This paper was accepted for publication in Advances in Service-Oriented and Cloud Computing: Workshops of ESOCC 2015. The copyright was transferred to Springer.
The final publication is available at Springer via http://dx.doi.org/10.1007/978-3-319-33313-7_25

BEACON: A Cloud Network Federation Framework *

Rafael Moreno-Vozmediano[1], Eduardo Huedo[1], Ignacio M. Llorente[1], Rubén S. Montero[1], Philippe Massonet[2], Massimo Villari[3], Giovanni Merlino[3], Antonio Celesti[3], Anna Levin[4], Liran Schour[4], Constantino Vázquez[5], Jaime Melis[5], Stefan Spahr[6], and Darren Whigham[7]

[1] Universidad Complutense de Madrid, SPAIN.
{rmoreno,ehuedo,llorente,rubensm}@ucm.es
[2] Centre D'excellence en Technologies de L'information et de la Communication (CETIC), BELGIUM.
philippe.massonet@cetic.be
[3] Università di Messina, ITALY.
{mvillari,gmerlino,acelesti}@unime.it
[4] IBM Israel - Science and Technology Ltd, ISRAEL.
{lanna,lirans}@il.ibm.com
[5] OpenNebula Systems, SPAIN.
{cvazquez,jmelis}@opennebula.systems
[6] Lufthansa Systems, GERMANY.
stefan.spahr@lhsystems.com
[7] Flexiant Limited, UK.
dwhigham@flexiant.com

Abstract. This paper presents the BEACON Framework, which will enable the provision and management of cross-site virtual networks for federated cloud infrastructures in order to support the automated deployment of applications and services across different clouds and datacenters. The proposed framework will support different federation architectures, going from tightly coupled (datacenter federation) to loosely coupled (cloud federation and multi-cloud orchestration) architectures, and will enable the creation of Layer 2 and Layer 3 overlay networks to interconnect remote resources located at different cloud sites. A high level description of the main components of the BEACON framework is also introduced.

1 Introduction

There is a strong industry demand for automated solutions to federate cloud network resources, and to derive the integrated management cloud layer that

* This research was supported by the European Union's Horizon 2020 Research and Innovation Program under the Grant Agreement No 644048.

enables an efficient and secure deployment of resources and services independent of their location across distributed infrastructures. From big companies and large cloud providers interested in unifying and consolidating multiple datacenters or cloud sites to SMEs building hybrid cloud configurations, federated cloud networking is needed to support the automated deployment of applications across different clouds and datacenters.

Many big companies (e.g. banks, hosting companies, etc.) and also many large Government institutions maintain several distributed datacenters or server farms, for example to serve to multiple geographically distributed offices, to implement HA (High Availability), or to guarantee server proximity to the end user. Federated cloud networking is needed to unify and consolidate datacenters in a virtual way, so that different distributed datacenters can be exposed as a single cloud-like virtual datacenter, and networks of different datacenters can be interconnected in a virtual overlay. Some large cloud providers offer different, geographically dispersed regions, so that users can choose to deploy their infrastructures and services in one particular region attending to different criteria, such as proximity, prices, or available resources. Usually these regions are isolated from other regions inside the same provider, to achieve fault tolerance and stability, and there is no interaction or cooperation between them. Federated cloud networking is needed to support distributed services, and provide the overlay networks needed to interconnect servers on different regions, so freeing the service administrator from manually configuring these remote connections. Many SMEs have their own on-premise private cloud infrastructures to support the internal computing necessities and workloads. These infrastructures are often oversized to satisfy peak demand periods, and avoid performance slowdown. Hybrid cloud (or cloud-bursting) model is a solution to reduce the on-premise infrastructure size, so that it can be dimensioned for an average load, and it is complemented with external resources from a public cloud provider to satisfy peak demands. Federated cloud networking is needed to improve this kind of hybrid configurations, so that local and remote resources can be seen as they belonged to the same cloud, and communication channels between these resources can be automatically configured.

