PROCEEDINGS OF SCIENCE

Creating and Managing Dynamic Cloud Federations[*]

Giuseppe Andronico[†], Marco Fargetta, Salvatore Monforte and Maurizio Paone
Istituto Nazionale di Fisica Nucleare, Sezione di Catania. Italy
E-mails: giuseppe.andronico@ct.infn.it, marco.fargetta@ct.infn.it, salvatore.monforte@ct.infn.i, maurizio.paone@ct.infn.it

Massimo Villari
Università degli Studi di Messina. Italy
E-mail: mvillari@unime.it

Cloud computing has evolved from a promising approach to the service provisioning to the reference model for all new data centres to build. Additionally, an increasing number of companies are choosing to migrate their business in the cloud "ecosystem" adopting the solutions developed by the biggest public Cloud Service Providers (CSPs). Smaller CSPs build their infrastructure on technologies available and to better support user activities and provide enough resources to their users, the federation could be a possible solution. In this work, we present different federation models, showing their strengths and weakness together with our considerations. Beside the highlighted existing federation we show the design of a new implementations under development at INFN aiming at maximising the scalability and flexibility of small and/or hybrid clouds by the introduction of a federation manager. This new component will support a seamless resources renting on the base of acceptance of federation agreements among operators.

Additionally, we will discuss how the implementation of this model inside research institutes could help in the field of High Energy Physics with explicit reference at LHC experiments, digital humanities, life sciences and others.

International Symposium on Grids and Clouds (ISGC) 2015,
15 -20 March 2015
Academia Sinica, Taipei, Taiwan

[*] Project CHAIN-REDS – European Community 7th FP Grant Agreement n. 306819, and Project PRISMA – IT PON program cod. PON04a2_A

[†] Speaker.

1. Introduction

Small and medium clouds, to increase the adoption of their services, have to provide improvement for some cross cutting aspects, such as *availability*, *governance*, *interoperability*, *performance*, *portability*, *privacy*, *regulatory*, *security* and many others. These aspects are difficult to manage for single cloud operators and could be even worst when different operators trying to collaborate in order increase the portfolio of reseources and services made available. Therefore, implement models for Intercloud Interoperability and Federation enabling the operator to facilitate their activities supporting the above aspects are an important role in the current scenario.

Practical approaches to federation does not supply with any clearly defined real example leading to some sort of semantic clash on what federation means. In other words, some clouds declare to be federated because of a shared file-system or other distributed or replicated service. This is not true and in order to understand the idea behind our approach it is important to keep in mind the following assertion: federation and resource sharing are two distinct concept with different meaning. So what is federation and what we want federated clouds act as? Let's start simply from a federation definition taken from a common dictionary:

> *Act of joining states or other groups with an agreement in common affairs they will be governed under one central authority.*

Translating this sentence into the cloud world is the idea underlain the proposed approach to federation, nothing is shared among federated clouds members, they have their own resources, users and autonomy but given the federation agreement they belongs to, each member supply the federation with its own resources.

Some standard organisations have started to recognise the need for a standardisation in the are of federation cloud. This is very important because standard can facilitate the co-operation of different operators. Among these organisation an important role has the *IEEE Standard Association* which is working on Intercloud Interoperability and Federation with its project named *"P2302 - Standard for Intercloud Interoperability and Federation (SIIF)"* [9]. The SIIF project aims at developing standard methodologies for cloud-to-cloud interworking and it will be interesting to evaluate how its outcome will impact the federation of clouds but some time is still needed.

In this work we describe a model of cloud federation able to provide scalability and flexibility to a group of small clouds. Create a federation among small cloud operators with heterogeneous and different administration domains and technologies raises many problems. However, it provides business benefits exceeding the drawbacks because the federation as a whole, and so each member, can compare with big cloud players, thanks to the possibility of accessing seamless resources according to federation agreements among the federated operators. The work is in a preliminary stage, but it represents a starting point for investigating and formalising a model able to consider all implications in accomplishing and managing Dynamic Cloud Federations. The model, aimed at small cloud operators, allows them to easily join and leave the federation minimising all possible issues due to the evolving configurations. Moreover, the added-value of this work is in providing a concrete model that looks at heterogeneous cloud systems, in order to include in the federation different cloud middleware (e.g. OpenNebula, CloudStack, etc.).

The paper is organised as follows: Section 2 describes a brief survey on cloud federation models useful positioning our work respect to the State of the Art. In Section 3, we make clearness on the concept of federation, distinguishing it from interoperability and orchestration, presenting the general idea of federation we are dealing with in this paper. Our model is presented in Section 4. Finally, some indication on the goal for INFN and the aspected impact in the area of the physics reasearch is discusse in section 5. Section 6 concludes the work providing highlights for the future.

2. A survey on Cloud Federation Models

Cloud federation refers to mesh of clouds that are interconnected by using agreements and protocols necessary to provide a universal decentralised computing environment. Introducing the federation concept is raising many challenges in different research fields on cloud computing (see [17, 18, 8, 1]). Most of the works in the field concerns the study of architectural models able to efficiently support the collaboration between different cloud providers focusing on various aspects of the federation.

