



EGI-Engage

e-Infrastructures integration report

D4.6

Date	04 March 2016
Activity	JRA2.3
Lead Partner	INFN
Document Status	FINAL
Document Link	https://documents.egi.eu/document/2672

Abstract

This deliverable reports the status of the integration of e-Infrastructures with the EGI infrastructure during the first year of the project. The report covers the activities of the integration with the EUDAT EC project, the Canadian Advanced Network for Astronomical Research (CANFAR), the gCube/D4Science data and a computational e-Infrastructure and the integration of accelerators into the infrastructure. The activities also include the update of the federation model of the EGI Federated Cloud necessary to facilitate the integration and collaboration with new cloud infrastructures.



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:	E. Fernandez	EGI.eu/JRA2	24/02/2016
Moderated by:	Małgorzata Krakowian	EGI.eu/NA1	
Reviewed by	Eric Yen	ASGC/PMB	2016-02-26
	Yannick Legre	EGI.eu/NA1	2016-02-15
	Alvaro López García	IFCA/JRA2	2016-02-18
Approved by:	AMB and PMB		2016-03-04

DOCUMENT LOG

<i>Issue</i>	<i>Date</i>	<i>Comment</i>	<i>Author/Partner</i>
v.1	29/12/2015	Initial TOC	E. Fernandez/EGI.eu
v.2	01/02/2016	Version after internal review	P. Fabriani/ENG D. Scardati/EGI.eu M. Verlatto/INFN, G. Taffoni/INAF, E. Fernández/EGI.eu M. Viljoen/EGI.eu
v.3	19/02/2016	Revision with input from external reviews	E. Fernández/EGI.eu
FINAL	24/02/2016	Final version	E. Fernández/EGI.eu

TERMINOLOGY

A complete project glossary is provided at the following page: <http://www.egi.eu/about/glossary/>

Contents

1	Introduction.....	5
2	EGI-EUDAT	6
3	CANFAR.....	8
3.1	Scope of cloud federation	8
3.2	The EGI CANFAR federation.....	8
4	gCube/D4Science.....	11
4.1	Evaluation of jOCCI.....	12
4.2	Virtual Appliances Creation	12
4.3	Prototype design and implementation	12
5	Accelerated computing.....	14
5.1	Accelerated computing support for HTC platform.....	14
5.2	Accelerated computing for Cloud platform	16
6	EGI Cloud Federation	18
6.1	Federating Clouds.....	18
6.1.1	EGI Federated Cloud realms	21
6.2	Integration of providers into the EGI Cloud	21
7	Future plans.....	24
7.1	EUDAT	24
7.2	CANFAR	24
7.3	gCube/D4Science.....	24
7.4	Accelerated Computing	25
7.5	EGI Cloud Federation.....	25

Executive summary

The activities on this period were focused on EUDAT, CANFAR and gCube/D4Science, accelerated computing and the integration of other infrastructures with the EGI Federated Cloud.

EUDAT collaboration selected relevant user communities who are already collaborating with both infrastructures from relevant European Research infrastructures in the field of Earth Science (EPOS and ICOS), Bioinformatics (BBMRI and ELIXIR) and Space Physics (EISCAT-3D) and produced a pilot for that successfully demonstrated the interoperability between EGI and EUDAT at the EGI Community Forum 2016 (Bari, IT). The integration work will continue in the next months by evolving this pilot to fully implement a universal use case that covers the user needs with respect to the integration.

The CANFAR integration related activities officially started at month 6 of the Project, with the definition of a roadmap for the federation focused on two main areas: Authentication and Authorization Infrastructure (AAI) and Data Federation. The work started with the analysis of core infrastructure services offered by EGI and CANFAR: AAI, Accounting, Monitoring services, Operation tools and services, Service registry and market place. Developments on the AAI services have started to produce an interoperable service between both infrastructures.

The gCube/D4Science integration activity was driven by the identification of D4Science specific use cases whose implementation on top of EGI Federated Cloud would maximize the impact on the management and usage of the D4Science infrastructure, therefore the work was focused on the achievement of quick scalability for a typical D4Science process such as the execution of data analysis (e.g. signal forecasting) against one or more input datasets. After the design of the overall architecture and components for the integration a prototype implementation of the components was done; and now these components are planned to be included in the next gCube release.

The accelerated computing task aims to provide GPGPU support both on HTC and Cloud platforms. For the HTC platform, a first prototype of GPGPU-enabled CREAM-CE with support for LRMS was developed and now is ready for integration in the next release. For the cloud platform, the task investigated the current technologies to prepare specialised VM images support in the EGI Federated Cloud. A fully integrated site is now providing GPGPU-enabled cloud computing resource in EGI Federated Cloud.

The EGI Cloud federation model was updated to accommodate different types of integration with the EGI services, in cloud realms. The currently existing EGI cloud federation becomes the Open Standards Realm and a new OpenStack Realm, supporting integration also with EGI services but with different interfaces, is now supported. Several integration activities with other infrastructures are identified and tracked in a dedicated queue in the EGI RT system.

1 Introduction

EGI-Engage WP4 objective is to further expand the EGI capacities and capabilities with the technical development of existing solution and by the integration of other e-Infrastructures. On the one hand, the technical development activities of the WP will further evolve the EGI Federated Cloud infrastructure and will introduce an Open Data Access platform. On the other hand, the integration activities will ensure support for a broad number of use cases and data commons needs by partnering with e-Infrastructures both located in Europe and worldwide and introducing the support for accelerated computing facilities (General-Purpose Graphical Processing Units or GPGPUs). This document reports the activity of these integration activities of WP4 in the last 12 months.

EGI-Engage identified three target e-Infrastructures and technical solutions for integration: EUDAT, the Canadian Advanced Network for Astronomical Research (CANFAR), and gCube and the D4Science infrastructure (D4Science).