Different types of federation architectures for clouds and datacenters have been proposed and implemented [9] (e.g. cloud bursting, cloud brokering or cloud peering) with different level of resource coupling and interoperation among the cloud resources, from loosely coupled, typically involving different administrative and legal domains, to tightly coupled federation, usually spanning multiple datacenter locations within an organization. In both situations, an effective, agile and secure federation of cloud networking resources is key to impact the deployment of federated applications. An integrated cloud management platform able to leverage a federated cloud network will be able to deliver to applications a reliable and secure access to a large geographically dispersed pool of resources.

This paper presents the BEACON Framework[8], funded by an European project (H2020 Program), which will enable the provision and management of

[8] http://www.beacon-project.eu

cross-site virtual networks for federated cloud infrastructures, to support the automated deployment of applications and services across different clouds and datacenters. BEACON is fully committed to open source software. Cloud networking aspects will be based on OpenDaylight[9], a collaborative project under The Linux Foundation, and specifically it will leverage and extend the OpenDOVE[10] project with new rich inter-cloud APIs to provision cross-site virtual networks overlays. The new inter-cloud network capabilities will be leveraged by existing open source cloud platforms, OpenNebula[11] and OpenStack[12], to deploy multi-cloud applications. In particular, different aspects of the platforms will be extended to accommodate the federated cloud networking features like multi-tenancy, federated orchestration of networking, compute and storage management or the placement and elasticity of the multi-cloud applications.

2 Architectures for Cloud Network Federation

Most cloud federation scenarios can be classified into three main federation architectures: datacenter federation (peer cloud architecture), cloud federation (hybrid cloud architecture), and multi-cloud orchestration (cloud broker architecture). In this section, we describe these three main federation architectures, and introduce some security considerations both at application level and architecture level. The BEACON framework will support these different federation architectures, and will enable the creation of different kind of cross-site virtual networks (e.g. Layer 2 or Layer 3 overlay networks), according to the user needs, to interconnect remote resources located at different cloud sites.

2.1 Datacenter Federation and Interconnection

Datacenter federation architecture (see Figure 1) corresponds to a tightly coupled federated cloud scenario [10], also called peer cloud federation, consisting of several private cloud premises (or datacenters) usually belonging to the same organization (or closely coordinated), and normally governed by the same Cloud Manager (CM) type, such as OpenNebula or OpenStack. In this scenario, each CM instance can have full control over remote resources (e.g., placement control, full monitoring, or VM life-cycle management and migration control). In addition, other advanced features can be allowed, including the creation of cross-site networks, the support for cross-site migration of Virtual Machines (VMs), the implementation of high-availability techniques among remote cloud instances, the creation of virtual storage systems across site boundaries, etc. The interaction between CM is usually implemented using private cloud interfaces (administration level APIs) and data models (e.g., OpenNebula XML-RPC[13] or OpenStack

[9] http://www.opendaylight.org
[10] http://wiki.opendaylight.org/view/Open_DOVE
[11] http://www.opennebula.org
[12] http://www.openstack.org
[13] http://docs.opennebula.org/4.4/integration/system_interfaces/api.html

component APIs[14]). On top of the CM there could be a SM to simplify service definition, deployment and management.

Within this architecture, the Network Manager (NM) is responsible for managing virtual networks, both inside and among datacenters. The NM can be integrated with the CM (e.g. OpenNebula Network Manager) or can be a separated component (e.g. OpenDove). NMs in different datacenters interact and cooperate using (possibly private) inter-cloud northbound APIs and protocols (e.g. OpenDayLight Controller REST API[15] or the OpenNebula VirtualNetwork XML-RPC API[16]) that enable the instantiation and management of cross-datacenter networks, mainly based on SDN (Software Defined Networks) and NFV (Network Functions Virtualization) technologies.

