Geoprocessing in Hybrid Clouds

Theodor Foerster, Bastian Baranski, Bastian Schäffer & Kristof Lange

Institute for Geoinformatics, University of Münster, Germany

{theodor.foerster; bastian.baranski;schaeffer; kristof.lange}@uni-muenster.de

Abstract. Meeting specific Quality of Service parameters in distributed architectures is one of the key requirements to build an operational infrastructure. This applies especially to SDIs, which offer geoprocessing functionality. This paper describes Hybrid Clouds as a means to meet these requirements in a efficient way by scaling the processing base load on internal (Private Cloud) and peak loads to external (Public Cloud) Cloud Computing infrastructure. The paper describes an architecture for Hybrid Clouds and a scenario performing image processing at a data center by the means of a Hybrid Cloud.

1. INTRODUCTION

Geoprocessing on the web is considered to be the next evolution step in Spatial Data Infrastructures (SDIs) (Kiehle et al., 2006). Increasing network bandwidth and processing capacities along with common communication protocols for Web Services such as established by ISO, OASIS and OGC have paved the way (Kradilis, 2007). On this basis, comprehensive architectures for distributed geoprocessing have been demonstrated (Schaeffer et al., 2009). To meet specific Quality of Service (QoS) parameters in such an architecture, Grid Computing has been identified as an applicable mean (Baranski, 2008). The utilization of the relatively new Cloud Computing paradigm seems to be the next logical step in this direction.

A Cloud infrastructure can be hosted either by the organization itself (Private Cloud) or by an external provider (Public Cloud) (Armbrust et al., 2009). A Hybrid cloud approach for Geoprocessing has not been proposed yet, but is considered to be beneficial to optimally use internal and external cloud infrastructure.

In this paper we will describe an architecture of a Hybrid Cloud for geoprocessing based on the Web Processing Service (WPS) interface specification (OGC, 2007). Furthermore we will present a prototypical implementation of the Hybrid Cloud architecture. The implementation is based on the 52°North Web Processing Service implementation (52N WPS) that utilizes the OpenNebula software package. OpenNebula is a tool for combining your local virtualized infrastructure (Private Cloud) with 3-rd party infrastructures (Public Cloud) to enable highly scalable hosting environments. The implementation is demonstrated for a data center, which needs to process incoming satellite imagery on-demand, that results in a use case in which critical processing steps during the analysis of satellite data are dynamically outsourced to the Amazon EC2 service platform. Finally, we will summarize our experiences in implementing the Hybrid Cloud approach and evaluate the existing gaps and bottlenecks in utilizing Cloud Computing in the geospatial domain.

The remainder of this paper is structured as follows. Section 2 will describe the related concepts in the field of Geoprocessing and Cloud Computing. In Section 3 we will present the architecture of our proof-of-concept of the Hybrid Cloud approach. This architecture is demonstrated in Section 4 by a scenario for image processing at a data center. These data centers need to be operational also during peaks of processing load. Finally, in Section 5 the findings are discussed and an outlook for future work items is presented.

2. RELATED WORK

This section presents the related concepts as applied in this research. First Cloud Computing is introduced. To demonstrate the benefits of Cloud Computing and Hybrid Clouds in particular the problems of distributed Geoprocessing are reviewed.

2.1 Cloud Computing

Cloud Computing is one of the latest trends in the mainstream IT world (Gartner, 2008). Several companies such as Amazon, Google, Microsoft and Salesforce have already carried out significant effort in that direction e.g. by building up their own public accessible cloud infrastructure. The term Cloud Computing indicates a future in which applications, storage and computations are no longer located on single computers, but on distributed facilities operated by third-party storage and computational utilities (Foster, 2008).

Cloud computing allows providers to scale their applications according to specific rules by providing computational power and storage dynamically in a cost efficient and secure way over the web. Thus, a client user is able to access the resources without having to manage the underlying complexity of the technology. In simple application scenarios, the client user does not need to know, that he/she is interacting with a cloud. Thereby cloud computing provides the following characteristics:

- Efficiency
- Outtasking
- Scalability.

Defining specific QoS parameters can be done in advance or ad-hoc. Specific rules can define for instance to allocate additional resources. In principle, there are two different ways of deploying an application on a cloud by using internal resource (private cloud) or by using external resource (public cloud).

A revision of cloud computing with regard to current SDIs is presented in Schaeffer et al. (2010). The authors analyzed the concepts of SDIs in respect to common gaps which are addressed by Cloud Computing.

