D-SA1.4.1 – VERCE platform: updated operation and monitoring report including application deployment evaluation

20/11/2015

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-SA1.4.1.pdf
Authors[1]*:* S. Claus, A. Gemünd, S. Hachinger, J. Matser, G. Moguilny, H. Schwichtenberg, D. Weissenbach.
Location: http://www.verce.eu/Repository/Deliverables/RP4
Type of document: Deliverable
Dissemination level: Public
Status: Draft
Due date of delivery: 31/03/2015 – revised on 20/11/2015
Reviewer: A. Michelini (INGV, WP2/NA2)
Keywords: iRODS, INCA, GlusterFS, Unicore 7, GRAM, GridFTP, gUSE, X.509, GSI, VOMS, DCI bridge, KVM, OpenNebula, WS-PGRADE, Liferay, Tomcat, QCOW2, MongoDB, FDSN, Cloud, Grid, HPC

History	*Author*		*Date*	*Comments (and sections)*
1	G. Moguilny	(IPGP)	24/03	Initial skeleton
2	S. Claus	(SCAI)	30/03	iRODS update (section 2)
3	S. Hachinger	(LRZ)	30/03	INCA update (section 3)
4	A. Gemünd	(SCAI)	01/04	Integration of new resources (section 1)
5	S. Claus	(SCAI)	01/04	Training preparation (section 4)
6	A. Gemünd	(SCAI)	02/04	Services operated by VERCE (section 5)
7	H. Schwichtenberg	(SCAI)	07/04	Edited summary and sections 1 to 5
8	J. Matser	(KNMI)	08/04	GCMT FDSN service
9	D. Weissenbach	(IPGP)	10/04	Resource tables (Appendices)
10	G. Moguilny	(IPGP)	13/04	Conclusion
11	A. Gemuend	(SCAI)	19/11	Revision of section 5
12	S. Hachinger	(LRZ)	20/11	Proofreading and further refinement of section 5
13	A. Gemünd	(SCAI)	20/11	Refinement of executive summary
14	S. Hachinger A. Gemünd	(LRZ) (SCAI)	20/11	Revision of section 1

[1] Alphabetical order

Contents

List of Figures

List of Tables

Executive Summary

SA1 is responsible for leading the operations and deployment activities of the VERCE platform. The activities of SA1 are aimed at providing a smooth operation of the infrastructure and services, to facilitate the usage and to provide support in case of issues.

The main objectives of the work package are to:

- Operate and manage a Virtual Organisation in coordination with SA3.
- Provide and manage the VERCE research platform, combining Data, Grid and HPC resources.
- Provide tools and services for unified access to the 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.

In this document we provide an updated report to D-SA1.4, where we introduced the re-installation of the VERCE services as virtual machines in Cloud environments. Since then, SA1 has improved its services based on feedback from daily operations and in particular on feedback from the training sessions VERCE has organised. Next to typical service activities maintaining installations and services, SA1 focused on supporting the ongoing developments of SA3 and JRA2 as described in the latest VERCE Roadmap. SA1 participated actively in the task force to implement the Misfit calculation.

In Section 1, we present changes regarding the resources integrated with the VERCE platform (section 1). The changes include a system that was replaced by one of the resource centers and one newly integrated local resource. The VERCE software stack and monitoring tools have been deployed, and SA1 worked closely together with SA3 and the respective sites on issues concerning the accessibility of the compute resources via the Science Gateway.

Operational experience, in particular gained during the training events, had shown that the stability of the iRODS data platform needed to be increased. Primary changes from a catalogue of measures are described in Section 2.

In the following Sections, we give a short overview about the maintenance activities around the monitoring system INCA (Section 3) and the preparation steps required for the training in March 2015 (Section 4).

In view of the end of the funding period in 2015, Section 5 presents an overview of the current VERCE services and their respective resource requirements. Thanks to the consistent usage of virtualisation technologies, it's possible to run the entire suite of services in any public or private Cloud environment.

1 Integration of new resources

All VERCE partners provide compute resources to the VERCE community. These resources are normally accessible under the policies of the respective site. In order to ensure accessibility via grid tools and specifically via the VERCE Science Gateway, additional middleware and mechanisms for handling X.509 EUGridPMA certificates [2] have to be installed. To a large part, these tasks involve the IT administrators in the various institutes. Usually, technical and security issues arise in this context, making the integration a cumbersome process. SA1 tries to mitigate problems by providing direct technical support to the local IT groups. Thus, we shall be able to connect most VERCE compute sites to the Science Gateway in the next months.