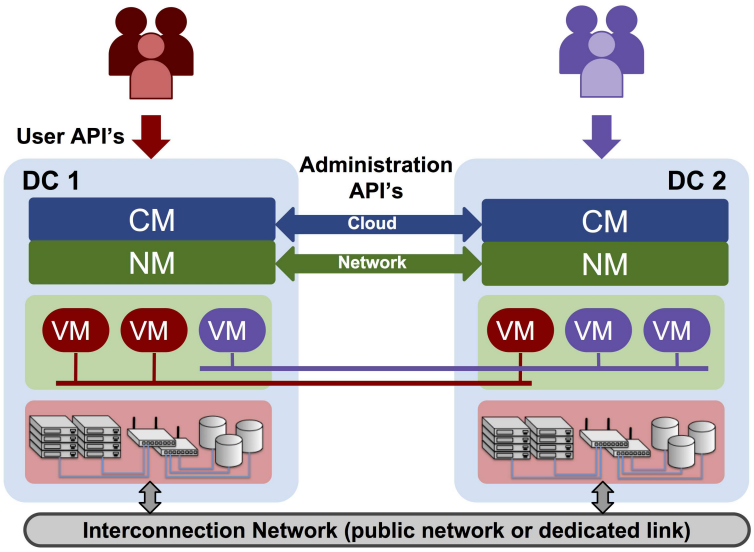

Fig. 1. Architecture for datacenter federation and interconnection.

These cross-site networks are commonly implemented as Layer 2 (L2) or Layer 3 (L3) overlay virtual networks on top of the physical interconnection network, which can be a public network (i.e., a L3 insecure network, such as Internet) or a dedicated high-performance link (usually a private L2 or L3 network). In this context, the most challenging situation is deploying a cross-site secure L2 virtual network over an insecure L3 public connection.

[14] http://developer.openstack.org/api-ref.html
[15] http://wiki.opendaylight.org/view/OpenDaylight_Controller:REST_Reference_and_Authentication
[16] http://docs.opennebula.org/4.12/integration/system_interfaces/api.html#actions-for-virtual-network-management

2.2 Cloud Federation and Interconnection

Cloud federation architecture (see Figure 2) corresponds to a loosely coupled federated cloud scenario that combines multiple independent cloud (both public and private clouds). A typical realization of this architecture is a hybrid cloud [11, 7] or inter-cloud federation, also called cloud bursting model, which combines the existing local cloud infrastructure (e.g., a private cloud managed by a CM, such as OpenNebula or OpenStack) with external resources from one or more remote clouds, which can be either public clouds (e.g. Amazon EC2, FlexiScale, Digital Ocean, etc.), or partner clouds (managed by the same or a different CM).

The main goal of this hybrid model is to provide extra capacity to the local cloud to satisfy peak demand periods, and transforming the local datacenter in a highly scalable application hosting environment. This architecture is loosely coupled, since the local cloud has no advanced control over the virtual resources deployed in external clouds, beyond the basic operations allowed by these providers. The interaction between the local CM and the various remote clouds is usually implemented using public cloud interfaces (user level APIs) and data models (e.g. Amazon AWS EC2 API[17] or OCCI[18]). As in the previous architecture, on top of the CM there could be a SM.

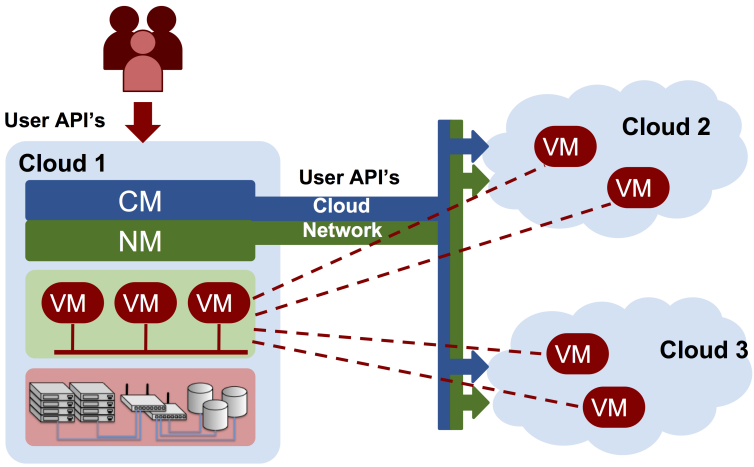

Fig. 2. Architecture for cloud federation and interconnection.