The FP7 European founded project RESERVOIR [15], which operates at IaaS introducing an abstraction layer allowing to develop a set of high level management components that are not tied to any specific environment. Therefore, several sites can share physical infrastructure resources creating a kind of federation, with the condition that all the involved clouds have a homogeneous environment. The experience acquired in RESERVOIR leads up to the latest EU initiative known as FI-Ware [6]. In particular, the EC is encouraging a federated framework based on Fi-Ware platform called XI-FI Federation [7]. Indeed, XI-FI federates homogeneous FI-Ware systems based on OpenStack framework. It is noteworthy the work has been done in the area of formalisation of federation cloud components, which are: *Federate Security*, *Federate Resources*, *Monitoring Resources* and *Define Scalability Rule*. However, XI-FI Federation maintains a static approach for making up the early phases of federation. XI-FI needs to formalize a-priory agreements among the cloud parties interested in joining the federation.

In the work of [2] the authors describe an architectural solution for federation by means of the Cross-Cloud Federation Manager (CCFM), a software component in charge of executing three main functionalities: i) discovery, which allows to exchange information on federated Clouds, ii) match-making, which performs the best choice of the provider that can loan its resources, and iii) authentication, to create a secure communication channel among federated Clouds.

Despite the obvious advantages, the implementation of a federated environment is not trivial at all. Even the OpenStack framework [13] is looking at the possibility to federate two or more OpenStack clouds. In particular, OpenStack initiative, is investigating on *Inter Cloud Resource Federation Models* as described into [14], where the InterCloud Resource Federation Alliance is formalized. In brief, the idea presented is to give partners investing in a joint venture the opportunities to make a bigger cloud entity with massive resources capacity. OpenStack foundation realized that security is one of the main challenge in cloud federation, as from the first item within the list of issues to be overcome in the presented assessments:

> *Security: as Tokens management, Single Sign On features, Resource Access Across Clouds, Data Export Control, etc.*

Therefore, researchers are looking at the possibility to federate users and policies as presented in Cloud Infrastructures [10] exploiting Virtual Organization Membership Service (VOMS) originally conceived for GRID computing. It is also interesting to see works trying to federate Keystones, the Identity and Access Management systems of OpenStack like reported in [16] and [3]. The complexity of Inter-cloud Architectures is well described in [4] where an architectural framework for cloud based infrastructure services provisioned on-demand is presented.

3. Reference Scenario

To face the issues concerning *cloud federation*, two aspects have been isolated and investigated separately by the scientific community. One aspect focuses on cloud interoperability, which mostly consist in the action of devising protocols able to access cloud services on different software systems (e.g., OpenStack, OpenNebula, EC2, etc.), thus the main effort is devoted to the design communication protocols and resources dissemination policies (e.g the EGI federated cloud [5]). This activity has involved several standardisation organisation and produced standards like OCCI [11] and others. The other aspect deals with the definition of the entities operating in the *Cloud Federation*, and the actions these entities need to perform in order to manage the system (e.g.: Federation joining, service negotiation, SLA monitoring, etc.).

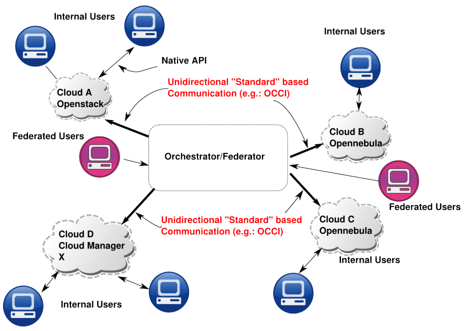

Figure 1: Centralized approach to the federation: the users requests are translated and forwarded to external CSPs by a central entity.

In our opinion, the approach focused on interoperability requires a centralised entity which receives requests from a CC that have to access the federated cloud resources and translates them into requests to external CSPs (see Figure 1). Actually, this model is not a cloud federation following the definition presented in this paper, since users are aware of the different CSPs and there is not cooperation among CSPs. Generally, this approach will imply that users need to adopt different software interfaces to access either their own internal resources or the external ones offered by "federated" sites. Hence, users are divided in *internal users* and *federated users*: the former access cloud services through native APIs, whereas the latter interact with the central entity through federation specific APIs. This simplification of federation presents some issues that can be critical in specific scenarios. Internal users cannot extend their cloud resources by taking advantage of federated CSPs, because they need another external software system that, in turn, will access resources not related to their own cloud. Additionally, internal users may have applications developed on cloud manager specific APIs thus in order to exploit the federated resources such applications have to be rebuilt. Nevertheless, each cloud may provide different services or interfaces (e.g. event notification or monitoring service), which cannot be available on the federation system.

