



EGI-Engage

EPOS EGI pilot integrating computational resources and services

D6.19

Date	20 February 2017
Activity	SA2
Lead Partner	INGV
Document Status	FINAL
Document Link	https://documents.egi.eu/document/3023

Abstract

The EPOS Competence Centre (CC) of the EGI-Engage project drives collaboration between EGI and the European Plate Observing System (EPOS) service developers and providers, in order to collect, analyse and compare Earth Science community needs with EGI technical offerings. The work is driven by three specific use cases that were selected by the CC members in the first project year. This document describes the three pilot demonstrators that have been developed during the second year based on the use cases. The pilots highlight how EPOS could benefit from EGI services and provide the first suggestions about how services from the EGI infrastructure could be connected to the EPOS Research Infrastructure as Distributed Integrated Core Services (ICS-D). A roadmap for future enhancements of the pilots is also defined in the document for each of the pilots, together with plans for exploiting and disseminating the work done so far to the wider EPOS community.



This material by Parties of the EGI-Engage Consortium is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

The EGI-Engage project is co-funded by the European Union (EU) Horizon 2020 program under Grant number 654142 <http://go.egi.eu/eng>

COPYRIGHT NOTICE



This work by Parties of the EGI-Engage Consortium is licensed under a Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>). The EGI-Engage project is co-funded by the European Union Horizon 2020 programme under grant number 654142.

DELIVERY SLIP

	<i>Name</i>	<i>Partner/Activity</i>	<i>Date</i>
From:	Daniele Bailo	INGV/WP6	20/02/2017
Moderated by:	Gergely Sipos	EGI Foundation /WP6	
Reviewed by	Luca Postpischl Jan Michalek	INGV UIB	23/02/2017 24/02/2017
Approved by:	AMB and PMB		7/03/2017

DOCUMENT LOG

<i>Issue</i>	<i>Date</i>	<i>Comment</i>	<i>Author/Partner</i>
v.1	20/02/2017	Full draft ready for external review	Daniele Bailo/INGV Mariusz Sterzel/CYFRONET Alessandro Spinuso/KNMI Diego Scardaci/EGI F. - INFN
v.2	21/02/2017	Updated abstract, added document metadata	Gergely Sipos / EGI.eu (SA2 coord)
v.3	28/02/2017	Updated version based on reviewers' feedback	Daniele Bailo/INGV
FINAL		Final version	

TERMINOLOGY

A complete project glossary and acronyms are provided at the following pages:

- <https://wiki.egi.eu/wiki/Glossary>
- <https://wiki.egi.eu/wiki/Acronyms>

AH: Anthropogenic Hazard
 EPOS: European Plate Observing System
 GUI: Graphical User Interface
 ICS: Integrated Core Services
 ICS-C: Integrated Core Services Central Hub
 ICS-D: Integrated Core Services Distributed
 LtOS: EGI Platform for the long-tail of science
 TCS: Thematic Core Services
 VOMS: *Virtual Organization Membership Service*

Contents

1	AAI use case	7
1.1	Introduction	7
1.1.1	Scientific use case description	8
1.1.2	E-infrastructure requirements	9
1.2	Service architecture.....	9
1.2.1	High-Level Service architecture	10
1.3	Release notes	11
1.3.1	Requirements covered in the release	11
1.4	Feedback on satisfaction	11
1.5	Plan for Exploitation and Dissemination.....	11
1.6	Future plans	12
2	Earthquake simulation use case (MISFIT).....	13
2.1	Introduction	13
2.1.1	Scientific use case description	14
2.1.2	E-infrastructure requirements	14
2.2	Service architecture.....	15
2.2.1	High-Level Service architecture	15
2.2.2	Integration and dependencies	16
2.3	Release notes	18
2.3.1	Requirements covered in the release	18
2.4	Plan for Exploitation and Dissemination.....	18
2.5	Future plans	20
3	Satellite Data use case	21
3.1	Introduction	21
3.1.1	Scientific use case description	22
3.1.2	E-infrastructure requirements	22
3.2	Service architecture.....	23
3.2.1	High-Level Service architecture	23

3.3 Release notes 24

 3.3.1 Requirements covered in the release 24

3.4 Feedback on satisfaction 24

3.5 Plan for Exploitation and Dissemination..... 24

3.6 Future plans 25

Executive summary

The EPOS Competence Centre of the EGI-Engage project drives collaboration between EGI and the European Plate Observing System (EPOS) service developers and providers in order to collect, analyse and compare Earth Science community needs with EGI technical offerings. The work is driven by three specific use cases that were selected by the CC members in the first project year. This document describes the three pilot demonstrators that have been developed during the second year based on the use cases. The pilots highlight how EPOS could benefit from EGI services and provide the first suggestions about how services from the EGI infrastructure could be connected to the EPOS Research Infrastructure as Distributed Integrated Core Services (ICS-D). A roadmap for future enhancements of the pilots is also defined in the document for each of the pilots, together with plans for exploiting and disseminating the work done so far to the wider EPOS community.