Due to the heterogeneity of network managers (NMs) in different clouds, each cloud can provide different capabilities to interconnect with external resources, regarding the possibility of creating L2 or L3 overlay networks, VPNs, secure channels, or even high level network functions like balancers. In some clouds,

[17] http://aws.amazon.com/ec2
[18] http://occi-wg.org

VMs are seen as independent resources (e.g., Amazon EC2-Classic platform), that can be accessed using a public IP, so the final user is responsible for configuring the appropriate communication channels (e.g. overlay tunnels or VPNs). Other clouds provide private networking to interconnect VMs inside the cloud (e.g. Amazon EC2-VPC platform) and also some kind of VPN capabilities to implement a L3 overlay between local network and remote resources. However, methods to instantiate and configure these VPNs differ from one provider to another. Regarding the creation of L2 overlay networks between independent clouds, currently there are not any cloud technology offering this kind of capabilities, so this is one of the most important challenges in cloud federation and interconnection.

2.3 Multi-cloud Orchestration and Interconnection

Multi-cloud orchestration architecture (see Figure 3), also called cloud brokering architecture [6], usually consists of a central broker or orchestrator, which has access to several public independent clouds. This orchestrator can deploy virtual resources in the different clouds, according to criteria specified by the user, such as location restrictions, cost restrictions, etc., and should also provide networking capabilities to enable the interconnection of different resources deployed in geographically dispersed clouds. There could be also decentralized brokering schemes, with several brokers interacting to each other. We assume that, as in the previous architectures, the orchestrator is basically a multi-cloud SM, which is responsible for managing application and network services across clouds.

Similar to the cloud federation architecture, this architecture is also loosely coupled, since the orchestrator interacts with the different clouds using public cloud interfaces (user level APIs, such as Amazon AWS EC2 API[19] or OCCI[20]), which usually do not allow advanced control over the virtual resources deployed.

Regarding networking issues, the orchestrator must be able to deal with different network managers with different network capabilities, hence it is responsible for creating the required interconnection topologies (e.g. L2/L3 overlay networks) on top of these heterogeneous cloud network services. These overlay networks will be based on virtualized network functions (VNFs) and services, such as bridges, routers, load balancers or firewalls, deployed on the different clouds involved.

2.4 Security Considerations

In BEACON we can have a privileged environment where to enforce and test new security features. Indeed, from the security perspective federated cloud networking provides the opportunity to monitor the virtualized compute, storage and network resources across a federation. This provides opportunities to detect attacks at the federation level that could not be detected at the individual

[19] http://docs.aws.amazon.com/AWSEC2/latest/APIReference/Welcome.html
[20] http://occi-wg.org/about/specification