EUDAT is an EC project that seeks to deploy a Collaborative Data Infrastructure driven by research communities needs. The EGI-EUDAT integration activity collaborates with EUDAT towards the harmonisation of the two infrastructures, including technical interoperability, authentication, authorisation and identity management, policy and operations.

The Canadian Advanced Network for Astronomical Research (CANFAR) is a digital infrastructure for Astronomy and Astrophysics (A&A) based on cloud storage and cloud processing middleware and on tools and services developed by the International Virtual Observatory Alliance (IVOA). The collaboration seeks to create a uniform platform for international astronomy research collaboration.

D4Science is both a Data and a Computational e-Infrastructure that hosts several Virtual Research Environments (VREs) for different communities and leverages the gCube toolkit. gCube is an open-source software toolkit used for building and operating Hybrid Data Infrastructures enabling the dynamic deployment of VREs by favouring the realisation of reuse oriented policies. The integration activities have focused in extending the gCube framework to use EGI Federated Cloud resources through implementing OCCl client capabilities.

Many EGI sites already provide accelerated computing technologies (GPGPUs or MIC co-processors) but these are not directly supported by the EGI platforms, and it is needed to interact with the local provider to enable dedicated access to those resources. This task seeks to enable these technologies to be accessible directly through EGI platforms (both HTC and Cloud).

Since the task is not limited to the mentioned infrastructures, contact with external partners has been established for the integration of other solutions and e-Infrastructures into EGI. As a result of these contacts, the EGI Cloud federation model has been updated to accommodate different types of integration.

2 EGI-EUDAT

EUDAT¹ is a collaborative Pan-European infrastructure providing research data services, training and consultancy for researchers, research communities, research infrastructures and data centres. EUDAT's vision is to enable European researchers and practitioners from any research discipline to preserve, find, access, and process data in a trusted environment, as part of a Collaborative Data Infrastructure (CDI) conceived as a network of collaborating, cooperating centres, combining the richness of numerous community-specific data repositories with the permanence and persistence of some of Europe's largest scientific data centres.

The EGI-EUDAT collaboration started in March 2016 with the main goal to harmonise the two infrastructures, including technical interoperability, authentication, authorisation and identity management, policy and operations. The main objective of this collaboration is to provide end-users with a seamless access to an integrated infrastructure offering both EGI and EUDAT services and, then, pairing data and high-throughput computing resources together.

To define the roadmap of this collaboration, EGI and EUDAT selected a set of relevant user communities who are already collaborating with both infrastructures. These user communities are able to bring requirements and help to assign the right priorities to each of them. In this way, the integration activity has been driven by the end users from the start. The identified user communities are relevant European Research infrastructure in the field of Earth Science (EPOS and ICOS), Bioinformatics (BBMRI and ELIXIR) and Space Physics (EISCAT-3D).

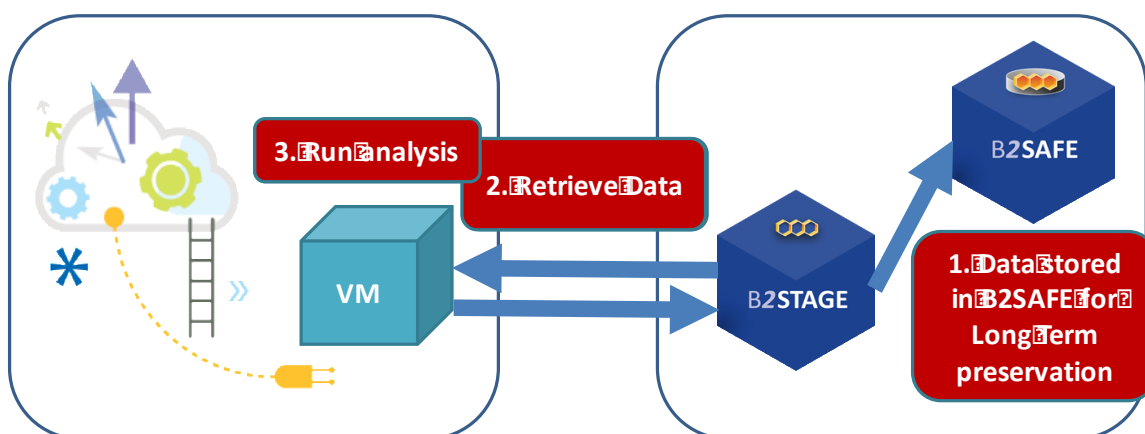
The first outcome of this activity has been the definition of a universal use case that covers the user needs with respect to the integration of the two infrastructures previously identified. This use case permits a user of either e-infrastructure to instantiate a VM on the EGI Cloud Federation for the execution of a computational job consuming data preserved onto EUDAT resources. The results of such analysis can be staged back to EUDAT storages, and if needed, allocated with Permanent Identifiers (PIDs) for future use. To implement all the steps of this use case the following integration activities between the two infrastructures has to be fulfilled: (1) harmonisation between the authentication and authorisation model, (2) definition and implementation of the interfaces between the involved EGI and EUDAT services.

The first step to implement the universal use case has been the development of a pilot that was demonstrated at the EGI Community Forum 2016 (Bari, IT). This pilot involves accessing both EGI and EUDAT services using a X.509 digital certificate proxy with its Distinguished Name (DN) registered on both infrastructures. Once the proxy certificate is created, the pilot uses this credential to:

¹ www.eudat.eu

- instantiate a virtual machine in the EGI federated cloud with standard Globus client tools preinstalled on the virtual machine
- transfer data from the virtual machine to EUDAT B2STAGE via the globus-url-copy command
- transfer the same data back from EUDAT B2STAGE to the virtual machine and verify integrity by comparing the checksums of the data
- transfer the data from B2STAGE to B2SAFE for long-term preservation

- **Access to EGI and EUDAT services with a single user identity**
- **Data Staging between EGI Federated Cloud and EUDAT services**



This pilot constitutes the first steps in proving that essential interoperability is possible between EGI and EUDAT. The integration work will continue in the next months evolving this pilot to fully implement the universal use cases adapting it accordingly to requirements that will be further collected by the potential users.

