



## **D-SA1.3.1 – VERCE platform: updated operation and monitoring report including Grid Appliance server deployment**

**31/03/2014**

---

*Project acronym:* VERCE  
*Project n°:* 283543  
*Funding Scheme:* Combination of CP & CSA  
*Call Identifier:* FP7-INFRASTRUCTURES-2011-2  
*WP:* WP5/SA1, Management and operation of the research platform  
*Filename:* D-SA1.3.1.pdf  
*Authors<sup>1</sup>:* M. Carpen, A. Gemünd, G. Moguilny, V. Muraleedharan, C. Ramos Garcia, I. Saverchenko, H. Schwichtenberg, D. Weissenbach.  
*Location:* <http://www.verce.eu/Repository/Deliverables/RP3>  
*Type of document:* Deliverable  
*Dissemination level:* Public  
*Status:* Draft  
*Due date of delivery:* 31/03/2014  
*Reviewer:* Anton Frank (LRZ)  
*Keywords:* Data platform, iRODS, MongoDB, OTRS, Inca, monitoring, Science Gateway, liferay, gUSE, GT5, cloud, integration of sites, user support, deployment, Unicore, Grid, HPC, X.509 certificates.

---

<i>History</i>	<i>Author</i>	<i>Date</i>	<i>Comments (and sections)</i>
1	G. Moguilny (IPGP)	14/03	Initial skeleton
2	V. Muraleedharan (IPGP)	20/03	End user support (section 5)
3	V. Muraleedharan (IPGP)	23/03	Data platform with iRODS (section 6)
4	C. Ramos Garcia (KNMI)	25/03	The Science Gateway (section 3)
5	A. Gemünd (SCAI)	28/03	Access to the platform (section 4)
6	A. Gemünd (SCAI)	28/03	Cloud evaluation (section 8)
7	A. Gemünd (SCAI)	28/03	Integration of new sites (section 2)
8	M. Carpen (CINECA)	31/03	Some Unicore details in section 1
9	I. Saverchenko (LRZ)	31/03	Inca monitoring (section 7)
10	D. Weissenbach (IPGP)	31/03	Updated lists of resources (Appendices)
11	G. Moguilny (IPGP)	01/04	Summary and conclusions
12	A. Gemünd, H. Schwichtenberg (SCAI)	14/04	Addressing reviewers comments
13	Th. Garth (ULIV)	16/04	English improvement
14	P. Danecek (INGV)	18/04	Corrections including iRODS clarifications

---

<sup>1</sup>Alphabetical order

## Copyright notice

COPYRIGHT © VERCE PROJECT, 2011-2015. SEE [www.verce.eu](http://www.verce.eu) FOR DETAILS ON VERCE.

VERCE, *Virtual Earthquake and seismology Research Community e-science environment in Europe*, is a project co-funded by the European Commission as an Integrated Infrastructure Initiative within the 7th Framework Programme. VERCE began in October 2011 and will run for 4 years.

This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, and USA.

The work must be attributed by attaching the following reference to the copied elements:

COPYRIGHT © VERCE PROJECT, 2011-2015. SEE [www.verce.eu](http://www.verce.eu) FOR DETAILS ON VERCE. Using this document in a way and/or for purposes not foreseen in the license requires the prior written permission of the copyright holders. The information contained in this document represents the views of the copyright holders as of the date such views are published.

## Contents

<b>Executive Summary</b>	<b>4</b>
<b>1 Resource updates</b>	<b>5</b>
<b>2 Integration of new sites</b>	<b>5</b>
<b>3 Science Gateway</b>	<b>6</b>
<b>4 Access</b>	<b>6</b>
<b>5 End user support</b>	<b>7</b>
<b>6 Data platform with iRODS</b>	<b>8</b>
<b>7 Monitoring with Inca</b>	<b>9</b>
<b>8 Cloud evaluation</b>	<b>10</b>
<b>Conclusion and perspectives</b>	<b>12</b>
<b>Appendices</b>	<b>13</b>
<b>A Computing resources available</b>	<b>13</b>
<b>B Software components</b>	<b>14</b>
<b>C Domains</b>	<b>15</b>

## List of Figures

1	Current iRODS federation status of sites. . . . .	8
2	Main components of VERCE data platform at UEDIN. . . . .	8
3	Default Inca index page. . . . .	10
4	VERCE Inca summary view. . . . .	10

## List of Tables

1	Hardware resources. . . . .	13
2	Data and Software resources. . . . .	14
3	Components (recommended by SA2) / resource . . . . .	14

## Executive Summary

The main objectives of the WP5/SA1 work package are to:

- Operate and manage in coordination with SA3 a Virtual Organisation.
- Provide tools and services for a unified access to the VERCE research platform combining Data, Grid and HPC resources.
- Provide and manage the distributed research platform.
- Provide and manage a set of application-tailored workbenches and enactment gateways, in coordination with SA2, to support specific use-case scenarios.
- Provide and manage a development testbed on which SA2 and JRA2 develop, integrate and evaluate the next releases of the VERCE platform.