Two new compute clusters have been made available for VERCE: **DEP-SCAI-01** is a new institutional resource, for which access to project users can be provided upon request. It is an Intel Xeon 2650 based cluster (4096 cores) with Mellanox FDR Infiniband network and the Fraunhofer parallel file system [3]. The new GALILEO system at CINECA, denoted as **HPC-CINECA-04**, replaces the PLX HPC system ("HPC-CINECA-02"). Such replacements typically occur every 3 to 5 years at supercomputing sites. GALILEO is a new PRACE Tier-1 resource, and is among the fastest supercomputers available to Italian industrial and public researchers. It is equipped with up-to-date Intel accelerators (Intel Phi 7120p [4]). This strong system, which will surely be of great use for VERCE scientists, will be made available via the VERCE Science Gateway as soon as possible.

1.1 Science Gateway integration

The VERCE Science Gateway is based on the WS-PGRADE/gUse framework [5]. The abstraction layer to access compute resources in this framework is the DCI Bridge service[6]. Many site and infrastructure specifics have to be mapped to common denominators in the configuration of this service. The completion of this configuration for some resources remains future work.

For **HPC-CINECA-04**, CINECA provides external access based on the UNICORE 7 middleware. The DCI Bridge supports basic features of an older UNICORE version only, and its adaption is more difficult than expected. This work is currently being started together with the validation of all VERCE software components on the system, such that production tests can be conducted in the next reporting period.

DEP-SCAI-01 provides Globus GRAM and GridFTP site services, which are supported by gUSE. Thus, the integration is straightforward. Due to its suitability for MPI workloads (high-speed Infiniband network, fast parallel file system), the selection of software deployed on this resource is similar to HPC sites. The system is suitable to execute the VERCE forward modelling workflow. The site has already passed the respective VERCE validation tests.

The integration of **DEP-INGV-01**, another VERCE local resource, has been delayed due to site specific network privacy issues. In particular, the necessity to provide externally available services, that at the same time have access to the internal network of the cluster and its file systems, proved difficult to realise in compliance with the network policies of INGV. The INGV VERCE partners were able to arrange a VERCE exception with their authorities and SA1 site validation tests are currently carried out. INGV has all VERCE components in place (see Appendix B) in order to support the forward modelling workflow.

For the Data Intensive Research Cluster in VERCE, **DEP-UEDIN-01**, a virtual server providing access through the Globus middleware (GRAM and GridFTP) has been set-up from scratch by SA1 and was

[2] https://www.eugridpma.org/
[3] http://www.beegfs.com/content/
[4] http://ark.intel.com/de/products/75799/Intel-Xeon-Phi-Coprocessor-7120P-16GB-1_238-GHz-61-core
[5] http://guse.hu/about/home
[6] http://guse.hu/about/architecture/dci-bridge

further on integrated with the Science Gateway for the submission of workflows. The selection of components to be installed on **DEP-UEDIN-01** has been guided by the nature of the system, i.e. focused on local analysis of big data instead of pure processing power. Software which will allow for the execution of the Misfit workflow has been deployed on this resource, such as `dispel4py` (see JRA2) and `ObsPy` (see SA2 VERCE software components in Appendix B).

The integration of resources belonging to the European Grid Infrastructure (EGI) via the VERCE Virtual Organization will be attempted in in the next reporting period; preparations have already started. Another job-submission middleware (gLite [7]) has to be integrated with the gateway, requiring some adaptions by SA1 and SA3. The access to EGI clusters, in addition to PRACE HPC resources, makes more computing power available via the VERCE Science Gateway's unified interface.

2 Updated Data platform with iRODS

The VERCE data platform builds the foundation for the (meta)-data management of distributed VERCE data. It is maintained and operated by SA1 as the underlying data infrastructure to support the workflow tool and provisioning of application data (see D-SA1.3, D-SA1.3.1 and D-SA1.4). To maintain a failure-free operation, the storage infrastructure had to be restructured. Additionally, a revision of the data replication policies (see subsection 2.2) has been the second main activity within the reporting period of the deliverable D-SA1.4.1. Both activities serve the purpose of guaranteeing a 24/7 operation of the complete VERCE platform during training and daily operation.

