Deployment of a Federated Cloud Infrastructure in the Black Sea Region

Hrachya Astsatryan, Andranik Hayrapetyan, Wahi Narsesian
Institute for Informatics and Automation Problems of the National Academy of Sciences of Armenia, Yerevan, Armenia
e-mail: hrach@sci.am, {andranik, vner75}@ipia.sci.am

Ramaz Kvatadze, Nugzar Gaamtsemlidze
Georgian Research and Educational Networking Association, Tbilisi, Georgia
e-mail: {ramaz, ngamtsemlidze}@grena.ge

Peter Bogatencov, Nicolai Iliuha
Research and Educational Networking Association of Moldova, Chisinau, Moldova
e-mail: {bogatencov, nick}@renam.md

Vladimir Florian, Gabriel Neagu, Alexandru Stanciu
National Institute for Research and Development in Informatics, Bucharest, Romania
e-mail: {vladimir, gneagu, alex}@ici.ro

ABSTRACT
Over the last few years distributed computing infrastructures (DCI) - Grid, HPC, Cloud, etc. -, have become the backbone of those fields of science that require to solve complex computational problems. At the same time, the commercial application of such technologies has led to new categories of offerings (e.g. on-demand offerings) and will likely play an important role in the "software as a service" landscape. However, the complexity of these infrastructures is often discouraging new and inexperienced users and impedes the use of DCI technologies in new application domains.
The main aim of the article is to introduce a federated Cloud infrastructure for the South East Europe, including Romania, Armenia, Georgia and Moldova.

Keywords
Virtualization, Cloud management platform, federated Cloud infrastructure, on-demand Grid services, integrated Grid and Cloud infrastructure, software-as-a-service.

1. INTRODUCTION
In resent years, state-of-the-art research infrastructures have been deployed in the Black Sea region via a number of targeted initiatives funded by the European Commission. Various types of research infrastructures which are available in the region make possible to handle large data sets and to provide significant computational resources.
In the first stage an integrated Grid Infrastructure [1] has been deployed in the South East Europe [2], which later became a part of European Grid Infrastructure [3]. Grid computing architecture is a collaborative, network-based model that enables the sharing of computing resources: data, applications, storage and computing cycles. The computational resources of National Grid Infrastructures in Armenia, Georgia, Moldova and Romania are presented in the Table 1.

Table 1: Computational Resources

Country	Cores
Armenia	496
Georgia	40
Moldova	84
Romania	7710
Total	**8330**

In the second stage, the countries started to deploy other types of infrastructures, such as supercomputers, Clouds, etc. The article introduces a federated Cloud infrastructure which can provide different solutions for universities, scientific and research communities and more. The platform will make possible to solve the following problems:

- increasing the effective usage of computational resources;
- providing additional different services for scientific and research communities;
- close collaboration between different countries to solve common regional problems.

The objective is to create an infrastructure which uses resources provided by geographically distributed computing clusters. These sites are operated by independent organizations, which have total control on managing their own resources, including the setup and enforcement of special administrative policy regarding authorization and access, security, resource usage quota, monitoring and auditing of the local infrastructure. The resource providers must delegate the control for a part of their infrastructure in a safe and efficient way, so a federated infrastructure can be built based on the resources available on these distinct administrative domains. This is a major challenge regarding the implementation of a federated Cloud infrastructure as it is sought to be achieved by the BlackSea-Cloud [4].

2. Integrated Cloud Infrastructure

Recently there has been a great change from Grid computing to the Cloud computing [5]. The main aim of Cloud computing is to build systems in which users interact with applications remotely over the internet. This approach has many advantages for both the application providers and users. For the application providers this system will prevent them from buying and maintaining expensive hardware, and gives them the opportunity to use highly scalable resources in the Cloud to meet their needs. From the users perspective they will not need to install any software on their side, and they will access the services over the internet from any location in the world, also they will not concern about data loosing. Also, the pay-and-go approach for the Cloud will be a great opportunity for both software developers and users to pay only for an exact need, in other words, to pay only for the time needed to use or deploy or test the service.

The increasing demand for the Cloud computing and the requirements to build a reliable infrastructure encouraged the development of different Cloud computing middlewares which sustains the variety of applications by system level abstracts and common functionalities. Recently there is an increasing number of open-source solutions to build private, public and hybrid Clouds, such as OpenNebula [6], Eucalyptus [7], and OpenStack [8]. These packages help universities and institutions to build their own public/private Cloud and to offer different dedicated services and solutions for their research and scientific communities. The above mentioned open-source toolkits have been studied that enable the deployment of Infrastructure-as-a-Service Clouds, in order to identify the appropriate solution to support the integration of distributed Cloud resources that are independently managed by their providers.