The process of integrating a new site in the VERCE platform is now more clearly structured and documentation is available for site administrators/resource providers. Sites that want to provide an independent Science Gateway instance including the Forward Modelling application can now find the installation described in the Deployment wiki of the project as well.

VERCE's central Science Gateway instance is now accessible at <http://portal.verce.eu>. Preparatory steps like requesting a X.509 certificate and applying for membership in the VERCE.eu Virtual Organisation, have been more clearly documented. Users going through these steps can contact SA1 through new support channels: by mail through [support@verce.eu](mailto:support@verce.eu) or by using the new public ticketing system, based on OTRS, at <http://support.verce.eu>.

SA1 operates the VERCE data platform based on iRODS to give the users access to their data at all VERCE sites. Transport of the data between sites is now integrated in iRODS through a GridFTP storage interface module from the Verce partner CINECA. A testing instance is provided for the development of iRODS microservices that integrate with workflow metadata.

The Inca monitoring environment continuously evolves to reflect development of the VERCE infrastructure (<http://inca.verce.eu>). The integration of the VERCE HPC resources has been finalised, so all HPC and Grid resources are now monitored.

The Cloud evaluation is a work-in-progress. An application using the ObsPy Python Framework for Seismology as an elastic Cloud service has been accepted from the EGI Federated Cloud (FedCloud) TaskForce.

## 1 Resource updates

The updated list of computing resources available to VERCE can be found in Appendix A on pages 13 and 14. The updated list of components per site is described in Appendix B on page 14. Verce domains and subdomains used are listed on Appendix C page 15.

In order to simplify the integration between the VERCE Science Gateway and UNICORE, the most recent release has been deployed and tested on the CINECA development environment. The UNICORE 7 server contains a lot of new features compared to the previous version 6.4.2-p2 (including data-triggered processing, data staging support wildcards for UNICORE protocols and job restart). It also improves support for proxy certificates, which is required for authentication and job submission from the VERCE Science Gateway.

## 2 Integration of new sites

As described in previous deliverables, VERCE builds upon the European HPC infrastructure PRACE and the European Grid infrastructure EGI. VERCE users can run computations on resources provided in these infrastructures, as well as on 'local' resources provided separately.

In both cases, SA1 supports and coordinates with the site administrators (i.e. resource centres) to ease the integration. The first aim is to clarify access and use policies. Secondly, VERCE requires some software to be available on the resources of the institution. To be able to integrate the heterogeneous set of available resources, VERCE makes use of open source Grid middleware.

In its current incarnation, the VERCE software expects a site to provide the following features

1. Resources of the site organised as a cluster using a local resource management system (LRMS) like Torque, Open Grid Scheduler or SLURM accepting batch jobs from users.
2. Grid middleware acting as entry-point to the cluster, with the middleware ports opened to external networks, while all other systems can be firewalled. Currently supported middleware include Globus, Unicore 6 and gLite. The most tested solution is Globus GRAM in version 5.2, which is also used in the demonstrators. Globus GRAM is easy to set up and [well documented](#).
3. GridFTP access to storage that can be accessed both from inside the cluster and outside the cluster to be able to transfer input data and output results. As server component it is suggested to use the easy to use [globus-gridftp-server](#).
4. (optional) GSIssh access to a front-end node for manual job submission and testing / debugging.
5. IGTF compliant host certificates for services that provide outside interface (GSI services).
6. Sites should accept EUGridPMA certificates and X509 proxies (cf. section 4).
7. VERCE environment (codes and scientific libraries) deployed:  
Available through the initialisation of environment modules (cf. [D-SA1.2](#)). The parts of the stack that should be available on the site depends on the applications that the site wants to offer to the users. Additional information is available on the VERCE deployment wiki for installation, and on the VERCE end-user documentation wiki for use. For further details and support, sites can contact SA1 through [support@verce.eu](mailto:support@verce.eu).
8. Inca reporters to participate in operational monitoring (cf. section 7).