2.2 Distributed Geoprocessing

Distributed Geoprocessing describes the execution of any geoprocess model, which is distributed on the web. The geoprocess model can run on a single Geoprocessing Service or be composed of a chain of Geoprocessing Services. In the context of the OGC, Geoprocessing Services are realized using WPS interface (OGC, 2007). Following the Model as a Service paradigm (Dumitru et al., 2009) by exposing complex geoprocess models on the web via standardized interfaces implies two general problems in production environments:

- the complexity of geoprocess models requires large computing capacities
- the number of user cause high computing loads even for simple geoprocess models.

To fulfill these needs, a scalable architecture is required. Cloud Computing identified in this research as a suitable architecture, which is scalable but also cost-efficient.

3 ARCHITECTURE FOR HYBRID CLOUDS

From an organizational perspective using a Private Cloud is applicable to scale the process effort on the organization's own infrastructure. This however might not always meet the required QoS parameters due to a lack of sufficient computational resources in the organization's own infrastructure. Therefore, to meet specific QoS parameters in peak times (e.g. the 99% availability of a web service as required by INSPIRE) and to generally optimize operational costs (e.g. the hardware utilization rate), organizations need to be able to dynamically outsource (parts of) their business processes and operational functionality to Public Clouds - the so-called Hybrid Cloud. The architecture for Hybrid Clouds is presented in Figure 1. The Hybrid Cloud entry point manages the underlying Private and Public Cloud and allocates additional resource if required by using first the private cloud and if additional resources are required, which cannot be provided by the private cloud, the public cloud (e.g. Amazon EC2) is accessed.

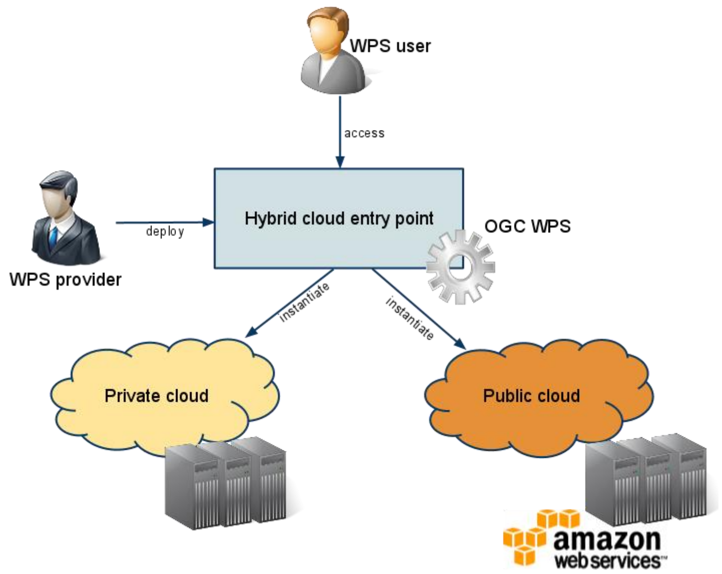

Figure 1: Architecture for a Hybrid Cloud.

The architecture identifies two different users of the Hybrid Cloud. The service provider (in this case providing WPS) configures the Hybrid Cloud with several rules to meet certain QoS parameters. The service user (accessing a WPS) uses the WPS service without noticing the cloud behind it.

For the given scenario the Cloud Computing characteristics (Section 2.1) are used to tackle the problems of distributed Geoprocessing (Section 2.2). This scenario is described in Section 4.

4 APPLYING THE HYBRID CLOUD FOR OGC WPS

Hybrid clouds are beneficial for companies, which host their own infrastructure for daily work such as work stations and several servers for providing customer services. In the given scenario, the company deploys a

Hybrid Cloud for cost reduction and for improved efficiency (i.e. cost reduction but still ensuring a specific level of QoS). The Hybrid Cloud is configured to ensure a specific level of QoS for performing on-demand analysis processes on up-to-date imagery. The on-demand analysis results in peaks of processing load in an irregular fashion. The Hybrid Cloud accesses first the Private Cloud and the Public Cloud if the process load exceeds the computational capacities of the Private Cloud (Figure 1).

For implementing the Hybrid Cloud, OpenNebula (www.opennebula.org) is used. OpenNebula is an open source tool for establishing any type of cloud in existing data center environments. It can be primarily used as a virtualization tool to manage a virtual infrastructure in the datacenter or cluster, for either installing a Private Cloud or Public Cloud. OpenNebula supports Hybrid Clouds by combine internal infrastructure with external cloud-based infrastructure to realize highly scalable applications (Figure 2). In particular, it deploys a mini-cluster in a private network and allows external nodes of for instance Amazon EC2 to join.