3 CANFAR

A&A community has gathered rich experiences in cloud computing within the CANFAR federated cloud deployed on Compute Canada resources and operated by National Research Council Canada. The Canadian cloud infrastructure represents a unique example of an A&A oriented infrastructure that joins together the Infrastructure as a Service (IaaS) and the standards and services developed by the IVOA² e.g. for user authentication and authorization, data sharing, access to data and archives, and finally data processing. CANFAR is a community cloud for Astronomy build on top of Compute Canada resources and based on a set of community services.

3.1 Scope of cloud federation

Scope of the cloud federation is to:

- Extend the portfolio of EGI federated cloud capabilities, through integration of new services based on IVOA standards and customization of generic EGI services (in particular clouds) to A&A requirements.
- Provide a new innovative cloud infrastructure (in particular data cloud) built for European Astronomers and Astronomical data centres.
- Provide close collaboration of e-Science infrastructures between EU and Canada.

Traditionally A&A has been at the forefront of implementing digital repositories, e.g. for sky observations with ground and/or space based telescopes. Typically such repositories are maintained within data centres with appropriate provision for tools/services for access and analysis. Data archives are expanding rapidly, e.g. through flagship high data volume generating A&A projects, and EuroVO has recently identified over 70 EU data centres. A&A data centres are employing Virtual Observatory (VOBs) to provide seamless unified access to distributed and highly heterogeneous data archives. The use of IaaS cloud computing facilities is thus becoming increasingly important. E.g. EuroVO emphasizes that relation to the VOBs should be taken into appropriate consideration in order to provide a complete data usage ecosystem for A&A communities.

3.2 The EGI CANFAR federation

The federation activities officially started at month 6 of the Project (September 2015) however some preliminary discussions have been done to identify the various aspects of the federation activity and roadmap (D4.1³).

² International Virtual Observatory Alliance <http://ivoa.net>

³ Deliverable D4.1 <https://documents.egi.eu/document/2549>

CANFAR is a cloud processing and cloud storage infrastructure that integrates authentication and authorization, monitoring, virtual storage environment and computing capabilities. The Federation model we propose in this document is based on two use cases:

- Authentication and Authorization Infrastructure (AAI) federation. CANFAR is now offering resources also to European users and groups. Users are registered in the CANFAR AAI service and they are issued CANFAR credentials. An interoperable AAI will allow European users to access CANFAR resources (data and computing) using their European (EGI Federated Cloud) credential. At the same time Canadian astronomers will be able to access EGI Federated Cloud resources using their Canadian credential. This will encourage collaboration between Canadian and European astronomers, projects and data centres. By implementing the AAI federation, users will exploit transparently resources provided via the integration with EGI Federated Cloud (EGI Federated Cloud).
- Data federation. CANFAR is offering virtual storage based on IVOA standards (VOSpace), which is used by Astronomers and data centres (CADDC) to store and share data. We would allow data access and sharing from A&A community and offer new capabilities to European data centres to share open data to astronomers and citizens using EGI Federated Cloud. It will allow data sharing (e.g. to replicate open and private data for data availability and preservation) and Virtual Machines (VMs) sharing between Canada and Europe. Finally, it allows A&A community to move computation close to data rather than moving large amount of data.

The final goal of this activity is to provide interoperable access to storage resources for both European and Canadian users.

Moreover, starting from project Month 6, we began to analyse and eventually federate the core infrastructure services offered by EGI and CANFAR: Authentication and authorization infrastructure, Accounting, Monitoring services, Operation tools and services, Service registry and market place.

A work plan has been defined and described in the federation Roadmap (D4.1): the main focus of the first six months of development activity (M6-M12) regards the Authentication and Authorization Infrastructure and a Storage service that allows astronomers to exchange data using IVOA standards (VOSpace⁴).

We analysed the AAI of the EGI Federated Cloud and CANFAR. Both clouds use X.509 certificate to identify users and to delegate user credential when using services. CANFAR however uses username/password authentication and users are registered to a Group Management Service (GMS) that is in charge of authorization procedures. GMS is queried by CANFAR services to verify user capabilities and capacities. The focus of the initial development is to implement an interoperable GMS service on top of the EGI Federated Cloud and to achieve group membership resolution from CANFAR to EGI GMS. In this way any European Astronomer will be able to access

⁴ <http://www.ivoa.net/documents/VOSpace/>

both CANFAR resource and the EGI resources dedicated to Astronomy. We also begin to study the CANFAR distributed storage system.

4 gCube/D4Science

D4Science is an infrastructure powered by the gCube⁵ system offering a number of services - currently integrating more than 500 software components - and Virtual Research Environments⁶ (VREs) for seamless access and analysis to a wide spectrum of data including biological and ecological data, geospatial data, statistical data and semi-structured data from multiple authoritative data providers and information systems.

D4Science is a Hybrid Data Infrastructure connecting >2000 scientists in 44 countries; integrating >50 heterogeneous data providers; executing >13,000 models & algorithms/month; providing access to over a billion quality records in repositories worldwide, with 99.7% service availability. D4Science hosts >40 Virtual Research Environments to serve the biological, ecological, environmental, mining and statistical communities worldwide. D4Science relies on a physical infrastructure counting more than 100 hosting nodes provided by the parties operating the infrastructure (namely CNR⁷, UoA⁸, FAO⁹, CITE¹⁰ and ENG^{11,12}).

D4Science Infrastructure usage - and thus resources needed - is far from being constant along time and spikes in the demand for computing capacity are often observed in the typical usage of the infrastructure. Cloud bursting appears as the natural deployment model to face this scenario. D4Science is implementing such model by relying on either research infrastructures (e.g. EGI Federated Cloud) or commercial public clouds (e.g. Amazon EC2).