SA1 supports sites in getting this set of features up and running. If a site does not yet have a local resource management system, administrators can also be supported during the basic setup of the cluster. Installation details should be reported back to SA1 to finalize integration of the site.

The original idea of providing the experience of a drop-in appliance has proven infeasible until now, mainly because of the large difference between the institutions resource configurations. The Grid middleware needs to be closely integrated with existing configuration, e.g. it needs different packages depending on the adopted LRMS, and requires access to job submission and logs. Therefore such an appliance would be difficult to preconfigure. Nevertheless, a virtual machine for client software has been created, which can be used to explore the resources in the VERCE infrastructure by utilizing the Grid middleware clients.

The experience acquired during the integration of INGV's departmental cluster can be regarded as representative for the integration of new sites that are not participating in any infrastructure such as PRACE or EGI.

The VERCE software stack can also be used to offer services solely for local users. For example, an institution could host its own instance of the VERCE portal which accesses only local resources (cf. section 3).

### 3 Science Gateway

One of the priorities of VERCE during this reporting period has been to provide the VERCE Science Gateway for user evaluation at the next training event in June.

Two instances of the Science Gateway have been set-up, a [development and test instance](#) for project-internal use and a [production instance](#) for end-users.

The documentation of the entire installation procedure is available for project members in the Deployment area in the projects' Redmine<sup>2</sup> and available on request, for externals, by mail at [support@verce.eu](mailto:support@verce.eu).

### 4 Access

As described in [D-SA3.1.3](#), the VERCE Science Gateway has reached its first production state and can now be used to easily access the VERCE infrastructure. As described in previous deliverables (e.g. [D-SA1.3](#)), the common authentication scheme supported by all resource providers is based on X.509 certificates. Thus, users that would like to be able to submit computations through the VERCE Science Gateway should possess a valid e-Science certificate issued by the corresponding national certificate authority belonging to the EUGridPMA. As described in [D-SA1.3](#), the portlets of the SCI-BUS Gateway Framework (gUSE) are used to integrate this in the Science Gateway.

Documentation of the request process for new certificates has been improved and is supported by SA1 for end-users. A map of the providers can be found on the web presence of the EUGridPMA<sup>3</sup>. The registration process differs among the member states, but can be summarised as registering at the National Certification Authority and being validated by a local Registration Authority (RA). After users have completed the process, they own a valid certificate compliant with the policies of the EUGridPMA, that can be used to authenticate with various services in the European e-Science landscape.

In the EGI, authorization is based upon membership in Virtual Organisations (VO). Researchers in possession of a valid certificate can easily register through the VO web interface<sup>4</sup>. After a confirmation by VO representatives, they are authorized to use the Grid resources supporting VERCE.

<sup>2</sup>[http://www.verce-project.eu/projects/deploy1/wiki/Install\\_the\\_HPC\\_use-case\\_Forward\\_Modelling](http://www.verce-project.eu/projects/deploy1/wiki/Install_the_HPC_use-case_Forward_Modelling)

<sup>3</sup><http://www.eugridpma.org/members/worldmap>

<sup>4</sup><https://verce-voms.scai.fraunhofer.de:8443/voms/verce.eu/>

To be authorized to use the Supercomputers of the participating computing centers, users should follow the processes described at their web presence <sup>5, 6</sup>.

The documentation of this enrollment process has also been integrated in the VERCE Science Gateway to be as easily accessible as possible. For more detailed instructions and direct support, users can contact SA1 through [support@verce.eu](mailto:support@verce.eu) or by using the new public Ticketing system at <http://support.verce.eu/> (cf. section 5), which has also been embedded in the Science Gateway.

Apart from evolving this existing process, SA1 continues to follow current developments in the area of Federated Identity (or Federated AAI). While some promising approaches are appearing (cf. D-SA1.3), there is currently no solution available that is supported by all infrastructure providers.

## 5 End user support

The VERCE platform is now extended by end user support services. For VERCE platform developers, VERCE is already using the features of Redmine. A professional ticket system for end users is deployed for operation tasks.

SA1 has identified the following requirements for a ticket system.

### Essential:

- Log tickets from users
- Change status of ticket on first response
- Close/Solve tickets within a certain timeframe
- Direct tickets to queues for classification in categories
- Report detailed statistics for logged tickets (Number of incidents, Average initial response time, Incident resolution time)

### Optional:

- Add tickets via email
- Respond to tickets via email
- Forwarding of tickets to ticket systems of infrastructures or technology providers

Redmine, xGUS, Request Tracker and OTRS have been considered to satisfy the requirements. OTRS was finally chosen, because apart from fulfilling all of the collected requirements, project partners have previous experience with it, and it has strong community support.