2.1 Storage restructuring

At the UEDIN site a Cloud infrastructure based on a cluster with several local disks is provided to VERCE. It gives VERCE an opportunity to demonstrate how such a Cloud resource can be provided in an efficient and operational manner in order to support a distributed file system based on iRODS. The first approach by SA1 was a solution based on the Hadoop Distributed File System (HDFS) to merge the multitude of cluster nodes and local disks and operate them as a single, high-performance storage node for iRODS (see D-SA1.4).

However, due to integration issues of HDFS and the GridFTP iRODS DSI module (see D-SA1.4), which is an elementary requirement to transfer files from the VERCE workflows to the iRODS environment, this storage had to be removed again. As a substitute, a GlusterFS storage volume [8] was used. Currently, this storage consists of 4 separate nodes with a "replica value" of 2. Thus, all files on the GlusterFS environment are persisted on two separate storage nodes. By choosing GlusterFS over HDFS, GridFTP is now supported and a higher data reliability can be reached.

2.2 Replication policies

Besides this lower-level and iRODS-independent data duplication, file redundancy on the iRODS level was further investigated as well. During this reporting period, our work concentrated on replicating iRODS managed data across separate VERCE sites. Tests with a single iRODS zone spanning multiple VERCE sites have already been disregarded in the last deliverable due to significant performance degradations (mainly due to a across-site database queries).

During the last months, a better performing approach has been successfully tested. Instead of using a single site-spanning zone, so called "resource groups" were used. Such groups might consist of local as well as remote iRODS resources. By leveraging iRODS' micro services, more specifically triggering a

[7] http://grid-deployment.web.cern.ch/grid-deployment/glite-web/
[8] http://www.gluster.org/documentation/About_Gluster/

call of msiSysReplDataObject within the acPostProcForPut rule (which is used for post-processing a put operation, i.e. additions to the iRODS catalogue), we are able to transparently replicate local files to remote resources within the same resource group. Of course, in scenarios with data transfer over low bandwidth public networks, this approach inhibits similar drawbacks as single, site-spanning zones. However, resource groups can explicitly be specified during an iPut operation (which is the iRODS command for storing a file) and thus a more fine-grained control over file replication is given. As well, only catalogues of local sites are queried (in contrast to the site-spanning iRODS zone where the central database is queried for each file operation).

3 Updated Monitoring with Inca

The INCA monitoring platform, as described in D-SA1.3 and D-SA1.3.1, has received minor updates in this reporting period. In addition, VERCE resources not yet monitored before have been integrated. Below, we describe these efforts in some detail.

While the INCA platform is generally stable and reliable, communication problems from/to the INCA system have appeared due to different SSL installations on the various resources. The issues could be fixed, mostly by regenerating certificates / keys with a longer key length on all sites. An update of INCA hardening it against recent SSL vulnerabilities is being built in collaboration with upstream (INCA developers) and will be deployed soon; for the meantime the INCA server has received a firewall configuration update.

INCA monitoring has been extended to the INGV institute cluster (DEP-INGV-01) by installing the necessary components on the site. The INCA monitoring, with small exceptions, has thus been extended to cover all VERCE resources. It will also cover the upcoming new Tier-1 system of Cineca, GALILEO (HPC-CINECA-04).

The INCA system database, storing the monitored state of the VERCE platform, now contains explicit availability-of-service information. The total number of executed and failed/passed tests of each sort is stored on a weekly basis, gathering the necessary information form the standard tables of the INCA database. The availability-of-service table shall allow SA2 to calculate Key Performance Indicators (KPI's) as defined in D-SA2.3.

We further note that the INCA system in the current reporting period helped to spot a number of temporary problems with VERCE systems and/or their network connection, which – after the respective INCA alert and notification of the administrators – could be fixed soon after they appeared. Thus it ensures an optimum availability of VERCE Grid services.

4 Training preparation and end-user support

From March 9th to 11th, a VERCE user training was held at LRZ. Prior to this 3-day workshop, some preparatory steps had to be taken.

- IGTF compliant X.509 training certificates with a validity of 7 days were used for all training attendees.
- VERCE portal and corresponding iRODS user accounts were created.
- The VERCE portal training accounts were added to the VERCE VOMS.
- A mapping from the iRODS accounts to the certificate subject strings was applied to enable GSI authentication for all iRODS related file operations.
- iRODS folder layout as well as access rights required for the different training scenarios/workflows were configured on the iRODS environment at UEDIN. Subsequent tests showed a fully functional setup.

No failures or any other events of significance occurred throughout the duration of the workshop.

5 Services operated by VERCE

In order to ensure sustainability of the VERCE platform, core services of the VERCE platform will have to be maintained also beyond the end of the funding period.