The implemented pilot demonstrators are the following:

- AAI: the envisaged model for the EPOS AAI, based on UNITY IDM and designed in collaboration by EPOS and EGI expert, has been implemented and a prototype is available. In addition, interoperability with EGI infrastructure has been proved integrating the EGI CheckIn service with such prototype. Then, the EGI infrastructure could act as ICS-D in the wider EPOS architecture.
- Earthquake simulation (MISFIT): this pilot shows how the back-end services of an existing application, MISFIT, in the field of Computational Seismology can be improved through the integration with the EGI Federated Cloud. Furthermore, it also demonstrates how EGI and EUDAT services can jointly serve a research community. Indeed, software previously developed by the VERCE project has been integrated with the computing services of the EGI Federated Cloud, using data from EIDA / ORFEUS organization via EUDAT data preservation services (B2SAFE & B2STAGE).
- Satellite Data: it is related to the services that will be offered by the EPOS satellite data TCS to the wide range of EPOS users. This TCS deal with the processing of the Earth Observation datasets collected by various satellites, including the Sentinels of the Copernicus programme to address several societal challenges. Aim of this use case is creating a collaboration between Earth Observation and EGI infrastructure experts to create an environment that facilitates the development of new services. The demonstrator consist of a pilot EPOS service deployed on top of the Geohazard Thematic Exploitation Platform, developed by Terradue on behalf of ESA and part of the EPOS satellite data TCS, that is now interconnected to the EGI Federated Cloud to exploit its computing and storage resources (work done in collaboration with the EGI-Engage task SA1.3¹). The selected demonstrator is P-SBAS InSAR Sentinel-1 TOPS developed by CNR-

¹ The e-Collaboration for Earth Observation (e-CEO) platform: https://wiki.egi.eu/wiki/EGI-Engage:WP5#TASK_SA1.3_Integration.2C_Deployment_of_Grid_and_Cloud_Platforms

IREA². P-SBAS stands for Parallel Small BAseLine Subset and it is a DInSAR processing chain for the generation of Earth deformation time series and mean velocity maps, which uses as Input SLC (Level-1) Sentinel-1 data.

² <http://www.irea.cnr.it/en/>

1 AAI use case

1.1 Introduction

EPOS main aim is to coordinate, collect, archive high quality Earth Science data across Europe. By a definition EPOS is a distributed Research Infrastructure where Data, Data Products, Software and Services (DDSS) are provided by different communities in the domain of the solid Earth sciences. In this framework, EPOS envisages the construction of a central hub called “Integrated Core Services” (ICS) which aggregates all DDSS from various disciplines³. From the technical viewpoint, DDSS are provided by a distributed network of endpoints (Thematic Core Services, TCSs), which use heterogeneous authorization mechanisms. Users access the ICS querying for some data/dataproducts/software/service, and ICS is delegated to fetch the resources on behalf of the user. ICS can be connected to one or more ICS-D, the latter is an external infrastructure that could provide additional resources, such as computing and/or storage, to the EPOS RI.

The main purpose of this use case is to provide a framework, to be used in the ICS hub, which enables any user to access the ICS with his/her preferred authorization mechanism (e.g. OAuth, eduGAIN, X509 certificates etc.), delegate this identity to ICS to fetch resources from the various endpoints (TCSs). These endpoints may implement heterogeneous authorization mechanisms, and the system should realise harmonised access to them. Furthermore, such use case shows how the ICS hub could be integrated with the AAI mechanisms of the EGI infrastructure enabling its usage as ICS-D within the EPOS RI.

The prototype mimics the functionality of a target ICS-C/D platform. After evaluation of different existing solutions to define the EPOS AAI infrastructure, the prototype has been chosen to be based on the solution implemented by the EGI Access Platform⁴: the Unity IDM⁵ framework and special proxies generated by robot certificates (so called Per-User Sub-Proxies) to access X509-based resources (typically for heavy job execution on HTC/HPC and cloud systems).

Tool name	Prototype of the EPOS Authorisation Authentication Infrastructure (AAI)
Tool url	https://epos-aa1.cyfronet.pl/home/home
Tool wiki page	Not applicable
Description	The service aims to provide the Authentication and Authorisation Infrastructure service prototype for EPOS. For this purpose, the AAI

³ EPOS-IP WP6 & WP7 teams. (2015). ICS-TCS Integration Guidelines - Handbook for TCS integration: Level-2. Zenodo. 10.5281/zenodo.34666

⁴ <http://access.egi.eu>

⁵ <http://www.unity-idm.eu/>

	backbone of the EGI LToS have been adopted and, the AAI prototype is thus based on Unity IDM technology. For the sake of simplicity, the EGI AAI is used as Identity provider for EPOS AAI prototype
Value proposition	EPOS infrastructure is currently under heavy development. The knowledge transfer from EGI to EPOS concerning AAI is meant to speed up development of the AAI service(s) and thus ease the integration between Integrated Core Services and Thematic Core services and underlying national research infrastructures.
Customer of the tool	Whole EPOS community, since Integrated Core Services, through Thematic Core services up to national (Earth science) research infrastructures.
User of the service	Large scientific community
User Documentation	Not applicable
Technical Documentation	http://www.unity-idm.eu , see also https://wiki.egi.eu/wiki/LToS
Product team	ACC Cyfronet AGH, Istituto Nazionale di Geofisica e Vulcanologia
License	Open Source, TBD
Source code	http://www.unity-idm.eu

1.1.1 Scientific use case description

User Story	<ol style="list-style-type: none"> 1. The user access to ICS-C portal for the first time, therefore s/he registers to EPOS ICS-C user database. 2. The user logs in with his/her credentials (different technologies IdP enabled, for instance X509 certificates, OpenIdConnect, eduGAIN). 3. Currently TCS AH and EGI LToS are supported 4. User does some simple but multidisciplinary data discovery (i.e. accessing to at least two types of data from different domain and TCSs, say seismological waveforms from seismology TCS and events from AH TCS) 5. S/he gets the complete list of results (e.g. data-objects, files in this case) and selects some of them to be downloaded 6. S/he obtains the data (e.g. download as zipped/tar format or simply in the native file format).
(Potential) User base	<p>The potential user base is composed by all users interested in the solid earth sciences, and in particular by: a) Data and service providers, b) Scientific user community, c) Governments, d) Private sector, and e) Society.</p> <p>Each of those stakeholder categories can interact with the system in a different way and therefore be identified as:</p> <ul style="list-style-type: none"> • Active users, those actively using the system. The majority of these users will be registered. We estimate this category to be the 20% of the total engaged. • Occasional users, those occasionally using the system in a “lightweight” mode. We estimate this category to be the 40% of the total engaged users in each year.