The ticketing system is now publicly available through a project URL (<http://support.verce.eu>) and SA1 has defined a first general classification of tickets as shown below. If required, this classification could be revised based on feedback from users and members.

- Account: To request access to specific resources
- Software Request: To request for new software, new version of an existing software, or specific question about a component
- Security Issue
- Service Availability
- Others

The ticketing system is now directly available from the VERCE Science gateway. The testers are already using it for account creation and other access related issues.

<sup>5</sup><https://userdb.hpc.cineca.it/>

<sup>6</sup>[http://www.lrz.de/services/compute/supermuc/access\\_and\\_login/](http://www.lrz.de/services/compute/supermuc/access_and_login/)

## 6 Data platform with iRODS

The VERCE data platform connects the data resources available at the partners' sites. The backbone of the platform is an iRODS federation, which currently consists of resources at CINECA, INGV, SCAI, UEDIN, ISTERre and IPGP. This retains the full administrative rights for each site respecting their data policies, while providing comfortable access and easy sharing mechanisms through a common interface. iRODS provides a set of data management tools and clients for users to upload, download and view data and results in iRODS. The status of the federation is shown in Figure 1.




Figure 1 – Current iRODS federation status of sites.

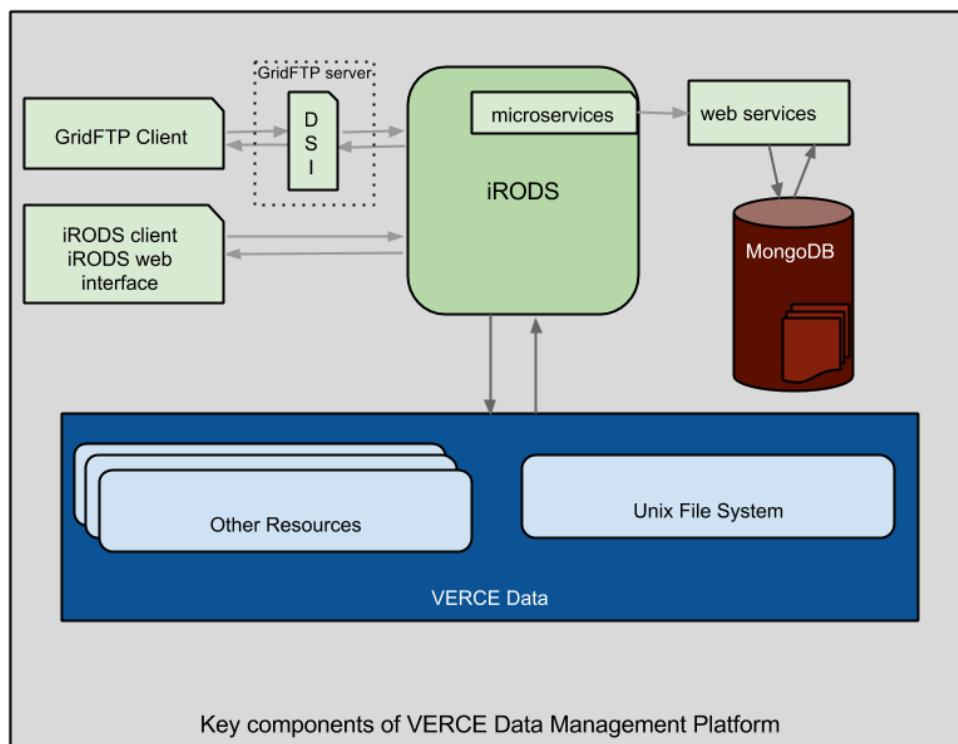


Figure 2 – Main components of VERCE data platform at UEDIN.

The data platform is already used to store data from forward modelling simulations. It will also be used to manage raw and processed data in the data-intensive cross-correlation use-case.

To test further extensions and newly developed catalog services, one iRODS installation is used as a test site. It uses storage resources at UEDIN and LRZ, which allows easy replication of high priority data, and is federated with the other installations at IPGP and SCAI. Data catalog, web services and microservices will be introduced at the other sites once they have been successfully tested on this development instance. To provide fast and easy access to data using metadata and provenance information, iRODS is coupled with a scalable MongoDB catalog in this case. Figure 2 shows how the different components come



together to form the data platform for VERCE.

Data has to be moved in and out of the platform to different resources. While iRODS provides its own parallel data transfer implementation, VERCE requires GridFTP for inter-operability with other services. With the help of the GridFTP-iRODS Data Storage Interface module (DSI) from CINECA, data can be easily transferred to and from the resources and platform using GridFTP without installing iRODS clients.

