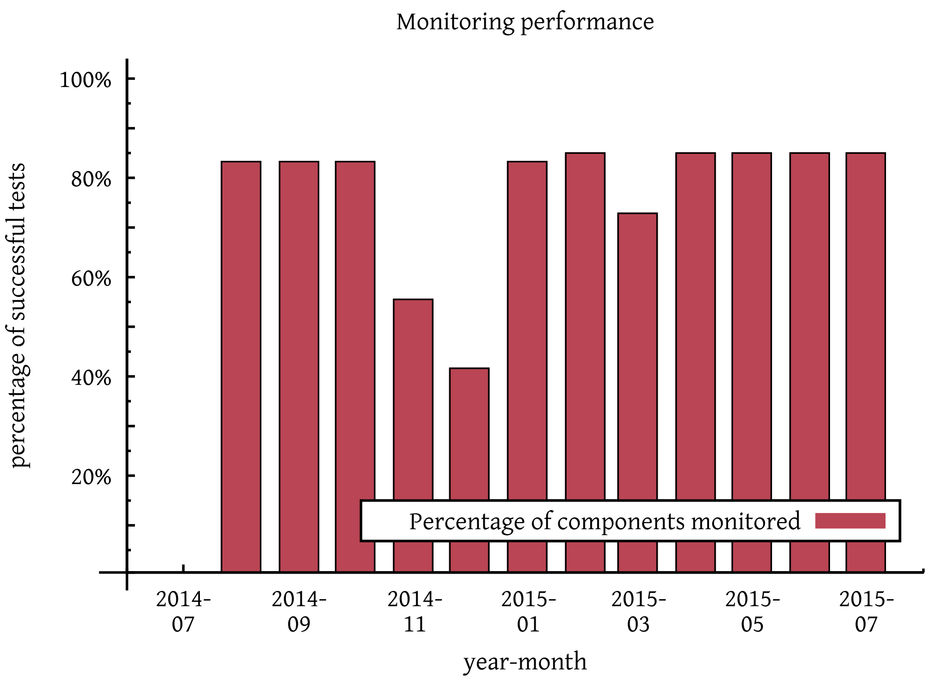

Figure 11 – Availability Monitoring

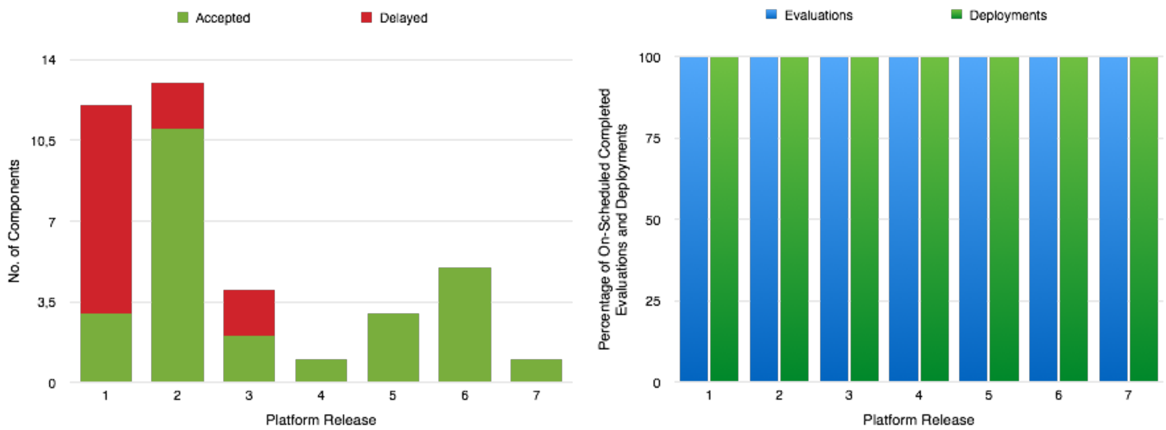

(a) Accepted and Delayed Components Per Platform Release (b) Completed Evaluations and Deployments within Scheduled Timeframe Per Platform Release

Figure 12 – Platform Releases over VERCE's Project Lifetime

2.4. Quality of Service

Table 5 shares the security measurements in this report period. Eight tickets from various beta users were received in this reporting period. Due to the tendencies of users to reuse existing tickets on new issues, each ticket typically contains more than one incident/issue. The average initial response time is 7.4 hours. The delay can be attributed to non-office hours and weekends. As such, the average response time is acceptable, i.e. within 24 hours. Typically, each incident requires an average of 0.5 hours before an acceptable solution is provided and accepted by the users. This is a result of the the professional support from experienced SA1 and SA3 operation members.

Table 5 – Quality of Service Measurements

Number of incidents	8
Average initial response time	7.4 hours
Incident resolution time	0.5 hours

3. Lesson Learned

The lessons learned in the previous release cycles were compiled and proposed as improvements to be carried out in every new release. Clearer recommendations introduced previously continued to provide a guideline to SA1 on the best matched or preferred resources. The constant communication with the work package members that had requested for component evaluation throughout the project lifetime.

The requirements of collecting data for the KPIs were clearly communicated to SA1 and led to the successful delivery of the required data. Consequently, less reliable services are identified and corresponding improvements are proposed and/or implemented. This enable a continuous improvement to sustaining a platform like VERCE.