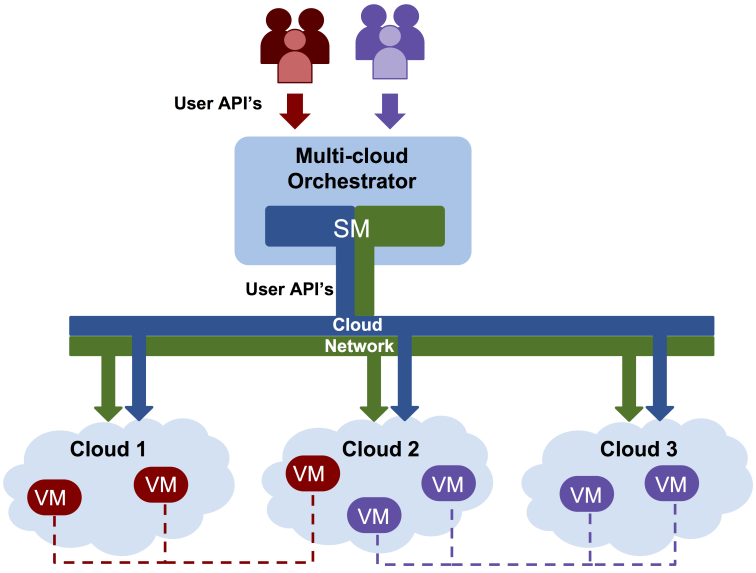

Fig. 3. Architecture for multi-cloud orchestration and interconnection.

cloud level. We can identify many security issues having a global picture of services deployed and executed in more federated Clouds. The security issues we are considering range from the Intrusion Detections, to vulnerabilities scanning, even to the distributed denial of service (DDoS). For example the DDoS attacks might be difficult to detect by monitoring activity within a single cloud. However DDoS attack patterns could be detected earlier by monitoring data from the cloud federation. Within the BEACON project we will identify opportunities for improving detection of threats thanks to the enhanced monitoring capabilities provided by federated cloud networking.

To summarize the work we are providing in BEACON, in Table 1 we classify our security considerations in four different categories for the BEACON architecture. The table considers security issues at the level of the cloud manager and the network manager on the vertical axis, and distinguishes between application level security and infrastructure level security requirements on the horizontal axis. Application level security deals with the security of the application when it is deployed in a federated cloud. Infrastructure level security deals with securing the cloud infrastructure services, i.e. the cloud manager and the network manager, and protecting them from unauthorized access from applications and users. We review the four categories of security issues identified and then conclude that the requirements from the BEACON case studies indicate that application level security needs to be studied at both the cloud manager and network manager levels.

The requirements from the different case studies of BEACON essentially refer to application level security considerations at both the cloud manager level

Component	Application level security	Infrastructure level security
Cloud manager	Applications should be able to request security services from the cloud manager, e.g., to perform vulnerability analysis on a given VM or to apply application level firewall rules to a given HTTP session.	The cloud manager services must be secured with respect to applications running in the cloud and system administrators.
Network Manager	Applications should be able to request security services from the network manager, e.g., to apply firewall rules on one or several network layers, vulnerability analysis at the network level or to apply network intrusion detection.	The network manager services must be secured from unauthorized access, e.g. access to the network controller must be controlled, the communication between the controller and the virtual switches must be encrypted...

Table 1. Application and infrastructure level security considerations.

and the network manager level. The application service manifest should specify required security services to be performed by the cloud manager and the network managers to ensure that the federated cloud meets the security requirements of the application. To guarantee security at cloud and network management at infrastructure levels, it is necessary to analyze the network managers provided by OpenNebula, OpenStack/Neutron and OpenDaylight/OpenDove, to see how they can be integrated and exchange security policies. It cloud be also interesting to analyze the issues related to the location of the network services, e.g. to decide which firewall NFV must be used when several instances are available. This question of which security function to use will also have to take into account live migration of VM within the cloud federation.

3 The BEACON Framework

The main goal of BEACON project is to define and implement a federated cloud network framework that enables the provision of federated cloud infrastructures, with special emphasis on inter-cloud networking and security issues, to support the automated deployment of applications and services across different clouds and datacenters. The implementation of these new federated cloud networking features, that will leverage on Software Defined Network (SDN) technology, include both, the configuration of overlay networks inside different cloud providers, and the interconnection of these overlays among geographically dispersed sites based on various cloud technologies.

One of the key points of this project is that it is fully driven by real industry uses cases proposed by different cloud actors, such as cloud providers, cloud technology developers, and cloud-user companies and institutions, which are represented by the different partners of the project consortium. These use cases

address the different federation architectures described in previous section, such as datacenter federation (peer cloud architecture), cloud federation (hybrid cloud architecture), and multi-cloud orchestration (cloud broker architecture).