SA1 is in discussion with SA3 and JRA2 members to decide on the implementation of data replication and privacy policies.

The web based data management tool embedded in the VERCE Science Gateway currently requires password authentication. In the future, this has to be replaced by X.509 certificate authentication to be consistent with the other services.

The next major task is to classify the file types that have to be cataloged by the platform and separate them from any other user data. This task will be achieved side-by-side with the support for the data intensive use-case.

## 7 Monitoring with Inca

Inca is an application for monitoring software components and services in a distributed infrastructure. For a detailed introduction see deliverables [D-SA1.2](#) and [D-SA1.2.1](#).

In VERCE the latest version of Inca (version 2.6) is deployed. Inca server components are installed and running at LRZ and are accessible under <http://inca.verce.eu>.

Inca monitoring environment continuously evolves to reflect development of VERCE infrastructure. Over the reporting period we have finalised integration of all VERCE HPC resources by setting up Inca monitoring on HPC-CINECA-03. As such all HPC and Grid resources are being monitored. These include:

- HPC: HPC-CINECA-02, HPC-CINECA-03, HPC-LRZ-01,
- GRI: GRI-IPGP-03, GRI-LRZ-02, GRI-SCAI-01.

Inca was configured to be inline with internal VERCE conventions. This helps to maintain a consistent overview of deployed and operational services and software components as described in project documentation.

The Inca suite structure that describes the configuration of executed tests has been revised to better reflect the organisation of the VERCE infrastructure. At this moment, the following suites are configured:

- *Production environment*: describes availability of module system and modules for recommended and deployed software components.
- *Services*: monitors availability and version of services and respective client tools on VERCE resources, including GRAM, GridFTP and GSISSH.
- *Interactive access/GSISSH*: tests functionality of interactive access service based on GSISSH in VERCE infrastructure. Tests from a single external host are sufficient to determine if GSISSH services on a resource are fully functional.
- *Data access/GridFTP*: tests functionality of data access service based on GridFTP in VERCE infrastructure. Here cross-site tests are performed to ensure that every resource is accessible from every other resource in VERCE infrastructure. Cross-site access is required by end-user workflows.

- *Computing endpoints*: tests functionality of computing endpoints by submitting jobs to gLite, GRAM and UNICORE services deployed on VERCE resources.
- *Monitoring*: collects information describing state of Inca monitoring environment.

Inca consumer provides a web-based interface to access and visualise collected monitoring data. Functionality of the web interface was extended to enhance usability and navigation through available monitoring results. Originally a user could access the latest monitoring results by making a selection from a list of configured resources and suites as can be seen in Figure 3.

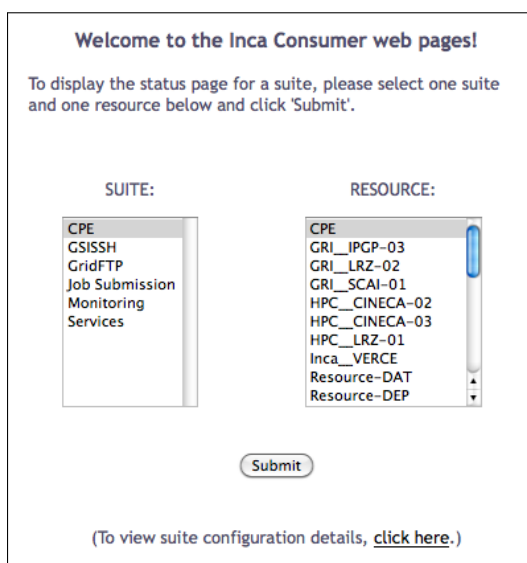


Figure 3 – Default Inca index page.