4. Conclusion

The seventh release of the VERCE software components successfully completed within the scheduled time frame. The only requested component was approved in this release and will be a part of the final VERCE platform. There are a total of twenty-two unique approved components, including all previous releases, on the platform. Due to the integration and evaluation activities of SA2, the VERCE architecture remains stable and robust throughout the project lifetime and enables the easy integration of new and updated components.

The measurements from the KPIs indicated a good overall performance of the VERCE platform across the most crucial areas. Some services have room for improvements and are correspondingly communicated to the operators.

In conclusion, the procedures introduced by SA2 are effective in managing the release. Improvements that were made in the earlier releases have helped SA2 to carry out its evaluation smoothly and effectively. A big thanks to all SA2 members for their active participation and motivation.

4.1. Achievements

Table 6: Achievements of WP6-SA2

DOW Objective	Achievement	Comment
Define a standard set of procedures for integration activities and the appropriate software-engineering processes	An overlapping PDCA cycle was established in the first reporting period. On top of the cycle, a management release process, once every six month, was created. Consequently, it enabled coordination among the WPs and led to the seven successful and on-time VERCE Platform releases.	Due to the very well-defined management release process, new components were always integrated into the VERCE Platform in a non-intrusive manner.

Analyse and rank codes and tools in terms of ease to install and deploy	Codes and tools that are recommended by the JRAs and SAs for deployment were analysed, i.e. check phase, within the six monthly PDCA cycle. The management release process provides a template to rank each analysed codes and tools in terms of ease to install and deploy. This enables a highly organised and efficient way to perform the required check. The installation procedures of codes, e.g. ObsPy and Specfem3d, were also improved upon our feedback.	The ease to install and deploy a tool/code is frequently dependent on the targeted systems. Consequently, there can be differing views by the evaluators on the rank.
Check performance and scalability of application codes on VERCE platforms to identify the best-suited platforms for integration and deployment.	Application codes were benchmarked on HPC platforms to check their performance and scalability within the PDCA cycle. The management release process includes also a step to evaluate the performance. This enables a highly organised and efficient way to perform the required check.	Since Specfem3D is the main and only code identified to be integrated onto the VERCE Platform (as per recommended by the reviewers), all HPC resources became targeted resources. Grid and departmental resources that can provide reasonable performance were also selected.
Evaluate existing services, tool and middleware offered by participating infrastructures.	Existing services, tool and middleware, i.e. components, offered by participating infrastructures, EGI, PRACE and EU-DAT, were first evaluated. Similarly, departmental components used by partners were also evaluated. Useful components, e.g. Globus, GridFTP, iRODS, gUSE, WS-PGrade, were successfully selected and integrated into the VERCE Platform.	The prioritised selection of existing services, tools and middleware offered by participating infrastructures ease the integration. Correspondingly, it allows the VERCE platform to be easily integrated to other resources within such infrastructures that are outside the VERCE consortium.

Integration of essential higher-level development and optimisation tools, e.g. ADMIRE, for data exploration, data integration, workflow construction and optimisation.	The first tool to be evaluated and recommended for integration was GridFTP, supported by IGE (a EU project where VERCE has a MoU with). It offers the best optimised stage-in and stage-out performance when compared to other open source tools. Additionally, it is a tool that is widely adopted by participating infrastructures. WS-PGrade and gUSE as opposed to ADMIRE is also selected and integrated for infrastructure and workflow management. This decision is made in view of the strong support that could be offered by SCI-BUS, a EU project where VERCE has a MoU as per recommended by the EU reviewers. Additionally, a new tool, dispel4py, was developed by JRA2, which was build on top of ObsPy, a core component within the VERCE platform. dispel4py manages misfit analysis workflow description, data exploration and optimisation. Data services like EIDA and FDSN were also exploited.	New developments within VERCE were offered recommendations to improve the integration process with the VERCE Platform before, during and after each release.
Assist SA1 to deploy application codes, client services and tools on the test-bed and finally on the production infrastructure.	Full detail documentations on the deployment process and testing were provided by SA2 to SA1 at the end of the evaluation process. Clear instructions on which are the recommended infrastructures and resources were also provided for each accepted components. Consequently, the deployment process can commence without delay and with ease. This contributed to the 100% on scheduled deployment for all release cycles.	

Analyse the statistical data and user feedback collected by SA1 and SA3 and assess the quality of data services and tools and perform corrective measures (including re-evaluation and re-integration) and/or enhancement or expansion of services and tools if necessary.	A set of KPIs were defined to quantify the statistical data and user feedback collected by SA1 and SA3. Corrective measures were also recommended to SA1 and SA3 to improve the quality of deployed services and tools, and to the JRAs on new and updated services and tools. Consequently, no tools and services were delayed in the evaluation process due to the corrective measures taken.	
Support to NA2 for the evaluation of the VERCE platform based upon the selected use cases.	Support to NA2 to evaluate the VERCE platform based on the selected forward simulation use case led to an improved version of Specfem3d that not only benefits VERCE but also the seismological community. The installation and testing procedures of ObsPy were also improved that benefited the ObsPy community. Members of NA2 were also allowed to access PRACE's Tier 0 resources to fully experience the benefit of top HPC resources and the ease the Science Gateway offers, i.e. to set up and perform a forward simulation.	Effort is made among the infrastructure providers to sustain the platform by continuing to maintain the VERCE software stack for a year to support the transition to EPOS.

References

[1] ObsPy Wiki, *ObsPy: A Python Toolbox for seimolgy/seimological observatories*, March 2014, `https://github.com/obspy/obspy/wiki`.