In the context of the EGI-Engage project, the activity was focused on a more stringent identification of D4Science specific use cases whose implementation on top of EGI Federated Cloud would maximize the impact on the management and usage of the D4Science infrastructure. The outcome of discussion between EGI.eu and CNR teams, was the identification of a set of use cases of interest, each one with a different focus: integration at IaaS level and bundling of D4Science-specific Virtual Appliances; elastic execution of D4Science processes on the remote Federated Cloud infrastructure by exploiting existing Virtual Appliances; identity and authorization federation.

As a consequence of the agreed set of priorities defined above, and considering the very tight time frame for this activity, which ends at project-month 12, preference was given to the achievement of quick scalability for a typical D4Science process such as the execution of data analysis (e.g.

⁵ The gCube Framework - <https://www.gcube-system.org>

⁶ L.Candela, D.Castelli, P.Pagano, Virtual Research Environments: an overview and a research agenda, Data Science Journal, Volume 12, 10 August 2013

⁷ Consiglio Nazionale delle Ricerche - <http://www.isti.cnr.it>

⁸ National and Kapodistrian University of Athens - <http://en.uoa.gr>

⁹ Food and Agriculture Organization of the United Nations - <http://www.fao.org>

¹⁰ Communication & Information Technologies Experts - <http://www.cite.gr>

¹¹ Engineering Ingegneria Informatica S.p.A. - <http://www.eng.it>

¹² Currently, ENG only hosts part of the testing infrastructure

signal forecasting) against one or more input datasets. The analysis is executed, when possible, in parallel on a set of gCube nodes equipped with an execution engine called SmartExecutor. As a result of the integration, the pool of machines running the SmartExecutor would be elastically managed on top of EGI Federated Cloud infrastructure.

A number of sub-activities were thus planned and performed during the reporting period. They are shortly reported hereafter, detailed information regarding this activity is available on D4.5¹³.

4.1 Evaluation of jOCCI

After registration to the fedcloud.egi.eu testbed¹⁴ Virtual Organisation (VO), the development team explored access to the various sites participating to the VO and got familiar with the Federated Cloud infrastructure. In particular, in the context of the experimentation of the jOCCI¹⁵ library, interactions with jOCCI developers and VO managers were held, related either to library usage, bug reporting and operational status of sites. For this activity, already-available virtual appliances have been used.

4.2 Virtual Appliances Creation

In parallel, two specific Virtual Appliances (VA) were created, including the gCube software stack needed to run the SmartExecutor and the DataMiner service. The VAs are based on Ubuntu 12.04 running the relevant software from gCube 3.9.0¹⁶. Contextualisation scripts have been designed to properly configure the VMs at boot time and to provide them with the needed security tokens and infrastructure endpoints to join the D4Science infrastructure. The VAs were successfully registered to EGI AppDB¹⁷ and endorsed by the fedcloud.egi.eu VO. Already at this stage, the D4Science development infrastructure was able to benefit from manually-started VMs (based on the above VA) running on EGI Federated Cloud.

4.3 Prototype design and implementation

The design of the overall architecture and components for the integration took place, including the definition of an abstract data-model, service APIs, deployment schema and user interaction scenarios. A prototype implementation of the components was done; at the time of writing, they're under integration and are planned to be included in the next gCube release scheduled by late February 2016, and deployed in the production infrastructure.

The Federated Cloud accounting facilities were explored with the goal of integrating Federated Cloud accounting services into D4Science. As a matter of fact, relying on an authoritative source of

¹³ <https://documents.egi.eu/document/2662>

¹⁴ <http://operations-portal.egi.eu/vo/view/voname/fedcloud.egi.eu>

¹⁵ <https://github.com/EGI-FCTF/jOCCI-api>

¹⁶ <https://www.gcube-system.org/software-releases>

¹⁷ <https://appdb.egi.eu>

usage data would avoid D4Science from producing duplicate and potentially incoherent accounting. Requirements and use cases were shared with the EGI accounting team. Although there's no short-term plan to provide support for this use case, EGI will evaluate it, also considering similar requirements coming from INDIGO DataCloud.

With the aim of supporting elasticity in terms of automatic provisioning and decommissioning of cloud resources (scale in/out) based on some combination of monitored parameters, either on EGI Federated Cloud or possibly on other infrastructures, the Occopus framework¹⁸ was considered for integration. The D4Science and Occopus teams shared initial requirements and features respectively to enable a potential adoption; an initial exploration of the solution was done. Occopus will be explored further in the coming period with the aim of exploiting it in the D4Science infrastructure.

With the perspective of running D4Science production-level use cases on top of EGI Federated Cloud resources and services, a new VO was registered and made available on the "vomsmnia.cnaf.infn.it" VOMS server since December 2015¹⁹.

¹⁸ <http://occopus.lpds.sztaki.hu>

¹⁹ https://ggus.eu/index.php?mode=ticket_info&ticket_id=117484

5 Accelerated computing

Accelerated computing systems deliver energy efficient and powerful HPC capabilities. Many EGI sites are providing accelerated computing technologies to enable high performance processing such as GPGPUs or MIC co-processors. Currently these accelerated capabilities are not directly supported by the EGI platforms. To use the co-processors capabilities available at resource centre level, users must directly interact with the local provider to get information about the type of resources and software libraries available and which submission queues must be used to submit tasks of accelerated computing.

This task has implemented the support in the information system, to expose the correct information about the accelerated computing technologies available – both software and hardware – at site level, developing a common extension of the information system structure, based on the OGF GLUE standard, in order to have the capabilities published uniformly by all the sites. Users will then be able to extract all the information directly from the information system without interacting with the sites, and easily use resources provided by multiple sites. The task has also extended the HTC and cloud middleware support for co-processors, where needed, in order to provide a transparent and uniform way to allocate these resources together with CPU cores efficiently to the users. In the following sections the detailed activities carried out to implement accelerated computing support respectively on HTC and cloud platforms are reported.