	<ul style="list-style-type: none"> • Sporadic users, those rarely using the system and for no specific purpose (e.g. they are simply curious, etc.). We estimate this category to be the 40% of the total engaged users in each year. <p>As for the number of users for each stakeholder category, a systematic study is being carried on, considering that: a) EPOS is a system under development and the ACTIVE user base will be increasing in time and b) the potential number of users can be enormous if we consider as “interested stakeholders” the total number of participants to huge geological meetings as EGU (12.500 registered users) or AGU (20.860 registered users).</p>
--	--

1.1.2 E-infrastructure requirements

HW Resources	This use case does not require any specific HW resources for processing. The amount of resources to host the AAI services will be provided by ICS-C resources. It is required however, high availability of the EGI CheckIn service, on which large number of EPOS users will rely.
SW Resources	The use case does not have particular SW requirement from EPOS. The SW requirements depend on the exact solution to be used and offered from EGI.
Cost of delivery	The service will be provided by EPOS ICS-C. Thus, the cost will be associated with the EPOS AAI services solely. However, due to dependence of the EGI CheckIn we expect some extra costs for covering it.
Operational aspect	The use case is foreseen as proof of concept only. Once integrated within the EPOS infrastructure, it will be operated by EPOS ERIC ICS-C.

1.2 Service architecture

In order to provide a robust and efficient architecture, after the description of the use case, requirements were gathered from the community, and general consideration were done:

1. Taking into account dispersed character of the EPOS infrastructure, including scientific data sets, high level services, computing facilities, etc. users should be provided with AAI system capable to interact with all the EPOS resources,
2. Several TCSs, RIs, etc., have already developed some AAI solutions. These will have to be connected to the general EPOS AAI in order to avoid break-up of already existing services.
3. To connect existing AAI approaches in EPOS into one ecosystem, AAI hub is needed which will assure interoperability between existing technologies. Hub technologies passing logins, passwords, etc. have to be abandoned by principle. Instead a modern, attribute based, solution is advised. Once user is authenticated within the infrastructure, all the authorisation should be done by attributes only. The Set of the EPOS specific attributes should be defined in the hub.

4. Current IdPs EPOS is aware of include: EduGAIN, LDAP, OpenID, x509 certificates.

The architecture in the picture below shows the EPOS AAI solution developed so far: it is based on the AAI backbone of the EGI LToS.

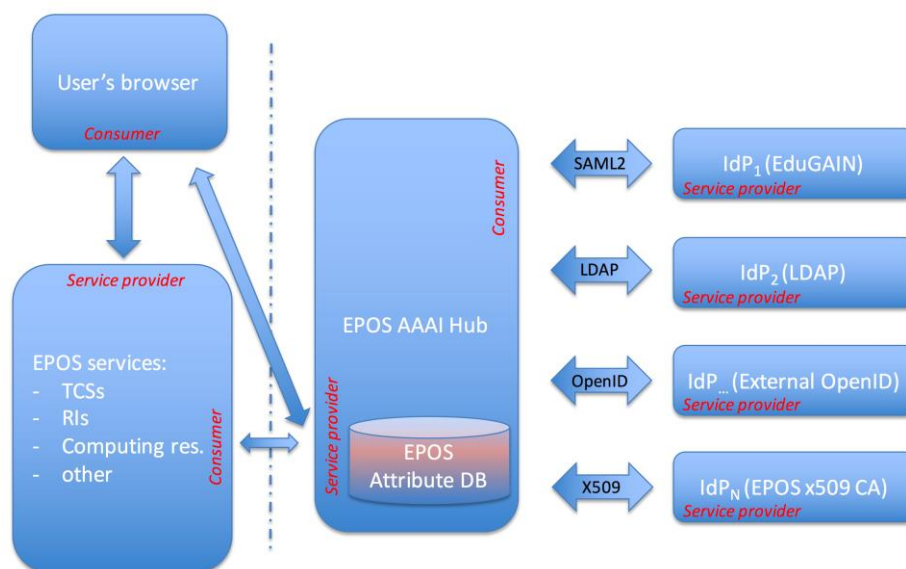


Figure 1. The proposed AAI scheme for EPOS

1.2.1 High-Level Service architecture

The prototype has been built according to the scheme sketched on Figure 2. In the current implementation, two IdPs have been integrated as a proof of concept including EGI CheckIn and TCS AH. Additional IdPs will be integrated within the EPOS ICS-C after initial integration with the ICS-C science gateway. The thick arrows on the scheme represent any further communication (between TCSs, ICS-C and EGI CheckIn) after the authorization has been completed. I.e. this includes job submission with data analysis, processing, advanced visualisations, etc.

Integration with the EGI CheckIn demonstrates how EPOS RI users could exploit EGI services and resources through the ICS in a transparent way. More coupled integration scenarios will be investigated in the future.