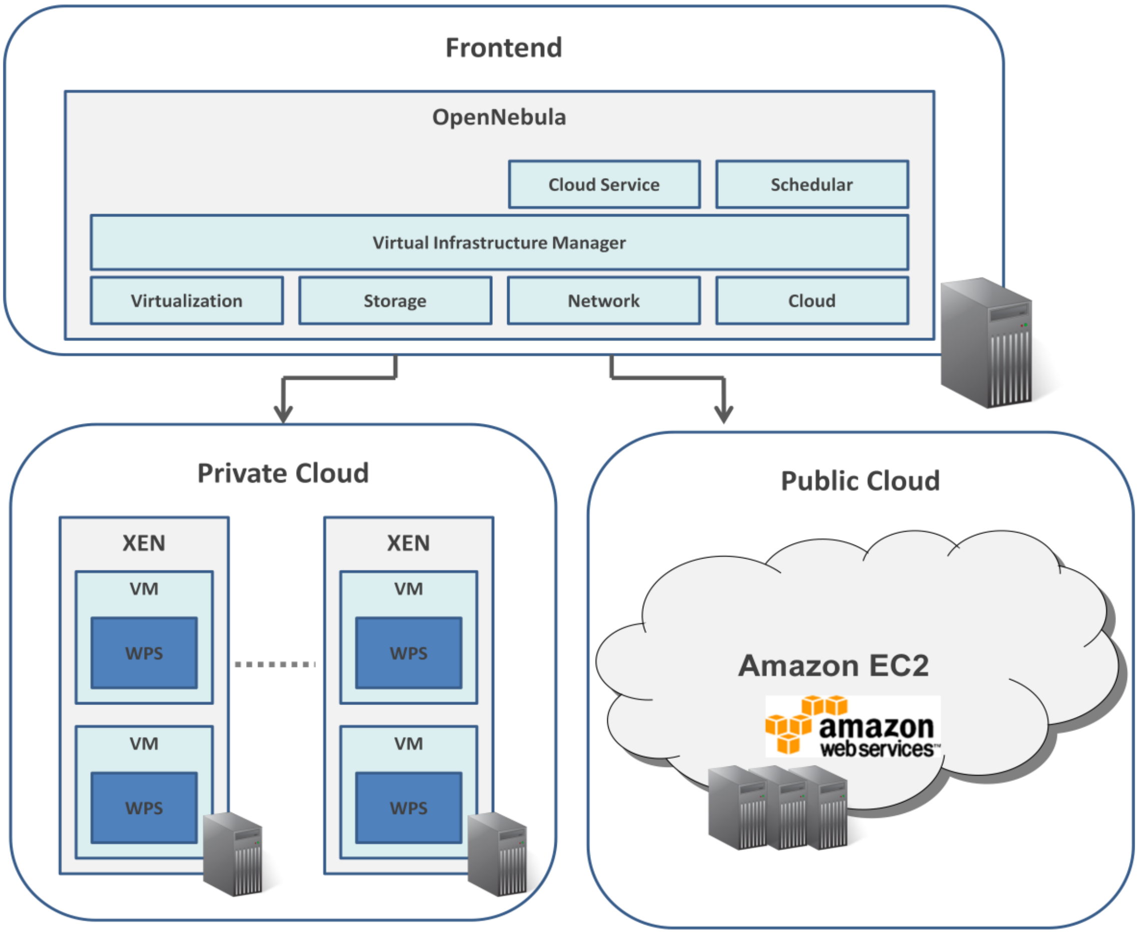

Figure 2: Implementation of a Hybrid Cloud.

OpenNebula is deployed within a local cluster of 4 servers (1 head node and 3 worker nodes). Open Nebula consists of a front end, which manages the local storage and interconnects the Virtual Machines (VMs) in the virtual network (a merger of the Private Cloud and the Public Cloud). Each worker node is configured with the XEN hypervisor, which enables the Virtual Infrastructure Manager at the frontend to start and stop any type of VM image at the worker nodes. Additionally, the Virtual Infrastructure Manager requires a VM image, which can be deployed on the worker nodes. In principle, one single VM image can be deployed on an unlimited amount of worker nodes at the same time. For accessing the Public Cloud (e.g. based on Amazon EC2), a cloud-provider specific image (such as the Amazon Machine Image (AMI)) is required at the frontend to deploy the desired functionality. Finally, a set of rules for the Scheduler at the Frontend has to be defined. This set of rules specifies, in which situation what type of VM should be deployed on how many nodes.