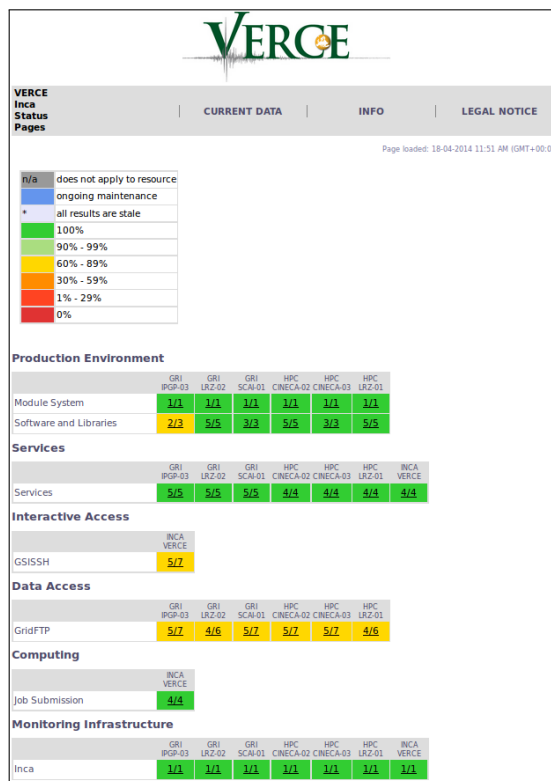


Figure 4 – VERCE Inca summary view.

This implies that all Inca users are familiar with Inca configuration, which should not be expected. To improve user experience a summary page providing an overview of VERCE infrastructure has been designed. This page is displayed to Inca users by default. Detailed views and reports are accessible through individual suites described on the page or via a "Current Data" menu on top of the page (Figure 4). Further Inca views and reports were improved to provide a clearer view on results of individual tests.

Documentation describing VERCE monitoring environment and Inca functionality was added to VERCE Wiki. This includes an extended description of resource access methods, reporter manager installation and usage guides and tools for validating system requirements. This documentation is essential for other partners to fully utilise the capabilities provided by Inca.

## 8 Cloud evaluation

The VERCE platform will be able to access European HPC, Grid, Data and Cloud infrastructures, coordinated by SA1. While the project initially emphasises HPC and Grid resources, SA1 has included an evaluation of Cloud last year. It evaluates possibilities to access resources of the EGI Federated Cloud (FedCloud) - an inter-operable integration of academic Cloud providers in Europe building on the existing EGI infrastructure, resources at SCAI based on Opennebula and the French Cloud initiative (see also D-SA1.3).

VERCE has enrolled a use-case in the EGI Federated Cloud, which has been accepted by the responsible TaskForce of EGI. To prepare more detailed information for the VERCE development team, SA1 and JRA1 have chosen an application from the VERCE validated software for a proof of concept. Questions to answer include how Cloud resources could be accessed by the VERCE Science Gateway (technical procedures to launch virtual machines, jobs, etc.), and from the perspective of seismologists, what advantages the Cloud can offer them compared to the existing infrastructures. Opportunities could for example arise through added flexibility for researchers. The possibility to administer the whole virtual machine, e.g. to install custom software and change configurations, may be helpful in some situations.

The VERCE use-cases require pre- and post-processing phases which can make use of resources outside of Supercomputers, such as Grid and Cloud resources. One of the main tools in VERCE for pre- and postprocessing is the ObsPy Python Framework for Seismology<sup>7</sup>. ObsPy was one of the first components that was evaluated and validated by SA2 and is a core component in VERCE. The proof-of-concept Cloud application allows users/researchers to discover seismic traces through a Web interface and to process them with Obspy on Cloud resources. For this purpose, data centers that provide a web service compliant with the web service standards of the International Federation of Seismograph Networks (FDSN)<sup>8</sup> are queried. Based on the user's data and parameter selection, the front end populates preconfigured/user-editable python scripts. With the help of a high-throughput yet light-weight messaging service, these python jobs are distributed to available worker VMs that carry out the processing.

The application is currently being developed and tested in the OpenNebula based Cloud environment at SCAI, but is not middleware-specific, and is planned to be tested on other available resources. The application is expected to be presented at the EGI Community Forum 2014.

Since the release of gUSE/WS-PGRADE v3.6.2 in February 2014, the framework, which is used for job submission from the Science Gateway, also supports the submission to certain Cloud resources. Time did not yet allow the evaluation of this feature of the framework, but it is planned for the following period. It is expected that prior work like virtual machine images (cf. [D-SA1.3](#)) can be re-used for this purpose. The implementation requirements are unclear as of yet, and therefore it is difficult to predict prospects and required effort.

---

<sup>7</sup><https://github.com/obspy/obspy/wiki>

<sup>8</sup><http://www.fdsn.org/webservices/FDSN-WS-Specifications-1.1.pdf>

## Conclusion and perspectives

The main achievements for this period are:

- best practice document for the integration of new sites;
- production portal to run the Science Gateway with Forward Modelling use-case;
- iRODS test environment in an OpenNebula virtual machine;
- Inca deployment now monitoring all HPC centres;
- improvement of the end-user support with a public web-based ticketing system;
- preparation of other use-cases for Cloud evaluation, with a web-GUI front-end providing query tools to find suitable seismological data.