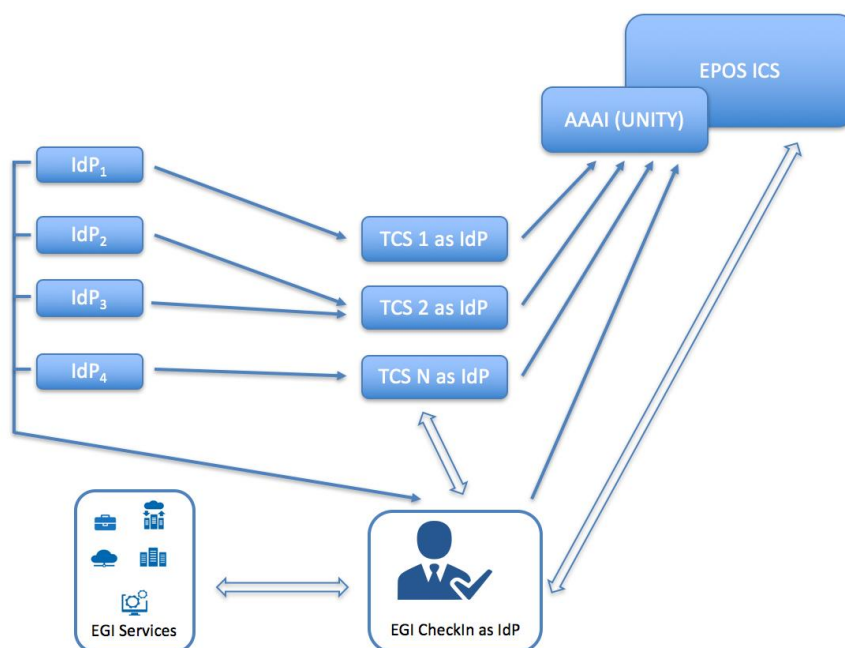


Figure 2. EPOS AAI prototype scheme

1.3 Release notes

1.3.1 Requirements covered in the release

- EPOS AAI Hub for ICS able to:
 - Authenticate users coming from different domains through different IdPs;
 - Register users into an internal database using a unique format;
 - Integrate third part services automatically forwarding the user identities stored in the internal database;
- As proof of concept, EGI CheckIn and TCS AH have been integrated in the Hub.

1.4 Feedback on satisfaction

The prototype has been presented to the EPOS IT teams (EPOS IP project WP6 and WP7 packages). Access to the prototype is open for anyone with valid EGI or TCS AH credentials. After the presentation at the TCS-ICS meeting in Prague in March 2016, the prototype has been accepted by EPOS WP6 and WP7 teams as a general solution of EPOS AAI. The first version, integrated with the EPOS ICS-C portal, is expected in a fall 2017, together with a second version of the ICS-C science gateway.

1.5 Plan for Exploitation and Dissemination

The AAI prototype will become an internal component of the ICS-C science gateway and will be transparent for the EPOS end users. The previous subsection described the road that leads there.

1.6 Future plans

The major planned step is integrating the prototype with the EPOS ICS-C science gateway. As mentioned, according to EPOS IT teams plans, this step shall be done by fall of 2017. Second stage will include integration of the remaining Thematic Core Services with the ICS-C platform. Final step will bind the computing resources provided by EGI with ICS-C and TCSs via ICS-D services.

2 Earthquake simulation use case (MISFIT)

2.1 Introduction

The use case enables the processing and the comparison of data resulting from the simulation of seismic wave propagation following a real earthquake and real measurements recorded by seismographs. While the simulation data is produced directly by the users and stored in a Data Management System, the observations need to be pre-staged from institutional data-services, which are provided by the international seismic-data archives of the FDSN⁶ consortium. Users can interactively select the data of interest, compose and execute processing pipelines and conduct eventually the MISFIT analysis between the synthetics and the observed-data streams. The final scope of the systems is to support the researchers with the study and the improvement of regional and global Earth Models. The target system is a comprehensive computational platform developed in the context of the EC funded project VERCE. It provides a science gateway and technological stack that is endorsed and currently supported by the EPOS-IP initiative.

The pilot aims at extending the computational capabilities of the VERCE platform with the EGI Federated Cloud resources through a required management middleware.

Tool name	EPOS Computational Earth Science – Seismology - VERCE Portal
Tool url	http://portal.verce.eu
Tool wiki page	http://www.verce.eu
Description	Virtual Earthquake and seismology Research Community e-science environment in Europe
Value proposition	The platform generally aims at facilitating the integration of computational resources to be exploited by seismological applications, adopting standard interfaces towards cloud and HPC. Integrating the EGI Federated Cloud would extend the possibilities of procurement, accounting and sustainability of the service.
Customer of the tool	Universities, International Research Centres and Earthquakes Monitoring Centres.
User of the service	The main users of the platform are researchers in computational seismology, teachers in Geodynamic, students and analysts. The latter, together with decision-makers benefit from the execution of rapid simulations and analysis, for the refined evaluation of seismic scenarios.
User Documentation	http://www.verce.eu/Training/UseVERCE.php
Technical Documentation	http://www.verce-project.eu (requires authentication)
Product team	Fraunhofer SCAI, INGV, ULIV, KNMI, CNRS

⁶ <http://www.fdsn.org/>

License	OpenSource
Source code	https://github.com/KNMI/VERCE

2.1.1 Scientific use case description

User Story	<ol style="list-style-type: none"> 1. Through the VERCE Science Gateway, user selects for which simulation results he/she would like to download the observed raw data. 2. User triggers the execution of a download workflow that will pre-stage the observed data in a dedicated storage space. 3. User combines observed data and simulation results then configures a processing pipeline. 4. User triggers the execution of the pipeline workflow which will ingest and process the observed and synthetic data 5. User selects the pre-processing results and specifies the parameters for the misfit analysis 6. User triggers the execution of the misfit processing workflow 7. Progress of the computations can be always monitored and the results and the associated metadata visualised at runtime or offline
(Potential) User base	<p>The current deployment of the platform counts ca. 110 users. Users mostly belong to the seismological community and are typically students, academics and researchers. We recognize that a certain number are also IT experts who are interested in the platform’s generics.</p> <p>Potential user base can be reached by training sessions that combine scientific lectures with the acquisition of technical skills aiming at clear research targets.</p>