In this chapter, we collect a list of the respective resources crucial for providing "VERCE core services" and of their characteristics. Thus, we clarify the efforts which are required to install/maintain the respective services. This shall serve as a planning aid, and also as an overview over the achievements of VERCE in establishing services up to the end of this reporting period.

5.1 Computational resources vs. resources for VERCE core services

The services VERCE provides to the user community split up into computational services on the one hand, and VERCE core services on the other hand.

Computational resources are required to carry out calculations and simulations, such as used in the forward modelling workflow. VERCE does not build compute clusters for this purpose, but relies on the availability of resources provided within European supercomputing/grid-computing incentives such as PRACE and EGI. In the course of the project, VERCE has achieved the availability of specialized tools and software with validated quality and compatability (c.f. SA2) on the systems that were selected for inclusion. Computing centres running the respective compute clusters have provided support much beyond their normal service to guarantee a validated consistency of the software stack (cf. SA2) as well as its availability in a carefully optimised form on each resource.

The VERCE platform is however not a mere aggregation of PRACE or EGI compute clusters, but provides core services integrating the different resources. To this end, interoperating servers have been set up within the infrastructures of the different project participants. These servers provide the necessary functionality for data management, execution of computational jobs, etc., running specialised software which is maintained and monitored by SA1. Furthermore, user community support, including technical assistance for the utilization of the platform with a focus on the scientific use cases, is provided via a respective support portal.

Computational resources will continue to be run by the respective resource providers, and the maintenance of the VERCE software stack installed on them should be a relatively straightforward task (which requires some manpower). The continued availability of the various compute clusters to VERCE users has to be clarified; results will be reported in the management report of work package NA1.

The maintenance of VERCE core services, on the other hand, will require a coordinated effort to ensure a continued high availability of the respective servers and to maintain various specialised software packages installed on them. Below, we first elaborate on the VERCE core services and then list the work needed to maintain them.

5.2 VERCE core services

The following core services are directly operated and supported by VERCE, and described in some detail further below:

- the Science Gateway and associated backend services
- a central instance of iRODS with corresponding metadata catalogue

- the provenance system developed within the VERCE project
- the User support portal with ticketing service
- the INCA monitoring system
- the GCMT FDSN web service

Science Gateway
The VERCE Science Gateway provides a uniform access point in the form of a web portal to the VERCE Platform. It is based on the gUSE/WS-PGRADE Science Gateway framework [9], allowing users to access third parties' data and compute resources (with the respective security mechanisms), without directly having to deal with complicated technological details.
Technically, the science gateway consists of two servers, hosting the frontend and the backend respectively. The frontend on the basis of WS-PGRADE is composed of Java Portlets implemented on the Liferay framework [10], which is running as a web application in an Apache Tomcat server. The backend consists solely of web services invoked via client API calls by the frontend. Both front- and backend are hosted in the framework of an Apache Tomcat server.

iRODS data management system
The large amount of data processed in the scientific workflows have motivated the VERCE consortium to set up and manage iRODS as an efficient system for distributed data management. Additionally, the Science Gateway offers interactive access to the iRODS "file system" through a graphical user interface, and the provenance system can automatically register files in certain folders.
The central component of the iRODS system is a server instance at UEDIN, which contains the full iRODS metadata catalogue. In addition to the central server, VERCE maintains an iRODS federation. As already mentioned in earlier deliverables, several VERCE resource centres participate in this federation with dedicated or shared machines. Resources and collections are shared between the federated iRODS servers by granting mutual access to users at other sites. To allow for continued operation of the platform, the federation does not necessarily have to be maintained: only one central iRODS server with the corresponding metadata catalogue and VERCE specific iRODS rules to connect the provenance service is strictly necessary.

Provenance system
The provenance system is responsible for collecting metadata about the workflows and data users instrument through the Science Gateway. It records selected parameters, pre-conditions, identifiers and important properties of data inputs, run-time information about the workflow and properties of all data produced during the run. In the provenance system developed within VERCE, provenance is represented using the W3C-PROV recommendation, providing flexibility and efficiency through the use of the JSON format. The system requires a scalable distributed database in which records are accumulated. For this to be flexible and extensible, with fast multi-faceted search, the NoSQL system MongoDB is used. The system is implemented using the Python twisted framework and runs as a stand-alone web service. Through the iRODS rule engine, all data uploaded to specific folders of the data management system is automatically indexed by the Provenance system, which is the reason why it is co-hosted on the same server.