Figure 4 depicts a high level view of the BEACON framework architecture, the main components, and the open source projects that will be extended and integrated to implement the BEACON architecture in the case of cloud federation. The proposed network federation model addresses the challenge of federating clouds based on different network technologies in their network backbone as well as in their cloud management platforms.

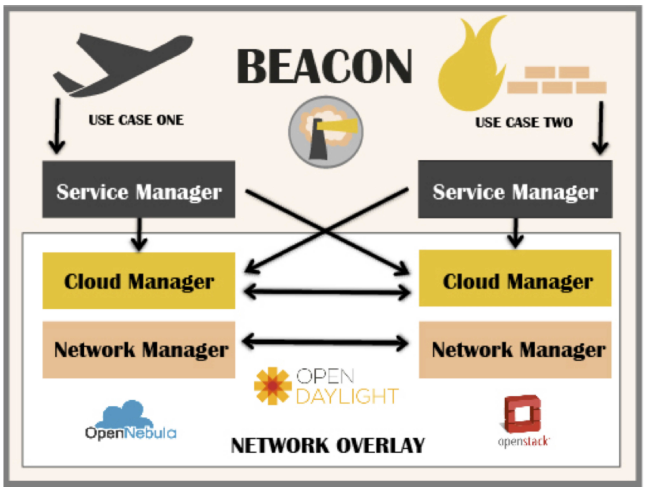

Fig. 4. BEACON federated architecture.

The three main components of the BEACON middleware are the Service Manager, the Cloud Manager and the Network Manager. The Service Manager is responsible for the instantiation of the service application by requesting the creation and configuration of VMs for each service component included in the service definition, using the Cloud interfaces exposed by the cloud manager. The Cloud Manager is responsible for the placement of VMs into physical hosts. It receives requests from the Service Manager through the cloud interface to create and resize VMs, and finds what is the best placement for these VMs that satisfies a given set of constraints (set by the Service Manager) and optimizes a site total utility function. The Cloud Manager is free to place, and move, the VMs anywhere, even on remote sites within the federation, as long as the placement is done within the constraints. The Network Manager is responsible for allocating network resources to manage federated cloud virtual network and overlay networks across geographically dispersed sites. The left and right parts of the figure show two different cloud stacks running on different cloud providers. Together they form a cloud federation with two cloud providers. The middle part of the

figure shows that the cloud manager and network managers of the two cloud providers communicate to share resources and manage the cloud federation. The top of the figure shows two application level case studies that are deployed on the cloud federation (a highly scalable airline application distributed over multiple cloud providers, and multi-cloud security use case). The bottom part of the figure shows the open source projects that are used to implement the federated architecture. The cloud provider on the left part of the figure is using OpenNebula to manage its cloud infrastructure. The cloud provider on the right is using OpenStack to manage its cloud infrastructure. The network managers of both cloud providers are both using OpenDaylight to manage the network resources and supports communications between the two cloud providers. This is an example of heterogeneous cloud federation because two different cloud middleware technologies, i.e. OpenNebula and OpenStack, are being used.

BEACON will develop and integrate OpenDaylight drivers for the overlay network managers of OpenNebula and OpenStack. They will be part of the BEACON framework. This will allow cloud providers, who use either OpenNebula or OpenStack, to form federations and share resources. By forming cloud network federations, the users of these cloud providers will thus automatically benefit from an increased pool of virtualized resources for their applications.