2.1.2 E-infrastructure requirements

HW Resources	<p>Apart from the existing infrastructure (Gateway, Middleware backend services, VOMS server, iRODS infrastructure), the use-case only requires additional Cloud VMs for carrying out the processing tasks. Of course, the requirements for the processing will depend on the amount of jobs and the data volume to be processed. We foresee that a low number of VMs (1-10) will be suitable for the evaluation purposes.</p>
SW Resources	<p>SW Product: ObsPy Technology Provider: The ObsPy Development Team (devs@obs.py.org) License: GNU Lesser General Public License, Version 3⁷</p> <p>SW Product: Dispel4Py Technology Provider: The Dispel4Py Development Team (rosa.filgueira@ed.ac.uk) License: Apache License, Version 2.0⁸</p>

⁷ <http://www.gnu.org/copyleft/lesser.html>

	<p>SW Product: WS-PGRADE gUSE (grid and cloud user support environment) Technology Provider: LPDS Sztaki (portalsupport@lpds.sztaki.hu) License: Apache License, Version 2.0</p> <p>SW Product: Globus Toolkit Technology Provider: The Globus Alliance License: Various⁹</p>
Cost of delivery	<p>Extension of the current verce.eu VO with Federated Cloud attributes. Enablement of cloud-friendly submission from WS-PGRADE workflow towards these cloud resources. Contextualisation of generic VMs suitable for the tasks or delivery of full-fledged dedicated VMs. Instances can be used by single users or in pools.</p> <p>The envisaged cost of delivery is estimated in about 20 person-months.</p>
Operational aspect	<p>The existing services (Gateway, Middleware backend services, VOMS server, iRODS infrastructure) are operated by SCAI which is member of NGI-DE.</p> <p>The EGI FedCloud resources to be used are operated by CNRS IPHC whose participation is supported by the French NGI.</p> <p>Until now, there have been no further negotiations with NGIs about additional support.</p>

2.2 Service architecture

The VERCE virtual research environment (VRE) requires communities driving the research, collections of relevant data, application software tuned and maintained to meet the latest research requirements and to exploit hardware advances, and teams of ICT experts maintaining the VRE’s advanced capabilities. Its usability and sustainability is a key to attracting these researchers, enabling their collaboration, creating effective interplay between ICT experts and domain-focused researchers, gaining access to the resources required and amortising costs over sufficiently broad communities. The design of the VERCE platform embraces this diversity.

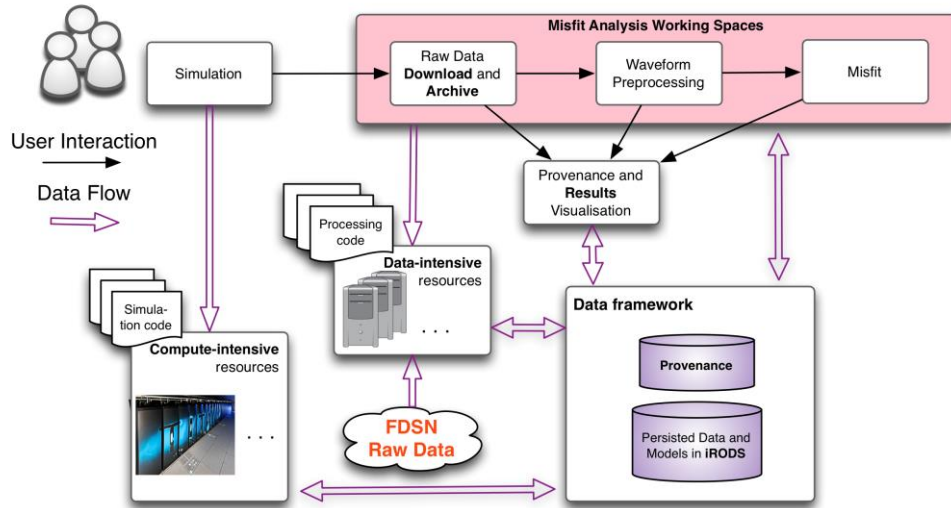
2.2.1 High-Level Service architecture

The platform offers a number of application specific services to perform operations such as simulation, raw data pre-staging from FDSN archives, pre-processing and MISFIT. These high level services are organised into independent workspaces, which operate through an interaction with platform’s webservices and workflows. The infrastructure serve these interfaces with

⁸ <http://www.apache.org/licenses/>

⁹ <http://toolkit.globus.org/toolkit/docs/5.2/5.2.3/licenses/>

computational and data resources. We provide below a high-level overview of the architecture including user interaction and data-flows.



The efforts of the team are focusing on extending the Data-Intensive component shown by the illustration with EGI Federated Cloud capabilities.

2.2.2 Integration and dependencies

This section provides a System-centric overview of the VERCE integrated software and dependencies.