Tasks for the upcoming period include:

- inclusion of additional iRODS storage nodes, with the catalogue process running on the data staging front end server,
- implementation of user driven replication policies for data which has to be reused, preserved or published,
- implementation of backup/restore mechanisms for production VMs, including the gateway services, users document store, provenance DB, etc,
- enhancement of the Inca monitoring,
- integration of new computational resources.
- definition of responsibilities for operational tasks and ticket service on different levels.
- collection of statistics and creation of reports based on tickets to help SA2 calculating Key Performance Indicators and Quality of Service.
- evaluating backup strategies for the VERCE platform.

# Appendices

## A Computing resources available

Table 1 – Hardware resources.

Designation	Resource Type	Short Name	Brief Description	Access Protocols	Accounting	Person in charge
<b>HPC-CINECA-02</b>	Computation	PLX GPU Linux Cluster	Model: IBM PLX (iDataPlex DX360M3).Architecture: Linux Infiniband Cluster, 274 IBM iDataPlex M3 Nodes, Processors: 2 six.cores Intel Westmere 2.40 GHz (548 processors, 3288 cores). GPU: 2 NVIDIA Tesla M2070/node for 264 nodes + 2 NVIDIA.Tesla M2070Q/node (for 10 nodes), RAM: 48 Gb/node.Internal Network: Infiniband with 4x QDR switches.Disk Space: 100 TB.Operating System: Red Hat RHEL 5.6, Peak Perf: 300 TFlop/s (142 TFlops sustained - Linpack benchmark)	ssh to login node	yes	Graziella
<b>HPC-CINECA-03</b>	Computation	BlueGene/Q	Architecture: 10 BGQ Frame with 2 MidPlanes each Front-end Nodes OS: Red-Hat EL 6.2 Compute Node Kernel: lightweight Linux-like kernel Processor Type: IBM PowerA2, 1.6 GHz Computing Nodes: 10.240 with 16 cores each Computing Cores: 163.840 RAM: 16GB / node; 1GB/core Internal Network: Network interface with 11 links ->5D Torus Disk Space: more than 2PB of scratch space Peak Performance: 2.1 PFlops	ssh to login node	no	Graziella
<b>HPC-LRZ-01</b>	Computation	SuperMIG/ SuperMUC	PRACE Tier 0 machine. SuperMUC (Summer 2012). Intel Xeon-Architecture, more than 110.000 cores. 3PFlop/s, 16cores/node with. 2GB/core memory and 40cores/node with 6.4 GB/core memory. 10 TB GPFS (aggregate BW 200 GB/s). linked via fully non-blocking. Infiniband network 1GB Ethernet to archive and backup	ssh to login nodes behind firewall, only accessible from IPs entered in firewall; then port 22 and globus ports range opened.	yes INCA, Dmon	Cerlane
<b>DEP-INGV-01</b>	Computation	selene	48 dual-proc AMD 6136 2.4GHz 8-core 64GB RAM (total 768 cores) 64 dual proc AMD 2378 2.4GHz 4-core 16GB RAM (total 512 cores) 36 TB scratch-grade storage	Account on request, then access via GSIsSh (ports 22 and 2222). GRAM via Globus gatekeeper (port 2119).	no	Daniele
<b>DEP-IPGP-06</b>	Computation	S-CAPAD Dell cluster	96 CPU nodes,16 data intensive nodes. 4 GPGPU nodes. 1 SMP node, parallel storage : GPFS on 576 effective TB, a high-performance network with InfiniBand QDR switches.		yes	Geneviève
<b>DEP-UEDIN-01</b>	Computation	EDIM1	Data intensive cluster. 120 nodes, each with 6.24TB of storage. CPU atom dual core, 4GB ram	SSH to login node behind firewall only from IPs within Edin. Informatics or EPCC	yes Ganglia	Paul
<b>DEP-ULIV-01</b>	Data	Linux Storage	CPU Intel Xeon dual core 2.8 GHz, 2GB RAM, 2.8TB storage		no	Andreas
<b>GRI-IPGP-03</b>	Computation	EGI Cluster	SiteBDII: sitebdii1.ipgp.fr - ldap://sitebdii1.ipgp.fr:2170/mds-vo-name=INSU01-PARISo=grid CREAM-CE: ce1.ipgp.fr - ce1.ipgp.fr:8443/cream-pbs-es DPM: dpm1.ipgp.fr SRM endpoint: srm://dpm1.ipgp.fr:8446/dpm/ipgp.fr/home/verce.eu/ GridFTP endpoint: gsiftp://dpm1.ipgp.fr:2811/dpm/ipgp.fr/home/verce.eu/ 136 CPU RAM: 1GB/CPU Storage: 8TB	gLite job/file operations, ssh to login on User Interface	yes	David
<b>GRI-LRZ-02</b>	Computation	Linux Cluster	> 2800 cores		yes (but only for internal purpose)	Cerlane
<b>GRI-SCAI-01</b>	Computation	EGI Cluster	32x Sun Fire X2200 M2 (Dual Opteron 2218, 2x2 cores), 8GB RAM, 250GB HDD 1 x Sun Fire X4600 M2 (Dual Opteron 8220, 4x2 cores), 78GB RAM, 2x140GB SAS HDD 8.5 TB NFS4 on ZFS shared filesystem (homes), 4.5 TB gLite DPM Storage Element (Grid) Connectivity: 2x Gbit Ethernet, Mellanox MT25204 10Gbps InfiniBand.	gLite, Unicore, Globus job/data access. GSISsh login to UI and dev-node (1, 2)	yes	André