User support portal
The VERCE user support portal provides a convenient web-based user interface for end-users to request help or report issues. It is based on the Open Ticket Request System (OTRS) [11], which offers advanced features such as ticket routing, queues, and extensive report generation facilities for the VERCE support

[9] guse.hu/about/architecture/ws-pgrade
[10] `http://www.liferay.com/products/liferay-portal/overview`
[11] `https://otrs.github.io/`

teams. It runs as a perl based web application on an Apache HTTPD server and uses a MySQL database to store data.

INCA monitoring system
The INCA system is used to monitor the state of resources and the availability and validity of the VERCE user environment on these resources. It implements a client-server model, where clients called reporter managers are deployed on the resources, and test components from a user-level perspective. They send the collected monitoring data to a central INCA server for processing, archival and presentation. Thus, INCA detects infrastructure problems by executing periodic, automated, user-level testing of software and services. The central INCA server consists of several daemons and a web application hosted in a Jetty servlet container.

GCMT FDSN web service
The GCMT FDSN web service provides the possibility to query the data of the Global Centroid-Moment-Tensor (CMT) catalogue maintained by the Global CMT project [12]. It is implemented as a RESTful web service compliant with the standards of the Federation of Digital Seismograph Networks (FDSN [13]). The implementation is based on an initial FDSN server implementation by Lion Krischer [14].

5.3 Tasks and resources

To continue operations of the aforementioned services, the hosting itself and the maintenance of the respective software need to be secured. Table 1 provides some technical details about the services and their current hosting arrangements. The listed centres have agreed to continue their provision of service after the end of the project. More detailed agreements will be reported in the deliverables of work package NA1.

Service	Type of hosting	Provider	Current capacity
Science Gateway frontend	IaaS VM	SCAI	4 CPU cores Intel Xeon E5-2650 v2 8GB RAM, 40GB HDD
Science Gateway backend services	IaaS VM	SCAI	8 CPU cores Intel Xeon E5-2650 v2 8GB RAM, 160GB HDD
iRODS data management & co-hosted provenance system	IaaS VM	UEDIN	2 CPU cores Intel Xeon E5-2470 2GB RAM, 250GB HDD, 3TB distributed storage (GlusterFS)
User support portal	IaaS VM	IPGP	
INCA monitoring system	IaaS VM	LRZ	2 CPU cores Intel Xeon E5-2660, 4GB RAM, 10GB HDD
GCMT FDSN web service	PaaS service	Red Hat OpenShift	'small gear' (1 core and 512 MB RAM) and 1 GB of storage.

Table 1 – Details about the current hosting arrangements of the services operated by VERCE

As the prevalent current hosting type (cf. Table 1) is to use Virtual Machines (VM) following the Infrastructure-as-a-service (IaaS) paradigm, practically no hardware management, but significant basic system-level management is carried out by VERCE. This includes regular administrative tasks such as

[12] http://www.globalcmt.org
[13] http://www.fdsn.org/webservices/
[14] https://github.com/krischer/FDSN_Event_Server

applying system upgrades and security patches to all systems to ensure basic security and stability. In the long term, this task could be handed over to local IT departments, constituting a shift towards managed service provision.

In addition to the management of system hardware/software, the following tasks are carried out regularly:

- Activation of entitled users on the Science Gateway and inclusion in the VERCE.eu VO
- Survey of monitoring system and reaction to reported events / issues
- User support through the ticketing system
- Prolongation of host certificates required for secure communication and authentication
- Update of Certificate Authority certificates and revocation lists
- Regular upgrades of Grid middleware
- Addressing changes in resources or their environments in the gUse backend and workflow properties
- Supervision of iRODS data management (sharing, permissions, space reclamation)
- Maintenance of user accounts in iRODS synchronized with portal user accounts and certificates
- Adaption of iRODS rules and micro services (e.g. replication policies, provenance service)
- Coordination and support of software deployment on compute sites

As the funding period is ending, the continued fulfilment of these tasks by suitable personnel shall be ensured.

Conclusion and perspectives

Besides the usual service activities, the following has been achieved during the reporting period:

- More VERCE compute resources are now available via the Science Gateway. This extends the possibilities for the users to use the Gateway even in times when no supercomputing resources are available and allows fall back scenarios during trainings.
- The iRODS based data infrastructure is more fail-safe now and thus serves as a stable backend for the Science Gateway.
- No failures occurred during the training with multiple users working simultaneously.
- Certificates are now more common in the VERCE community. For the training a solution with time-limited certificates was applied.
- All VERCE services are virtualised and may be hosted in commercial or public Cloud environments.
- More sites are now being monitored by VERCE with the INCA monitoring tools.
- A new service "GCMT FDSN" (see D-SA3.4.1) is now in operation.