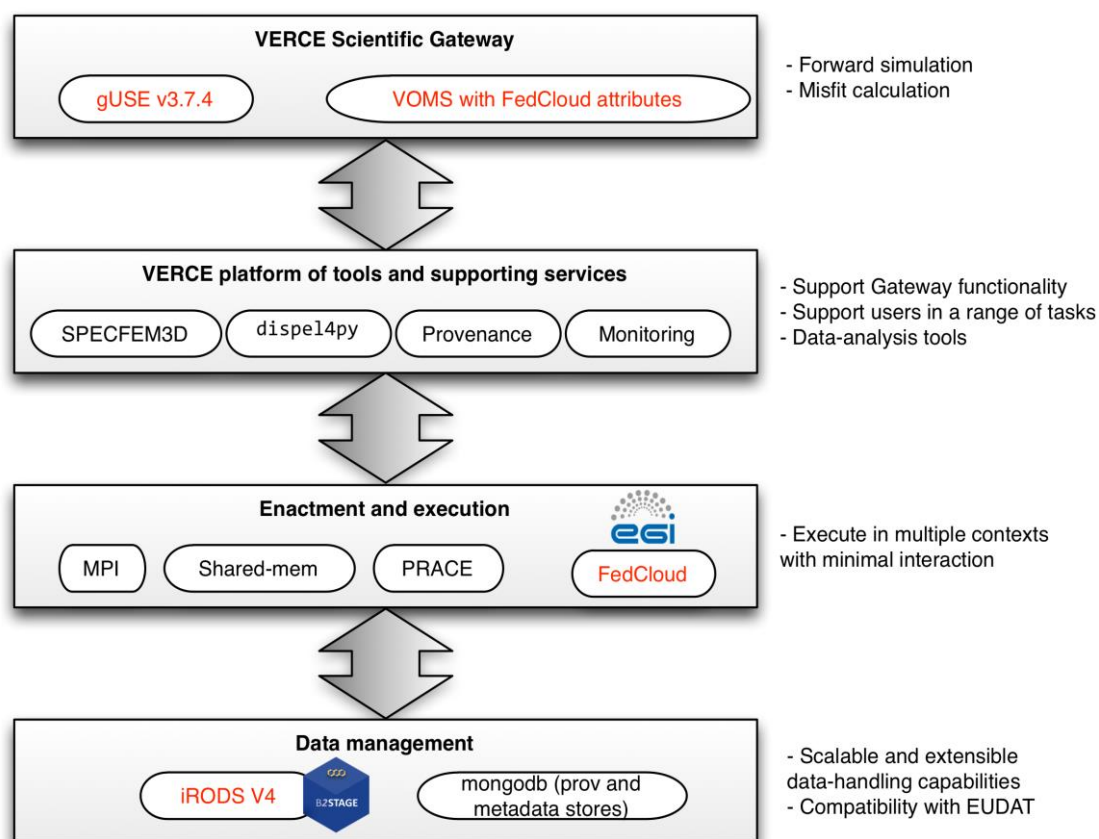


Figure 3. VERCE integration and dependencies: In red the updated components

VERCE Science Gateway

gUSE is a Liferay-powered science-gateway framework. The support of VOMS and x.509 certificates is a clear security requirement, which has been met both through gateway/portal technologies.

Platform tools and supporting services

- **SPECFEM3D:** a widely used tool for simulating seismic wave propagation¹⁰. The support of seismic wave propagation simulation for SPECFEM3D was requested by the seismology research community.
- **dispel4py:** Python framework for describing abstract data-intensive workflows. It is used for all the data processing tasks in shared-memory as well as MPI-powered clusters.
- **The VERCE Provenance service:** It is used on the gateway and it interacts with a number of architectural components, such as the corresponding data management elements shown under “Data Management”. It follows the W3C-PROV standard¹¹.

¹⁰ <http://geodynamics.org/cig/software/specfem3d>

- **Monitoring of execution:** a service offered on the gateway, interacting with the provenance solution.

Enactment and Execution

These are architectures and computational services that the VERCE solution is able to make use of, thus adding to sustainability and standardization. We extended this portfolio with the possibility of integrating resources provided by the **EGI FedCloud**

Data Management

iRODS¹² is a rule-based data-management system order to address requirements to do with storage and access at scale, with configurable access policies. We have updated the current setup to the iRODS v4. This brought improvements also to the GridFTP support for data-staging operations to and from iRODS, thanks to the integration of the **EUDAT B2STAGE** technology. MongoDB¹³ is a distributed and scalable document database used by the provenance and monitoring subsystems.

2.3 Release notes

2.3.1 Requirements covered in the release

- Setup the IPHC (@IN2P3/CNRS) cloud site to allow verce.eu VO members usage of its resources
- Setup the IPHC network to allow relevant traffic to pass (http/https/GUSE)
- Install GUSE images from EGI application database
- Add the IRES-IN2P3 (IPHC) cloud resource site to the VERCE development portal
- Upgrade the GUSE/WSPGRADE portal at SCAI to version 3.7.4, supporting the OCCI interface and, more generally, VMs with no automatically assigned floating IP addresses

2.4 Plan for Exploitation and Dissemination

As already mentioned, the Computational use case in the EPOS Competence Center builds upon results of previous work and aims at enhancing and providing computational resources to the VERCE platform, whose goal is facilitating the integration of computational resources to be exploited by seismological applications.

In such a framework, the exploitation and dissemination activities take advantage of already existing plans and user base, which need to be expanded.

¹¹ <http://www.w3.org/TR/prov-overview>

¹² <http://irods.org>

¹³ <http://www.mongodb.org>

The aim of the dissemination strategy is to maximize the exposure, impact and usage of the VERCE platform and of all the connected resources representing an extremely valuable tool for seismological and in general geophysical analyses in general. The key objectives of the strategy are:

- Increase the visibility of the use case and the platform by advertising them at national and international meetings, educational schools and public events;
- Regularly distribute and share updates about the platform;
- Transfer knowledge through training materials and dedicated events, educational resources and the organisation of workshops;
- Support the development of a strategy for the exploitation of the key outcomes of VERCE both for an increasing number of seismological applications and for other geophysical fields that will benefit from the VERCE accomplishments.