4 State-of-the-art in Cloud and Network Federation

Cloud federation has been an important research field and is still an open issue in cloud computing. In the literature, we can find many different realizations, and research works focussed on the different federation architectures. Regarding the tightly coupled peer cloud architecture, some of the most interesting initiatives are the RESERVOIR project [10], which enables the federation and interoperability of infrastructure providers, taking advantage of their aggregated capabilities to provide a seemingly infinite service computing utility, and the Contrail system [2], which provides collaboration, migration, and SLA management across multiple heterogeneous clouds that can be exploited as a single cloud. There are also various research works that show the advantages of hybrid cloud architectures [11, 7, 12, 15], which enable the transformation the local data center in a highly scalable application hosting environment, by combining the existing corporate infrastructure with remote extra resources from one or more public clouds. This is also the case of the StratusLab initiative, which use the hybrid capabilities of the OpenNebula Cloud Manager to support and provision scalable grid services. Finally, cloud brokering has been one of the most explored federation architectures, both in industry and academia. There are various commercial cloud brokers (e.g. RightScale[21], SpotCloud[22] or Kavoo[23], among others), open-source

[21] http://www.rightscale.com
[22] http://www.spotcloud.com
[23] http://www.kavoo.com

initiatives (e.g. Aeolus[24] or CompatibleOne[25]), and many other research works [13, 6, 4, 14] and projects [3, 5] on cloud brokering, that help cloud customers to cope with a variety of cloud interfaces, instance types, and pricing models, by providing intermediation, arbitrage, and aggregation capabilities. Regarding the networking capabilities of the above mentioned federated platforms (based on peer, hybrid, or broker architectures), most of them rely on public IP addressing to access compute instances deployed in different clouds, or use VPN tunneling mechanisms to improve security that usually are manually configured by the user. However, none of them provides any automatic method or interface to allow a user to instantiate and provision an overlay network across geographically dispersed clouds to interconnect virtual machines deployed in different clouds.

To provide federated networking capabilities, it is necessary a virtual network management system supporting seamless infrastructure, in which services can be deployed on demand across different network platforms and architectures. There are various solutions that provide tools for cloud network management, such as OpenDaylight [8], Contrail controller [2] and federated SDN controller for network virtualization overlays [1]. OpenDaylight is a collaborative project under The Linux Foundation created by leading industry partners with a goal to foster innovation and create an open and transparent approach to Software Defined Networking (SDN). An OpenDaylight controller provides flexible management of both physical and virtual networks. The network management capabilities implemented in OpenDaylight controller allow efficient integration with cloud computing platforms. For example, OpenDaylight is already integrated with Neutron, which provides SDN-based networking solution for OpenStack clouds. In order for OpenDaylight being able to manage heterogeneous networks spread over different cloud computing platforms, it has to be integrated with additional platforms, e.g. OpenNebula. With all the advantages the existing OpenDaylight solution brings to cloud network management, it does not provide a solution for federated cloud network management at its current state. Therefore, it lacks necessary federated cloud management interfaces both to the physical and virtual network elements. In order for the system being able to create and manage simultaneous virtual networks on demand with arbitrary topologies on a loosely coupled federated cloud systems, an additional extension must be defined and implemented in OpenDaylight controller that will allow its integration with federated cloud management systems. This integration should enable virtual network services across federated clouds. The Contrail Controller is a logically centralized but physically distributed SDN controller that is responsible for providing the management of the virtualized network. While the Contrail controller provides control plane, the forwarding plane of the Contrail system is represented by Contrail's virtual routers. Even though Contrail's virtual network management system is integrated with OpenStack, it is limited to the use of the specific virtual routers and does not support commonly deployed open virtual switches (vSwitch). In addition, in order for Contrail controller

[24] http://www.aeolusproject.org
[25] http://www.compatibleone.com

to provide full solution for federated virtualized cloud network management, it needs to be extended to support additional cloud platforms, such as OpenNebula for example. The federated SDN controller for network virtualization overlays is defined in [1]. It addresses the VXLAN and NVGRE overlays managed by federated SDN controller. This controller definition should be extended to support heterogeneous clouds, in order to be able to work in a federated cloud based on different cloud technologies. Also, the controller must include interfaces to the federated cloud management system, which exposes federated cloud services to applications.