The suggested platform has been deployed based on OpenNebula package. OpenNebula is the result of many years of research and development in efficient and scalable management of virtual machines on large-scale distributed infrastructures. Its innovative features have been developed to address the requirements of business use cases from leading companies in the context of flagship European projects in Cloud computing. OpenNebula is being used as an open platform for innovation in several international projects to research the challenges that arise in Cloud management, and also as production-ready tool in both academia and industry to manage Clouds.

Its open, architecture, interfaces and components provide the flexibility and extensibility that many enterprise IT shops need for internal Cloud adoption. These features also facilitate its integration with any product and service in the Cloud and virtualization ecosystem, and management tool in the data centre. OpenNebula provides an abstraction layer independently from underlying services for security, virtualization, networking and storage, avoiding vendor lock-in and enabling interoperability. OpenNebula is not only built on standards, but has also provided reference implementation of open community specifications, such as the OGF Open Cloud Computing Interface. This open and flexible approach for Cloud management ensures widest possible market and user acceptability, and simplifies adaptation to different environments.

The first level of the platform consists of resources provided by each partner Cloud infrastructures, such as:

- A Romanian Cloud infrastructure deployed at the National Institute for Research and Development in Informatics Bucharest, which is based on the KVM hypervisor, and uses Lustre as a distributed file system in order to provide a high performance storage for the virtual machines. The production Grid site RO-01-ICI is deployed using the virtual resource provided by this Cloud infrastructure.
- Moldavian experimental cloud infrastructure consists of one zone and deployed by using the physical resources of Research and Educational Networking Association of Moldova and Institute of Mathematics and Computer Science clusters (one physical server - 4 cores, 8 GB RAM, HDD 500 GB and one virtual machine - 2 cores, 2 GB RAM, HDD 160 GB).
- The Armenian experimental Cloud infrastructure which consists of two separate zones located in the Institute for Informatics and Automation Problems (IIAP) and in the Yerevan Physics Institute (YERPHI). Some specific images have been prepared for different scientific communities.
- The Georgian Research and Educational Networking Association is planning to deploy experimental Cloud infrastructure based on two servers.

It is suggested to combine or join together the national Cloud resources using the OpenNebula Zones (oZones) approach (see fig. 1) which allows centralized management of multiple instances of OpenNebula (zones), managing in turn potentially different administrative domains. This approach has a number of benefits such as:

- It offers the possibility to implement local policies regarding resource allocation and the partitioning of the Cloud infrastructure, in order to isolate the resources dedicated to the federated Cloud infrastructure from other resources which can be used at the provider's discretion;
- It provides scalability for the federated Cloud infrastructure, by including resources from each OpenNebula Zone, and providing a unified interface to access the images, virtual machines, networks available on each OpenNebula instance.

These zones can be effectively shared through the Virtual DataCenter (VDC) abstraction. A VDC is a set of virtual resources (images, VM templates, virtual networks and virtual machines) and users that manage those virtual resources, all sustained by infrastructure resources offered by OpenNebula. A VDC is supported by the resources of one zone, and it is associated with one cluster of the zone. The resources that the VDC can dispose of is a subset of that cluster. Users access the zone through a reverse proxy, so they don't need to know the endpoint of the zone, but rather the address of the oZones module and the VDC where they belong to.

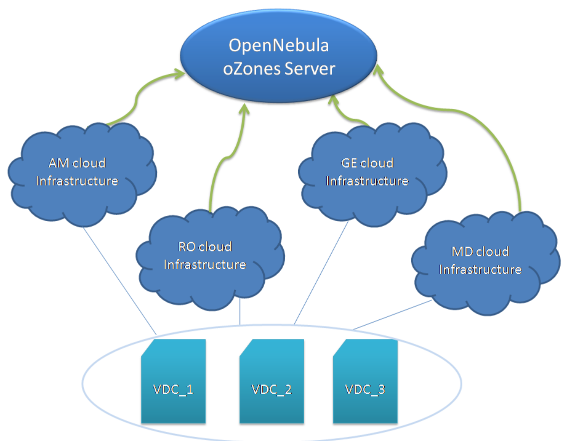

FIG. 1: OpenNebula oZones Hierarchy

After a successful combination of the resources the management of the virtual images may be carried out using a single interface regarding the virtual image's location.
The single point of access for the Cloud infrastructure must provide a web interface to enable users to conveniently manage virtual machines during their entire life-cycle, as well as to define and use other virtual resources such as storage and networks. In addition, the Cloud infrastructure should provide support on-demand provisioning of resources, such as the computing nodes and storage, which can be further contextualized to enable the creation of dedicated virtual appliances, complex frameworks, and platforms in order to support new application environments. The g-Eclipse framework [9] is planning to use to build different interfaces that allow users and developers to access arbitrary computing infrastructures (Grid or Cloud). Other programming APIs will be studied and implemented for the same purposes, by taking into consideration two approaches: one to have a desktop application and the second is to have a web interface.
The main approach is to combine the Grid and Cloud resources together as a single enhanced computational power, and offer the possibility to use Grid or Cloud resources on demand, as an example if the user requires a parallel computational resources he will submit a job on the Grid, but if the user needs any specific software or environment to solve some special problem he can use a dedicated Cloud service or virtual