5.1 Accelerated computing support for HTC platform

The activity on this field has not started from scratch. The previous EGI-InSPIRE project had in fact already faced the issue raised by many user communities to integrate in the EGI Grid Infrastructure their resource centres hosting accelerators cards, mainly GPGPUs. GPGPU (General-Purpose computation on Graphics Processing Units) is the use of a GPU (graphics processing unit) as a co-processor to accelerate CPUs for general scientific and engineering computing purpose. The GPU accelerates applications running on the CPU by offloading some of the compute-intensive and time-consuming portions of the code. The rest of the application still runs on the CPU. From a user's perspective, the application runs faster because it is using the massively parallel processing power of the GPU to boost performance. Experiments with the use of GPGPUs for scientific computing were already on-going within various NGIs, existing and potential EGI user communities. A GPGPU Virtual Team project was therefore set up in 2012 with the goal of collecting detailed requirements from existing and new EGI user communities and their support teams about using GPGPU services in the European Grid Infrastructure. Following this effort, a more structured EGI Technical GPGPU Working Group was set up in 2013 with the further goal of studying batch system integration best practices, and proposing and evaluating an appropriate GLUE schema suitable for describing generalised "Computational Accelerators". The results of

these studies were collected in a knowledge base and, for what concerns the information system aspect, embedded in the GLUE2.1 draft version.

In May 2015 the Accelerated Federated Computing session at the EGI Conference was organised with the goal of reviewing the old and new use cases proposed by the EGI-Engage user communities or Competence Centres (MoBrain, LifeWatch, Virgo, LHCb and MoDynGrid were represented). The benefit of using especially NVIDIA GPGPU cards was highlighted for many popular applications, together with the interest in accessing this kind of resources through both HTC and Cloud platforms.

This section will describe the activity carried out to enable GPGPU support in the EGI HTC platform, while the Cloud platform related activities will be reported in the next section.

A work plan was initially defined to develop a solution enabling GPGPU support in CREAM-CE for the most popular Local Resource Management Systems (LRMS) already supported by CREAM-CE and, for what concerns the information system, based on the GLUE 2.1 schema. It consisted in the following steps:

1. Identifying the relevant GPGPU-related parameters supported by the different LRMS, and abstract them to significant JDL attributes
2. GPGPU accounting is expected to be provided by LRMS log files, as done for CPU accounting, and then follows the same APEL flow
3. Implementing the needed changes in CREAM-core and BLAH components
4. Writing the infoproviders according to GLUE 2.1
5. Testing and certification of the prototype
6. Releasing a CREAM-CE update with full GPGPU support

At the same time, a test-bed with 3 nodes (2x Intel Xeon E5-2620v2) with 2 NVIDIA Tesla K20m GPUs per node was set up and made available at CIRMMP data centre. The test-bed was managed by the Torque 4.2.10 LRMS (source compiled with NVIDIA NVML libraries) with the Maui 3.3.1 scheduler, and was installed with CUDA 5.5 and two applications (AMBER and GROMACS) used by the MoBrain Competence Centre. The last version of the EMI3 CREAM-CE was finally installed on top of the LRMS to enable remote grid access to enmr.eu VO members.

Initial tests were carried out with local AMBER job submission using Torque with pbs_sched as scheduler (i.e. not using Maui). It allowed to specify, other than the number of requested GPU cards, the NVIDIA compute mode, that can be: default (shared mode available for multiple processes); exclusive_thread (only one compute thread is allowed to run on the GPU); prohibited (no compute contexts are allowed to run on the GPU); exclusive_process (only one compute process is allowed to run on the GPU). A first prototype of GPGPU-enabled CREAM-CE was therefore developed by defining two new JDL attributes: GPUNumber and GPUMode. It required it required the adaptation of the batch system adaptors (BLAH) and in CREAM-CE core. The remote job submission to this prototype through the glite-ce-submit client was successfully tested with the AMBER application in July 2015. After that, tests were made replacing the default scheduler in

Torque with the more popular Maui scheduler, but turned out that NVIDIA compute mode could not be set anymore with Maui. In parallel, an analysis of other popular LRMSes like LSF, Slurm, SGE and Condor showed that also for them the NVIDIA compute mode could not be set at job level, so it was decided not to implement the GPUMode JDL attribute in the final GPGPU-enabled CREAM-CE release. In November 2015 the same CREAM-CE prototype deployed at CIRMMP testbed was used to support the MoBrain Competence Centre application DisVis, encapsulated in a Docker container able to exploit the GPU cards capabilities.

Concerning the information system, the current draft of the OGF GLUE2.1 schema already containing GPGPU related objects and attributes was analysed and found that it could be easily supported by the CREAM-CE infoprovider mechanism. A discussion to clarify if other GPGPU related attributes should be included is on-going with the OGF GLUE coordinator.

Concerning the accounting, the APEL team was immediately involved for investigating how to address GPGPU accounting for Torque and other LRMSes. CREAM-CE accounting sensors in fact mainly rely on the LRMS log files. Unfortunately, there is no track of GPU usage in the log files of the LRMSes considered (Torque, LSF, SGE, Slurm and Condor). In fact, despite NVIDIA libraries allow to enable per-process accounting of GPU usage using Linux PID, there is no LRMS integration yet. The accounting issue was widely discussed at the EGI Community Forum held in Bari in November 2015. The outcome was that developing by our own the NVIDIA per-process accounting integration with one or more LRMSes is not an option that we can support and maintain in the long term. A common solution assuming that the use of GPU is exclusive and covering both grid and cloud scenarios is currently under study by the APEL team.

5.2 Accelerated computing for Cloud platform

The activity started with a review of the available technologies, focusing on GPGPU virtualisation in KVM/QEMU hypervisor. A testbed with an IBM dx360 M4 server with two NVIDIA Tesla K20 accelerators, Ubuntu 14.04.2 LTS with KVM/QEMU, and PCI passthrough virtualization of GPU cards was set up at IISAS laboratory. The CUDA version of NAMD molecular dynamics simulation was used for testing PCI passthrough virtualisation performances with respect to bare metal. It turned out that testing application run 2-3% slower in virtual machine compared to direct run on bare metal server. To avoid potential performance problems, hyperthreading had to be switched off.