In order to achieve these goals, the existing material (summarized at the end of the paragraph) is a useful starting point. In particular, the VERCE website provides a complete and detailed view of all the platform functionalities, resources and services and can be used to advertise its possible applications and its important societal implications. Moreover, the website has a dedicated section to access training presentations and videos that we distributed during the webinars and the face-to-face training events of the VERCE project and that explain step-by-step (with included demos) how to use the platform and its services and how to access and exploit the produced output. The user guide of the portal and the VERCE legacy video are also available for a thorough overview. Finally, a section of the VERCE website that links to all the project-specific publications and to other related papers is also available.

Based on this material, our exploitation and dissemination plan is organised as follows:

- Review the already available material in order to update it to the latest release of the portal and to its new functionalities (e.g., misfit calculation, simulation at global scale, ...);
- Disseminate and advertise this training material through all the available channels (e.g., mailing lists of interested users, newsletters, project website, etc.) in order to attract the attention of domain specific scientific communities and of more general geophysical users;
- Organise new interactive webinars addressed to a broader audience thanks to the framework of EPOS-EGI-EUDAT projects;
- Advertise the use case and the portal at forthcoming conferences and meetings;
- Advertise the use of the platform for solving scientific issues as the analysis of seismic waves generated by the 2016 Central Italy seismic sequence;
- Support and enhance scientific publications on national and international journals and magazines to highlight the impact and potentialities of the VERCE tool;

- Organise a large face-to-face training event open to a broad public with theoretical presentations, demos and interactive sessions.

The structure and contents of our dissemination plan will be addressed in order to reach different categories of audience often interconnected with each other:

- Domain scientists as well as members of the broader earth sciences community;
- Industry actors in the earth sciences, and in computing and information technologies;
- National public services, policy makers and governments.

Summary of the dissemination material already available:

- VERCE website <http://www.verce.eu>
- VERCE portal user guide <http://www.verce.eu/Training/UseVERCE/VERCEPortal-UserGuide-vers1.1.pdf>
- Trainings' presentations and videos <http://www.verce.eu/Training/UseVERCE.php>
- Comprehensive legacy video of the VERCE project
- <https://www.youtube.com/watch?v=t9Adh1042-M&feature=youtu.be>
- VERCE-specific and correlated publications:
<http://www.verce.eu/PublicDissemination/Publications.php>

The described use case will also be validated and used in the context of EPOS ICS-C. This process will start after October 2017.

2.5 Future plans

More communities within EPOS are looking with interest at the capabilities of the platform and at the outcome of the EGI use case. Currently a plan is developing to include Volcanological research applications in the same VERCE framework, with the scope of scaling all the generic components to a larger user-base and going towards the characterization of the EPOS Computational Earth-Science strategy.

3 Satellite Data use case

3.1 Introduction

This use case is related to the services that will be offered by the EPOS satellite data TCS to the wide range of EPOS users: a) Data and service providers, b) Scientific user community, c) Governments, d) Private sector, and e) Society.

The satellite data TCS services deal with the processing of the Earth Observation datasets collected by various satellites, including the Sentinels of the Copernicus programme to address several societal challenges related to climate, natural hazards and risk assessment and, efficient usage of resources.

Satellites are very extensive and complex sources of information and their exploitation requires advanced knowledge on Earth Observation systems. IT platforms can make access to satellite data possible for a broader research community, and therefore the EPOS satellite data TCS decided to host services on top of some ESA Thematic Exploitation Platforms (e.g. Geohazard TEP).

In the context of this use case, a pilot EPOS service will be deployed on top of the Geohazard TEP that is now interconnected to the EGI Federated Cloud to exploit its computing and storage resources. The work to integrate the Geohazard TEP with the EGI Federated Cloud has been done under the EGI-Engage task SA1.3¹⁴.

The selected demonstrator is P-SBAS InSAR Sentinel-1 TOPS developed by CNR-IREA¹⁵. P-SBAS stands for Parallel Small Baseline Subset and it is a DInSAR processing chain for the generation of Earth deformation time series and mean velocity maps, which uses as Input SLC (Level-1) Sentinel-1 data. Such processing chain generates an output in csv format: Line Of Sight (LOS) Displacement time series, Mean LOS Velocity, Temporal Coherence, Average scatterer elevation (Topography). The service can also generate wrapped and unwrapped interferograms that are delivered in geoTiff format. This pilot is planned to run in the EGI FedCloud from February/March 2017 to February 2018.

Tool name	P-SBAS InSAR Sentinel-1 TOPS
Tool url	N.A.
Tool wiki page	N.A.
Description	P-SBAS stands for Parallel Small Baseline Subset and it is a DInSAR processing chain for the generation of Earth deformation time

¹⁴ The e-Collaboration for Earth Observation (e-CEO) platform: https://wiki.egi.eu/wiki/EGI-Engage:WP5#TASK_SA1.3_Integration.2C_Deployment_of_Grid_and_Cloud_Platforms

¹⁵ <http://www.irea.cnr.it/en/>

	series and mean velocity maps, which uses as Input SLC (Level-1) Sentinel-1 data.
Value proposition	Describe how the new or changed service alleviates specific user pains and/or supports its intended customer(s) to exploit new opportunities
Customer of the tool	Scientific community representatives, Governments, and Industries.
User of the service	Large research groups, individual researcher, Government agency employers, Industry employers.
User Documentation	N.A.
Technical Documentation	N.A.
Product team	CNR-IREA
License	Please provide license
Source code	N.A.