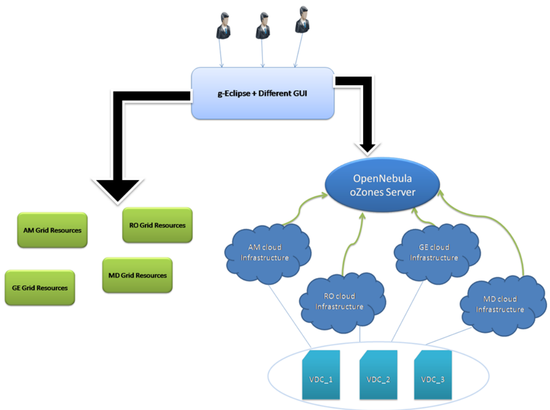

Fig 2. Skeleton of a regional platform.

In order to create an integrated Grid and Cloud infrastructures, we propose a two step process. The first phase in which Grid services are migrated from the physical infrastructure to the Cloud infrastructure deployed as specified above, and a second phase when Grid services are provided on-demand, being instantiated on the Cloud infrastructure when they are required and dynamically allocate resources for them according to the actual workload.
This feature will enable the creation of a dedicated Grid infrastructure that can be used for developing and testing new Grid application, as well as training and dissemination activities, without interfering with the production Grid services.
Regarding the infrastructure management it is planned to implement a centralized command and control system, that in some cases can be based on a configuration management system such as Puppet [10]. Because the federated Cloud infrastructure is based on several OpenNebula instances which are deployed by each resource provider, this means that each OpenNebula cluster would have to be managed by its own Puppet master. However, the principle of configuration management systems, that is to automate the system configuration processes, based on defined manifests or cookbooks, can be applied to each Puppet instance, and it can assure the proper management for each of OpenNebula Zones that provide resources for the federated Cloud infrastructure. This can have the same effect as a single Puppet instance deployed for the entire Cloud infrastructure.

3. CONCLUSION

The main and the most important outcome of this project is the optimal usage of these combined infrastructures and increasing the efficiency of the usage.
Another outcome of the project would be the simplification of the access to these resources for researchers and scientists, with dedicated solutions provided for them.
By sharing these resources, the partners of this project will benefit a lot by increasing the capability of their computational resources. It is also encouraged the sharing of the experience and knowledge between them.
After implementation of the idea mentioned above, there will be a possibility to provide a solution for various regional problems which connect these countries such as environmental problems, Big Data analysis, etc.

4. ACKNOWLEDGEMENT

This work was supported by the Organization of the Black Sea Economic Cooperation through the project "Experimental Deployment of an Integrated Grid and Cloud Enabled Environment in BSEC Countries on the Base of g-Eclipse".

REFERENCES

[1] I. Foster, C. Kesselman, Computational Grids Chapter 2 of "The Grid: Blueprint for a New Computing Infrastructure", Morgan-Kaufman, 1999.

[2] SEE-GRID eInfrastructure for regional eScience, http://www.see-grid-sci.eu.

[3] European Grid Infrastructure, http://www.egi.eu.

[4] Black Sea Cloud, http://www.blacksea-cloud.net/

[5] Katarina Stanoevska-Slabeva, Thomas Wozniak, Santi Ristol, Grid and Cloud Computing, Springer, 978-3-642-05192-0, 2010.

[6] OpenNebula Project Leads. http://opennebula.org

[7] Eucalyptus, http://www.eucalyptus.com/eucalyptus-cloud/iaas

[8] OpenStack, http://www.openstack.org/
M. Stümpert, H. Kornmayer, M. Knauer, g-Eclipse – A Middleware-Independent Framework for Accessing Existing Grid Infrastructures, Remote Instrumentation and Virtual Laboratories: Service Architecture and Networking, pp 277-292, Springer 2010, 10.1007/978-1-4419-5597-5_23

[9] James Loope, Managing Infrastructure with Puppet, O'Reilly Media Inc, 2011, p. 37.