After this initial phase, in July 2015 a cloud site with the OpenStack Compute, version 2015.1 (Kilo), Cloud Management Framework was deployed, with a controller node and two compute nodes hosting two NVIDIA Tesla K20 GPU cards with the configuration reported above. A VM image based on Ubuntu 14.04 and with GPU driver and libraries pre-installed was also created. Performance testing were then repeated in this OpenStack setup. Several discrepancies were found comparing performance of cloud-based VM with non-cloud virtualization and physical machine, that required not trivial tuning and optimisation, e.g. operating on setting CPU flavour in OpenStack Compute component, and adjusting the OpenStack scheduler. At the same time the

process of the integration of the site into the EGI Federated Cloud started, and was completed in October 2015. Moldyngrid, enmr.eu and vo.lifewatch.eu VOs were initially supported. In parallel, a new authentication module for simplifying the logging into the Horizon dashboard from EGI Federated Cloud users allowing them to use their keystone token was developed, together with various client tools for getting token, and installing NVIDIA drivers and CUDA. Two GPU-enabled flavours were made available: gpu1cpu6 (1GPU + 6 CPU cores) and gpu2cpu12 (2GPU +12 CPU cores).

A public live demonstration of the use of virtualised GPGPU servers on the EGI Federated Cloud with a real molecular dynamics application provided by the Moldyngrid community was successfully given at the EGI Community Forum held in Bari the second week of November 2015. Tutorial documentation for both users and site administrators was also made available on the project's wiki²⁰, guiding the user on how to use GPGPU on IISAS-GPUCloud site, how to create your own GPGPU server in cloud, and the cloud provider administrator on how to enable GPGPU PCI passthrough in OpenStack.

²⁰ <https://wiki.egi.eu/wiki/GPGPU-FedCloud>

6 EGI Cloud Federation

The EGI Federated Cloud is a multi-national cloud system that integrates institutional clouds into a scalable computing platform for data and/or compute driven applications and services. The architecture of the EGI Federated Cloud was defined in 2011-2012 and this was fully implemented by May 2014. Since then, the EGI Federated Cloud is a federation of autonomously operated, heterogeneous 'Infrastructure as a Service' (IaaS) type clouds, with all the participating cloud providers implementing and exposing the same set of interfaces towards cloud users and cloud administrators. These common interfaces allow providers to connect their sites to EGI central services and provide to users a common standard interfaces (OCCI) to manage and instantiate Virtual Machine images on the participating clouds in a homogeneous way. Implementations of these common interfaces are currently available for the OpenStack, OpenNebula and Synnefo Cloud Management Frameworks (CMFs). Connectors for other CMFs can be implemented as needed.

This initial concept of the EGI Federated Cloud is changing its scope and architecture and becomes a collaboration that enables various types of cloud federations to serve diverse demands of researchers from both academia and industry. The new EGI Federated Cloud brings together scientific communities, R&D projects, technology and resource providers to form a community that integrates and maintains a flexible solution portfolio that enables various types of cloud federations with IaaS, PaaS and SaaS capabilities. The collaboration is committed to the use of open source tools and services that are reusable across scientific disciplines. These tools and services form a flexible portfolio from which a scientific community can mix and match items to establish its own, customised cloud federation.

6.1 Federating Clouds

The EGI Federated Cloud provides the services and technologies to create federation of clouds (community, private or public clouds) that operate according to the preferences, choices and constraints set by its members and users. The EGI Cloud Federations are modelled around the concept of an abstract Cloud Management stack subsystem that is integrated with components of the EGI Core Infrastructure and that provides a set of agreed uniform interfaces within the community it provides services to.

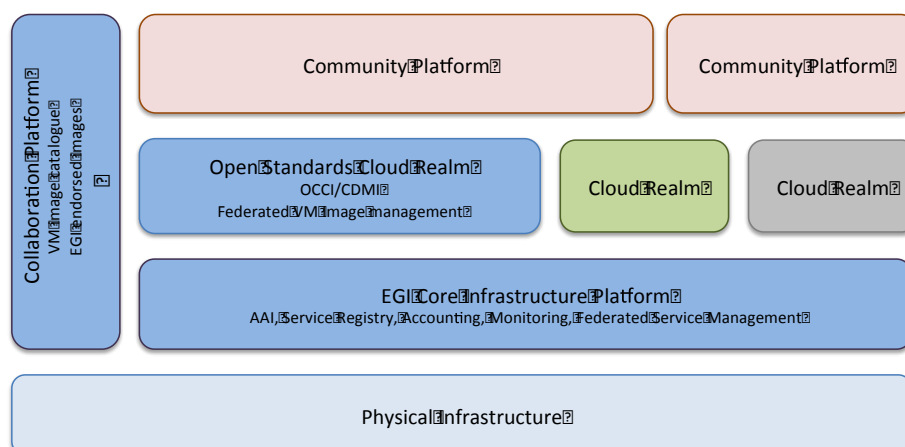


Figure 1. Federated Cloud realms

The EGI Cloud Federation (see Figure 1) is a hybrid cloud composed by public, community and private clouds, all supported by the EGI Core Infrastructure Platform services. The EGI Federated Cloud is composed by multiple “realms”, each realm having homogeneous cloud management interfaces and capabilities. A cloud realm is a subset of cloud providers exposing homogeneous cloud management interfaces and capabilities. The Open Standards Cloud Realm supports the usage of open standards for its interfaces and is completely integrated with the EGI Core Infrastructure Platform. A Community Platform provides community-specific data, tools and applications, which can be supported by one or more realms.

Despite the large diversity in the type of cloud realms, a relatively small number of identical building blocks (or federator services) can be identified in almost all of them. These services turn individual clouds into a federation. The table collects these common services to help architects identify topics they should focus on when designing a cloud federation. A more detailed description is available in the EGI wiki²¹.

Table 1. Services in a federated cloud.