In the next reporting period, the work of SA1 will continue with daily operation of the services. Daily support has to be provided to end users and developers of the VERCE work packages. Updates of the services, esp. in case of INCA, and hardening against SSL vulnerabilities have to be done. The addition of USGS data to the FDSN web service is another open task.

Besides the Globus middleware, Cineca uses UNICORE to access the compute resources. With SA1 support, UNICORE was setup again and is now supported by the Science Gateway. This integration has to be tested for production scenarios against the Science Gateway. SA1 also has to proceed with supporting the implementation of the Misfit use case.

Appendices

A Computing resources available

Table 2 – Hardware resources.

Designation	Resource Type	Short Name	Brief Description	Access Protocols	Accounting	Person in charge
HPC-CINECA-03	Computation	FERMI	IBM-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 PFlop/s	ssh to login node	yes	Piero
HPC-CINECA-04	Computation	GALILEO	IBM NeXtScale Nodes: 512 Processors: 2x8cores Intel Haswell 2.40 Ghz per node Cores: 16 cores/node, 8256 cores in total Accelerators: 2 Intel Phi 7120p per node on 384 nodes (768 in total) RAM: 128 GB/node, 8 GB/core Internal Network: Infiniband with 4x QDR switches Disk Space: about 2PB of scratch space	ssh to login node	yes	Piero
HPC-LRZ-01	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
DEP-INGV-01	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
DEP-IPGP-06	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
DEP-UEDIN-01	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	
DEP-SCAI-01	Computation		Intel Xeon 2650 based cluster (4096 cores) with Mellanox FDR Infiniband network and the Fraunhofer parallel file system			André
DEP-ULIV-01	Data	Linux Storage	CPU Intel Xeon dual core 2.8 GHz, 2GB RAM, 2.8TB storage		no	Andreas
GRI-IPGP-03	Computation	EGI Cluster	SiteBDII: sitebdii1.ipgp.fr - ldap://sitebdii.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
GRI-LRZ-02	Computation	Linux Cluster	> 2800 cores		yes (but only for internal	Cerlane
GRI-SCAI-01	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 3 – 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	gLite operations	gLite software stack	yes (EGI accounting)	André
DAT-KNMI-01	Data	ORFEUS Data Center	Virtual European Broadband Seismograph Network,European Integrated Data Archive	Web services, interactive sessions	ArcLink		
DAT-IPGP-04	Data	COHER-D	IPGP Data Archive, NIED (Japan)	direct access from IPGP network		no	Geneviève

B Software components

Table 4 – Components (recommended by SA2) / resource

	TestBed										
	HPC-LRZ-01	HPC-CINECA-03	HPC-CINECA-04	DEP-UEDIN-01	DEP-ULIV-01	DEP-IPGP-06	DEP-INGV-01	DEP-SCAI-01	GRI-SCAI-01	GRI-IPGP-03	GRI-LRZ-02
Codes											
SeisSol	✓	✓	◐						●		✓
SpecFEM3D	✓	✓	●			◐	●	✓			✓
Scientific libraries											
Matplotlib	✓	●		●		●	✓	✓	✓	✓	✓
ObsPy	✓			●		✓	●	✓	✓	✓	✓
Metis	✓	●				●		✓	✓	✓	✓
GMT	✓				●	✓	✓	✓	✓		✓
mpi4py	✓	✓	●	✓		◐	●	✓	✓		✓
NetworkX	✓	✓		✓		✓		✓	✓	✓	✓
h5Py	✓	●	◐			✓	●	✓	✓		✓
dispel4py			◐					✓	✓		
Paraview	✓		●			✓	●	✓			✓
FFmpeg	✓	●	●					✓			✓
Environment / Tools											
iRODS	✓	✓		✓			✓	✓	✓	✓	✓
iRODS-DSI		✓		✓			✓	✓	●	✓	
GSI-SSH	✓	✓					●	✓	✓	●	✓
GridFTP	✓			●			✓	✓	✓	✓	✓
GRAM	✓						✓	✓	✓		✓
JSAGA									✓	✓	
gLite CREAM									✓	✓	
gLite WMS									✓		

◐ Currently being installed ● Installed

✓ Available for VERCE users

An empty cell means not applicable or undecided as of yet