5 Conclusions and Future Work

This paper has analyzed three main types of federation architectures: datacenter federation (peer cloud architecture), cloud federation (hybrid cloud architecture), and multi-cloud orchestration (cloud broker architecture). The paper presented the BEACON federated cloud network framework that enables the provision of federated cloud infrastructures, with special emphasis on intercloud networking and security issues. The challenge is to design and develop a framework that can be integrated into different cloud middleware and yet provide support virtual networking and security for the different federation types mentioned above. Future work first involves integrating the BEACON federated cloud network framework into OpenNebula and OpenStack, and experimenting with OpenNebula and OpenStack based cloud federations. In a second phase experimentation will focus on the heterogeneous case where the BEACON framework provides interoperability between OpenNebula and OpenStack clouds within the same federation.

References

1. Balus, F., Stiliadis, D., Bitar, N.: Federated SDN-based Controllers for NVO3 (2012), available at http://tools.ietf.org/html/draft-sb-nvo3-sdn-federation-00
2. Contrail White Paper. Overview of the Contrail system, components and usage (2014), http://contrail-project.eu
3. Ferrer, A., Hernandez, F., Tordsson, J., Elmroth, E., et al.: Optimis: A holistic approach to cloud service provisioning. Future Generation Computer Systems 28, 66–77 (2012)
4. Guzek, M., Gniewek, A., Bouvry, P., Musial, J., Blazewicz, J.: Cloud Brokering: Current Practices and Upcoming Challenges. IEEE Cloud Computing 2, 40–47 (2015)
5. Kavoussanakis, K., Hume, A., Martrat, J., Ragusa, C., et al.: BonFIRE: the Clouds and Services Testbed. In: 5th IEEE International Conference on Cloud Computing Technology and Science (Cloudcom). pp. 321–326 (2013)
6. Lucas-Simarro, J., Aniceto, I.S., Moreno-Vozmediano, R., Montero, R.S., Llorente, I.M.: A Cloud Broker Architecture for Multicloud Environments, chap. 15, pp. 359–376. John Wiley & Sons (2014)

7. Montero, R., Moreno-Vozmediano, R., Llorente, I.: An Elasticity Model for High Throughput Computing Clusters. Journal of Parallel and Distributed Computing 71, 750–757 (2011)
8. Linux foundation collaborative projects. opendaylight - an open source community and meritocracy for software-defined networking (2013), available at http://www.opendaylight.org
9. R. Moreno-Vozmediano, R.S. Montero, I.L.: IaaS Cloud Architecture: From Virtualized Data Centers to Federated Cloud Infrastructures. Computer 45, 65–72 (2013)
10. Rochwerger, B., Caceres, J., Montero, R., Breitgand, D., Elmroth, E., Galis, A., Levy, E., Llorente, I., Nagin, K., Wolfsthal, Y.: The RESERVOIR Model and Architecture for Open Federated Cloud Computing. IBM Journal of Research and Development 53, 4–11 (2009)
11. Sotomayor, B., Montero, R., Llorente, I., Foster, I.: Virtual Infrastructure Management in Private and Hybrid Clouds. Internet Computing 13, 14–22 (2010)
12. Sturrus, E., Kulikova, O.: Orchestrating Hybrid Cloud Deployment: An Overview. IEEE Computer 47, 85–87 (2014)
13. Tordsson, J., Montero, R.S., Moreno-Vozmediano, R., , Llorente, I.M.: Cloud brokering mechanisms for optimized placement of virtual machines across multiple providers. Future Generation Computer Systems 28, 358–367 (2012)
14. Wang, W., Niu, D., Liang, B., Li, B.: Dynamic Cloud Instance Acquisition via IaaS Cloud Brokerage. IEEE Transactions on Parallel and Distributed Systems 26, 1580–1593 (2015)
15. Zhang, H., Jiang, G., Yoshihira, K., Haifeng, C.: Proactive Workload Management in Hybrid Cloud Computing. IEEE Transactions on Network and Service Management 11, 90–100 (2014)