Table 2 – Data and Software resources.

Designation	Resource Type	Short Name	Brief Description of Resource	Access Protocols	Installed Software	Accounting	Person in charge
GRI-EGI-01	Computation	EGI Virt. Org.	EGI Earth Science Research VO	gLite operations	gLite software stack	yes (EGI accounting)	André, David
GRI-EGI-02	Computation	EGI Virt. Org.	VERCE.EU VO Virtual European Broadband Seismograph Network, European Integrated Data Archive	gLite operations	gLite software stack	yes (EGI accounting)	André
DAT-KNMI-01	Data	ORFEUS Data Center		Web services, interactive sessions direct access from IPGP network	ArcLink		
DAT-IPGP-04	Data	COHER-D	IPGP Data Archive, NIED (Japan)			no	Geneviève

## B Software components

Table 3 – Components (recommended by SA2) / resource

TestBed										
	HPC-LRZ-01	HPC-CINECA-02	HPC-CINECA-03	DEP-UEDIN-01	DEP-ULIV-01	DEP-IPGP-06	DEP-INGV-01	GRI-SCAI-01	GRI-IPGP-03	GRI-LRZ-02
<b>Codes</b>										
SeisSol (1.0)	✓	✓	✓				◐	●		✓
SpecFEM3D (2.0)	✓	✓	✓			◐	●			✓
<b>Scientific libraries</b>										
Matplotlib (1.0.5)	✓	✓		●		●	●	✓	✓	✓
ObsPy (0.8.3)	✓	✓		●		✓	●	✓	✓	✓
Metis (4.0.3 or 5.0.2)	✓	✓	●			●		✓	✓	✓
GMT (4.5.9)	✓	✓			●	✓	●	✓	◐	✓
<b>Environment / Tools</b>										
iRODS (3.x)	★	✓	✓	✓			✓	✓	✓	★
iRODS-DSI (1.6)		✓	✓	✓			✓	●	✓	
GSI-SSH (GT 5.2.x)	✓	✓	✓				●	✓	●	✓
GridFTP (6.14)	✓	✓	✓	●			✓	✓	✓	✓
GRAM (5)							●	✓		✓
JSAGA (0.9.15)				●				✓	✓	
Unicore (6)	●	✓	✓							
gLite CREAM (1.14.0)								✓	✓	
gLite WMS (3.4)								✓		



Currently being installed



Installed



Available for VERCE users



★ backup replication of DEP-UEDIN-01 contents

An empty cell means not applicable or undecided as of yet

## C Domains

The **verce.eu** domain was bought at Gandi by IPGP before the start of the project. Today,

<a href="http://www.verce.eu">www.verce.eu</a>	(laboarix.ipgp.fr)	public website,
<a href="http://support.verce.eu">support.verce.eu</a>	(verce-support.ipgp.fr)	end user support,
<a href="http://portal.verce.eu">portal.verce.eu</a>	(dir-vm3.epcc.ed.ac.uk)	portal,
<a href="http://test.verce.eu">test.verce.eu</a>	(dir-vm5.epcc.ed.ac.uk)	test portal,
<a href="http://inca.verce.eu">inca.verce.eu</a>	(inca-verce.drg.lrz.de)	monitoring,

are used for the project.

The project also holds another domain, **verce-project.eu**, for the project management

<a href="http://verce-project.eu">verce-project.eu</a>	(sisprojs.ipgp.fr)	Redmine.
--	--------------------	----------