Federator service	Role within the federation	Existing technical solution within EGI
Service Registry	A registry where all the federated sites and services are registered and state their capabilities. The registry provides the ‘big picture view’ about the federation for both human users and online services (such as service monitors).	GOODB
Information System	A database that provides real-time view about the actual capabilities and load of federation	BDII

²¹ https://wiki.egi.eu/wiki/Federated_Cloud_Technology

	participants. Can be used by both human users and online services.	
Virtual Machine Image Catalogue	A catalogue of Virtual Machine Images (VMIs) that encapsulate those software configurations that is useful and relevant for the given community (typically pre-configured scientific models and algorithms).	AppDB
Image replication mechanism	A system that automatically replicates VMIs from the federation VMI catalogue to each of the member sites, as well as removes them when needed. Automated replication can ensure consistency of capabilities across sites and is very often coupled with a VMI vetting process to ensure that only properly working, and relevant VMIs are replicated to the cloud sites of the community.	vmcatcher/vmcaster
Single sign-on for users	Ensuring that users of the federation need to register for access only once before they can use the federated services. Single sign-on is increasingly implemented in the form of identity federations in both industry and academia.	IGTF X509 proxies with VOMS extensions
Integrated view about resource/service usage	A system that pulls together usage (accounting) information from the federated sites and services, integrates the data and presents them in such a way that both individual users and communities can monitor their own resource/service usage across the whole federation.	Cloud Usage Record, APEL Accounting repository and portal
Integrated interfaces or user environments	Having interfaces through which users and user applications can interact with the services offered by the various cloud providers. In case of an IaaS cloud federation these interfaces offer compute, storage and network management capabilities.	OCCI API and OpenStack API
Availability Monitoring	Use a shared system to monitor and collect availability and reliability statistics about the distributed cloud service providers and to retrieve this information programmatically.	ARGO monitoring system
Federated service management tools	A set of processes, policies, activities and supporting tools customized to the federated cloud.	EGI federated service management

6.1.1 EGI Federated Cloud realms

The new EGI Federated Cloud is expected to result into multiple cloud federations (community specific, private or public), with the currently existing EGI cloud federation becoming one of them (Open Standards Cloud Realm).

Based on the EGI federation services and custom external solutions, any scientific community can create a federated cloud. Each community or e-infrastructure that wants to build a cloud federation decides the services required to support their computational needs. Because these cloud federations are largely built from tools and services of the same solution portfolio, they can maintain the portfolio together; they can share best practices, and can offer user support and training in a collaborative fashion.

EGI currently operates two realms: the Open Standards Realm and the OpenStack Realm. Both are completely integrated with the EGI federator services described above but use different interfaces to offer the IaaS capabilities to the users: the Open Standards Realm uses OCCI standard²² (supported by providers with OpenNebula, OpenStack or Synnefo cloud management frameworks), while the OpenStack Realm uses OpenStack native Nova API²³ (support limited to OpenStack providers). This OpenStack Realm was introduced in the federation during November 2015 and most of the resource providers already in the Open Standards Realm using OpenStack have started to provide this API along with the existing OCCI interface.

Table 2. Current EGI Cloud Realms

Federator service	Open Standards Realm	OpenStack Realm
Service Registry	GOCDB	
Single sign-on	X.509 proxies with VOMS extensions	
Accounting	Cloud Usage Record	
Information discovery	BDII	
VM Image catalogue	AppDB	
VM Image distribution	vmcatcher/vmcaster	
IaaS interface	OCCI	OpenStack Compute API
Monitoring	ARGO (OCCI specific probes)	ARGO (OpenStack specific probes)

6.2 Integration of providers into the EGI Cloud

In order to track the different integration activities with the EGI Federated Cloud, a new queue in the EGI Request Tracking (RT) system was created. The queue, named “fedcloud-integration”, collects for each of the collaboration the contact points and its status following the same phases as the ones defined in the workflow of the Federated Cloud support use cases²⁴ that capture the

²² <http://occi-wg.org/about/specification/>

²³ <http://developer.openstack.org/api-ref-compute-v2.1.html>

²⁴ https://wiki.egi.eu/wiki/Federated_Cloud_Communities

process from the pre-assessment to complete integration in the production infrastructure. The selection of collaborations is mainly driven by the requirements of user communities involved in WP6. The following use cases and progress have been captured in this period:

- IHEP (Chinese Academy of Science): IHEP has several cloud deployments and is interested in the Federation of these within IHEP itself and with EGI. Initial contact has been established and IHEP will evaluate how the EGI model can fit their infrastructure
- NeCTAR: NeCTAR is an Australian-wide cloud federation based on OpenStack. Collaboration with EGI has started to integrate the resources following the requirements from key user communities such as the Human Brain Project (HBP), ELIXIR and SKA/LOFAR. Technical discussions on different aspects have taken place at several online meetings and during the Bari conference.
- CERN: CERN counts with a large OpenStack deployment and is interested in a loose federation profile with EGI using sustainable developments. Then CERN cloud is already integrated with the EGI accounting services. Native OpenStack support in EGI allows for further integration.
- KISTI: KISTI is a resource provider in Korea providing Cloud resources using OpenStack. Integration activity with EGI is stalled due to lack of OCCI support for their OpenStack networking configuration. Better network support in OpenStack is being developed by INDIGO project and expected to be released in the coming months.
- FogBow: The EU Brazil Cloud Connect project has developed a middleware named FogBow for the creation of cloud Federations. FogBow is partially integrated with EGI on providers that participate in EU Brazil Connect.
- Harness: Harness is a EU FP7 project providing support for non-conventional architectures on cloud. Harness plan included the evaluation of OCCI as an interface for making the resources of the project available for EGI. Harness project ended without producing any results on this area and the use case is now closed.
- Compute Canada: triggered by the CANFAR collaboration, Compute Canada has been identified as a possible e-Infrastructure to federate with EGI. The CANFAR developments will rely on this integration.
- IUCC: IUCC (Israel) is evaluating the EGI Federated cloud for federating their resources.
- GARR: The Italian NREN has expressed its interest on Cloud federation. The EGI cloud federation model was presented during a meeting and now is under internal discussion within GARR.
- BITP: Triggered by a use case in the federated cloud user support. BITP is a Ukrainian site that has successfully deployed EGI federated cloud tools and is expected to enter production status soon.