In the given scenario, the example VM images contain an installation of the WPS with the specific functionality for image analysis. Thus, the Hybrid Cloud can provide the desired functionality through WPS interface in a scalable fashion accessing internal and external infrastructure.

5 CONCLUSION

Cloud Computing allows service providers to scale applications in an on-demand way. Hybrid Clouds allow providers to scale their applications in a cost efficient way by using internal and external resources, if necessary. This paper presents Hybrid Clouds for scaling WPS-based image analysis processes. For the given scenario, OpenNebula is applied in combination with Amazon EC2.

Cloud Computing will have impact on business models of software vendors and service providers. Software vendors need to adjust their license models towards pay-per-use. Service providers will start to rent expensive hardware for specific tasks, to guarantee specific levels of QoS in an efficient way.

Hybrid Clouds are also important for organizations, when considering security of processes and data. In specific scenarios, processes and data will not be outsourced to external resources. Certificates are one solution to establish trust between the cloud consumer, the cloud provider and other 3rd parties. However these security issues are still subject to research.

Additionally, the standardization of Cloud Computing components is a major issue, that is already addressed e.g. by the so-called Open Cloud

Manifesto and different standardization organizations. Currently, as also shown in this research, for each cloud provider a specific image needs to be created. A standardized virtualization layer increases the interoperability of cloud computing solutions. Finally, a geo layer based on OGC standards might be integrated to ensure interoperability of geo services and data in all kind of cloud environments

In the future, in contrast to promising advertising slogans the real monetary benefits of Cloud Computing need to be investigated thoroughly to develop new business models as well as to improve existing business models. Additionally, the performance and the achieved level of QoS need to be measured to demonstrate the gained benefit of Cloud Computing.

6 REFERENCES

Armbrust, M., Fox, A., Griffith, R., Joseph, Anthony D., Katz, Randy H., Konwinski, A., Lee, G., Patterson, D. A., Rabkin, A., Stoica, I. & Zaharia, M. (2009). Above the Clouds: A Berkeley View of Cloud Computing. Technical Report. EECS Department, University of California, Berkeley.

Baranski, B. (2008). Grid Computing Enabled Web Processing Service. In E. Pebesma, M. Bishr, & T. Bartoschek (Eds.), *Proceedings of the 6th Geographic Information Days*, IfGI prints (Vol. 32, pp. 243-256). Presented at the GI-days 2008, Muenster, Germany: Institute for Geoinformatics. Available: http://www.gi-tage.de/archive/2008/downloads/acceptedPapers/Papers/Baranski.pdf.

Foster, I., Zhao, Y., Raicu, I. & Lu, S. (2008). Cloud computing and grid computing 360-degree compared. [Online]. Available: http://arxiv.org/abs/0901.0131.

Gartner (2008). Gartner Says Cloud Computing Will Be As Influential As E-business, Gartner Press Release. [Online]. Available: http://www.gartner.com/it/page.jsp?id=707508

Kiehle, C., Greve, K., & Heier, C. (2006). Standardized Geoprocessing - taking spatial data infrastructures one step further. In *9th AGILE International Conference on Geographic Information Science* (pp. 273-282). Visegrad, Hungary.

Kralidis, A. T. (2007). Geospatial Web Services: The Evolution of Geospatial Data Infrastructure. In A. Scharl & K. Tochtermann (Eds.), *The*

Geospatial Web, Advanced Infromation and Knowledge Processing Series (pp. 223-228). London, UK: Springer.

OGC. (2007). *OpenGIS Web Processing Service*. OGC implementation specification, Open Geospatial Consortium. Retrieved from http://www.opengeospatial.org/standards/wps.

Roman, D., Schade, S., Berre, A. J. , Bodsberg, N. R. & Langlois, J. (2009). Model as a Service (MaaS). Proceedings of Grid Technologies for Geospatial Applications - pre-conference workshop AGILE 2009. Hannover, Germany.

Schaeffer, B., Baranski, B., Foerster, T., & Brauner, J. (2009). A Service-Oriented Framework for Real-time and Distributed Geoprocessing. In *International Opensource Geospatial Research Symposium*. Presented at the OGRS 2009, Nantes, France: Springer Verlag.

Schaeffer, B., Baranski, B., & Foerster, T. (2010). Towards Cloud Computing in SDIs. In *AGILE 2010*. Guimarães, Portugal: Springer Verlag, in press.