3.1.1 Scientific use case description

User Story	<ol style="list-style-type: none"> 1. A user accesses the Geohazard TEP to exploit one of the services hosted by the platform. 2. The user requires some data via the selected services. 3. The user’s request triggers a computation. 4. The Geohazard TEP create a certain number of VMs into the EGI FedCloud to perform the computation 5. The Geohazard TEP will copy the input data needed for the computation into the EGI FedCloud storage 6. The Geohazard TEP retrieve the output data from the EGI FedCloud 7. The Geohazard TEP provides the requested data to the user.
(Potential) User base	<p>The potential user base is composed by all users interested in the solid earth sciences, and in particular by: a) Data and service providers, b) Scientific user community, c) Governments, d) Private sector, and e) Society.</p> <p>See the AAI use case for more details.</p>

3.1.2 E-infrastructure requirements

HW Resources	<p>The amount of needed HW resources requested by this pilot are the following:</p> <ul style="list-style-type: none"> • n. 10-12 virtual machines with 2-4 CPUs, 32 GB RAM, 100 GB of local disk. <p>Such resources will be provided by the EGI BEgrid-BELNET resource provider in the context of the SLA agreed between EGI and Terradue.</p>
---------------------	--

SW Resources	The SW resources needed are the P-SBAS InSAR processing chain, the ESA TEPs and related packages setup on top of them.
Cost of delivery	<p>Integration of the ESA Geohazard TEP into the EGI FedCloud could be done at different level.</p> <p>The one implemented in the EGI-Engage task SA1.3 is the “On-Demand” model: TEPs are deployed and operated by external organisation and use the EGI FedCloud to perform computations on demand:</p> <ul style="list-style-type: none"> • create VMs when a new computation should be performed • transfer input data to the EGI FedCloud • retrieve the output data • free all the resources allocated after the output data have been downloaded. <p>Different types of integration would be examined and assessed in the next future. For example, we will investigate how to create clone or big cache of the Satellite Data in the EGI infrastructure accessible by all the FedCloud sites to reduce the transfer time.</p>
Operational aspect	The ESA Geohazard TEP is currently in a pre-operation phase and is managed and operated by Terradue.

3.2 Service architecture

3.2.1 High-Level Service architecture

The P-SBAS InSAR Sentinel-1 TOPS demonstrator will run on top of the Geohazard TEP that has been developed outside this project. The Geohazard TEP implements the ESA Thematic Exploitation Platform architecture¹⁶, see the figure below.

¹⁶ <http://go.egi.eu/EP-OpenArchitecture>

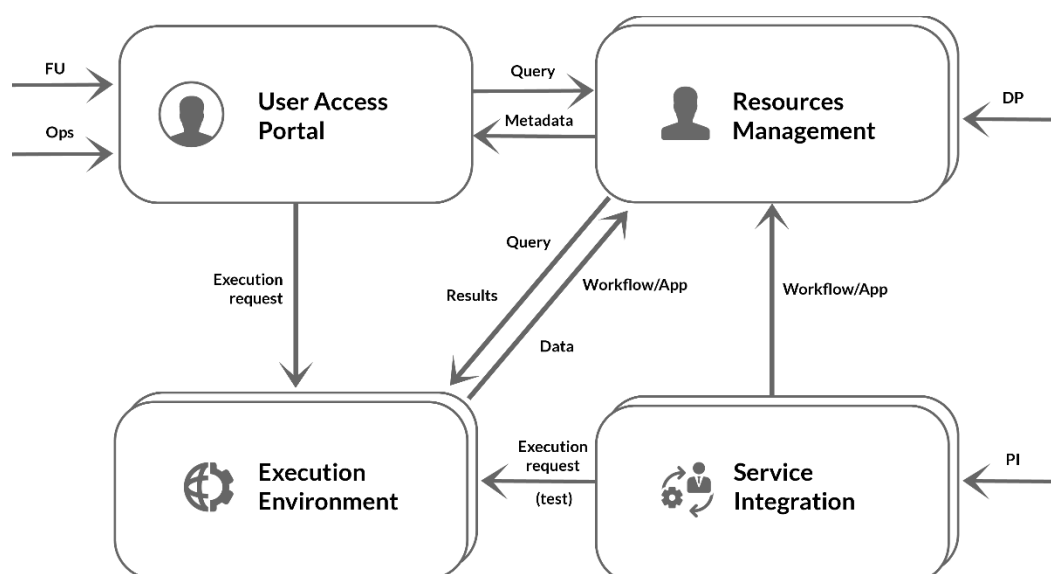


Figure 4. ESA Thematic Exploitation Platform architecture

3.3 Release notes

3.3.1 Requirements covered in the release

- Generation of Earth deformation time series and mean velocity maps from SLC (Level-1) Sentinel-1 data. The output will be in csv format and contain the following measurements: Line Of Sight (LOS) Displacement time series, Mean LOS Velocity, Temporal Coherence, Average scatterer elevation (Topography).
- Generation of wrapped and unwrapped interferograms delivered in geoTiff format.

3.4 Feedback on satisfaction

The demonstrator will run later this year for the EPOS community. A feedback on satisfaction will be requested to the users that will try the service.

3.5 Plan for Exploitation and Dissemination

The deployment of the P-SBAS InSAR Sentinel-1 TOPS demonstrator within the Geohazard exploitation platform is an activity that has been performed outside the EGI-Engage project. The integration of the Geohazard exploitation platform, instead, has been done in the context of the task SA1.3.

Aim of this use case is demonstrating how EGI, through its Cloud Compute service, could offer the provisioning of cloud compute and storage resources to the EPOS community serving as ICS-D for many EPOS TCS and the EPOS ICS.

Such result will be advertised to the EPOS community with the aim of creating a long term collaboration between EGI and the future EPOS ERIC.

3.6 Future plans

The P-SBAS InSAR Sentinel-1 TOPS demonstrator will run during the remaining months of the EGI-Engage project and after. In this period, we will invite people of the EPOS community to test and validate the service.

In addition, the integration of the Geohazard Exploitation Platform with the FedCloud will be further exploited scheduling and running new Satellite Data pilots according to the needs of the EPOS community.