- EBI: As part of the creation of the ELIXIR cloud, EBI resources are currently being federated with EGI Federated Cloud technology. The site is progressing with the integration and has completed the deployment of the modules to support EGI's AAI, OCCl and accounting.

7 Future plans

This section gives an overview of the future plans regarding the different activities reported above. The overall activity will also look perform a systematic analysis of the integration models from each case that can help as a guidance for new cases and a way to discover more efficient solutions.

7.1 EUDAT

The next steps of EUDAT integration involve building on the universal use case pilot that was introduced in Section 2 of this document. Staff from both EGI and EUDAT is working with user communities to provide services from both infrastructures into their specific use cases. Two use cases have been prioritised to drive the pilot implementation and to act as early adopters: EPOS and ICOS. We will then investigate ways for users of either infrastructure to benefit from interoperability in the following areas:

- Support the new AAI infrastructures based on Identify Federation that both infrastructures are designing and implementing according to the AARC guidelines;
- Adopt new high-level EUDAT API to transfer datasets between the 2 infrastructures instead of the low-level globus command;
- Support the generation and management of PIDs;
- Interconnect the EGI Federated Cloud with the EUDAT B2FIND, to discover datasets location, and B2DROP services, to share any intermediate result with the public.
- Review and updating of policies to enable and foster infrastructure interoperation
- Pan infrastructure service discovery and monitoring
- Combined user support and bug reporting

7.2 CANFAR

The CANFAR integration will continue following the roadmap defined in D4.1. A monthly teleconference is dedicated to the presentation and discussion of the high level services and their possible federation (authentication and monitoring have been discussed) and to enforce the collaboration of the two infrastructures. A weekly teleconference to monitor development and CANFAR services integration progress activity is in place.

7.3 gCube/D4Science

A number of further enhancements are foreseen and are being considered for realisation during the lifetime of the EGI-Engage project, also in collaboration with the BlueBRIDGE project. In particular:

- the evolution of the EGI accounting subsystem will be monitored in order to integrate possible public APIs so to benefit from an authoritative source of usage data, rather than tracking usage on the client side, producing potentially incoherent data.
- automate the creation or Virtual Appliances upon release of relevant gCube components and registration to EGI AppDB via the AppDB REST API. This would speed up the timely availability of updated appliances on AppDB.
- further explore the opportunity of adopting Occopus to support of elasticity in terms automatic provisioning and decommissioning of cloud resources across external cloud infrastructures. Different parameters are expected to drive the elastic behaviour, including current resources load, established quotas, pricing models, performance and QoS indexes, etc.

7.4 Accelerated Computing

In January 2016 the CREAM-CE developer's team has started to prepare a new CREAM-CE release working with CentOS7 and SL6 operating systems that will be included in UMD4²⁵. The modifications to enable GPGPU support (both for job submission and information system aspects) at least for Torque LRMS will be embedded in this release and therefore fully certified. The certification process is planned to be completed before the end of the JRA2.4 task, scheduled in May 2016. The new CREAM-CE will then be made available through UMD4 to all data centres that have shown interest in deploying a GPGPU-enabled computing element. The solution for GPGPU accounting designed by the APEL team is also planned to be implemented by the end of May 2016.

In the next months until the end of the JRA2.4 task, the cloud related activity will focus on supporting the deployment of new cloud sites hosting GPGPU servers, as well as new applications suited to be executed on the GPGPU-enabled cloud sites. For example, CESNET/MU partner of MoBrain Competence Centre has shown interest in adding their GPGPU nodes to the cloud testbed and preparing a GPGPU-enabled VM image with GROMACS and AMBER applications for testing them in the EGI Federated Cloud environment. Moreover, the cooperation with the APEL team will continue in order to address the accounting of GPGPUs in cloud environment. Concerning the information system, the team will collaborate with the EGI Federated Cloud task force and with the OGF GLUE coordinator in order to implement the correct publication of the GPGPU related attributes in the EGI Federated Cloud operations context.

7.5 EGI Cloud Federation

The EGI Federated Cloud promotes the use of standards for interoperability of cloud providers using different technologies. However, legacy applications and advanced use cases using specific provider's features may not be easily ported to the standard APIs. In order to attract communities and new providers with existing investments on OpenStack, EGI introduced the support for the

²⁵ Unified Middleware Distribution (UMD): <https://wiki.egi.eu/wiki/UMD>

OpenStack native interface in November 2015. It will also ease the migration to the EGI cloud and can attract communities, which may already have legacy applications using the OpenStack environment. The wide ecosystem of tools for this platform will allow EGI to offer more advanced features and interfaces for the existing resources, some preliminary work on using the OpenStack web dashboard with EGI AAI has proved that existing tools can be migrated with reasonable effort. This native support will also speed up the integration of OpenStack based federations (e.g. NeCTAR, Compute Canada) and providers (e.g. CERN).

The further development of the Open Data and Federated Cloud in WP4.1 and WP4.2 will also make EGI a more attractive platform for communities and will likely boost the integration activities with other peer e-Infrastructures. One of the main barriers for integration is the AAI architecture, which is currently under a major update within EGI. This new AAI architecture and the support for it at the level of the Cloud Management Frameworks can also act as a stimulus for more integration activities in this area.

The EGI Federated Cloud will continue to refine its federation model to accommodate the requirements coming from the different use cases and the requirements for providing a technical integration with offered by other e-Infrastructures. The task will closely work with user communities involved in WP6 to gather feedback from user communities that bring additional guidance on how to bridge any identified gaps and will coordinate the implementation of pilots and will liaise with the external partners.