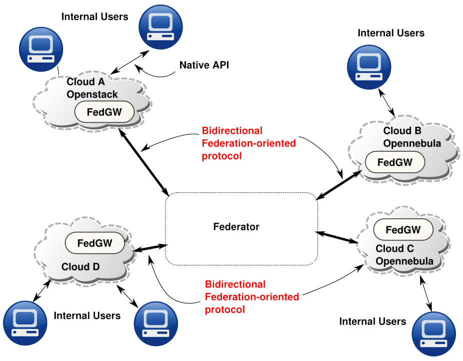

Figure 2: An user-transparent approach to federation: each cloud extends its resource using external federated resources

Differently, the approach focused on the entities definition, as described above, allows to design a system able to transparently extend each cloud including external resources. Figure 2 depicts a simple scheme of such a system. This model implies no distinction between internal and federated users: it defines only *cloud users*, which can access the resources offered by both their own CSP and external federated ones through cloud native interfaces. Most of the harmonization work among the federated CSPs is performed by software running on each site (represented by the FedGW graphical block in the Figure). The entity Federator will carry out operations like resource discovery, marketplace of image templates, and so on.

To better understand the roles played by each actor let us consider the scenario depicted in Figure 3, where clouds A, B and C are small CSPs (Cloud Service Providers), whereas clouds D and E are big enough to internally address any request. CSP D is distributed around the world and its internal interconnection is depicted (link between D and D').

Cloud brokers act as third part intermediary agents, that make their business selecting the best solution satisfying both the CSPs and SPs' requirements. Our interest is to provide clouds A, B and C the same type of business opportunities as for clouds D and E, in which neither brokers nor SPs might be aware of the capabilities each cloud operator supply with.

In this work we analyse the steps required to define a federation agreement under which the cloud operators A, B and C can cooperate maintaining different administration domains. The model presented next treats all the solutions in a general way, hence they can be used also for different cloud middleware (e.g. OpenStack, OpenNebula, CloudStack, etc.).

Since the federation does not involve neither SPs nor brokers it is necessary to setup a common federation framework where all rules and policies are respected. This transparency makes the task difficult considering issues such as compound SLAs (i.e. final SLAs towards SPs is made from a composition of more SLAs) or different network facilities. However, despite the complexity such a federated environment allows CSPs to make new business leveraging their internal infrastructure, but also external renting resource. Thus, each CSP is able to satisfy their customers demands and making profit of unemployed resources by providing them to other CSPs.

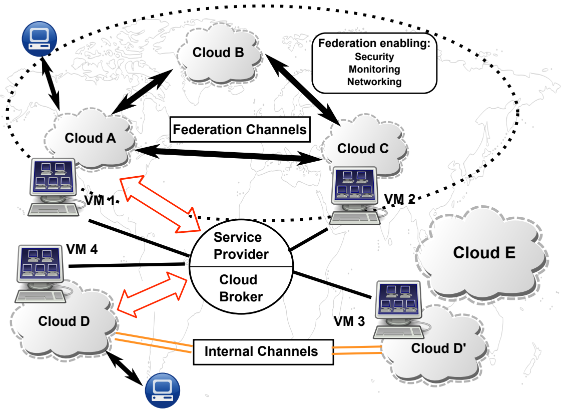

Figure 3: IaaS: Scenario with stand-alone Clouds(D and E), Federated Clouds (A, B and C) and Cloud Brokers, Service Providers, and Customers. Highlights of *In-Federation and Internal Channels*. Moreover Cloud Customers interacting respectively with Cloud A and Cloud D.

We remark the compelling work here is to investigate and formalise a model able to consider all implications required to accomplish all the goal of the federation described above. Section 4 provides all highlights to overcome the problems discussed, trying to minimise all issues due to the evolving configurations of networks, security and monitoring parts and so on.

4. Proposed Cloud Federation Model

In section 3 several approaches to cloud federation have been presented, each and every one having different peculiarities. Nevertheless, none of them comply with our interpretation of the federation leading us to define our own reference model.

According with our model cloud federation life cycle comprises of two distinct moments: join/exit and the resources access. The former is related to the activities performed by a CSP to create or destroy the environment needed by the federation members to communicate each others. The latter is related to the discovery, negotiation and usage of federated resources. The relation between the two moments as well as the actors involved is shown in figure 4.

To join a federation the CSP has to follow several steps as shown in the state diagram depicted in figure 5. In the first state the joining CSP contacts the federation manager sending information about the resources (e.g. cores, storage, etc.) which might potentially be available to the federation parties as well as usage policies on those resources. Federated resources are not dedicated for exclusive use by the federation members but these are upper bounds of the resources available and its real usage depends on the actual request during federation life cycle.

The federation manager, upon join request reception, checks whether the information provided about resources and policies matches with the federation rules or not. If the request is accepted the just joining member is instructed to create a *tenant* with the resources declared in the join request. At this point the CSP can be considered as being federated and ready to fulfil requests from/to others federation members.

A CSP can modify the amount of resources committed to the federation and the policy in any moment but the changes must be notified in advance to the federation manager, who will propagate

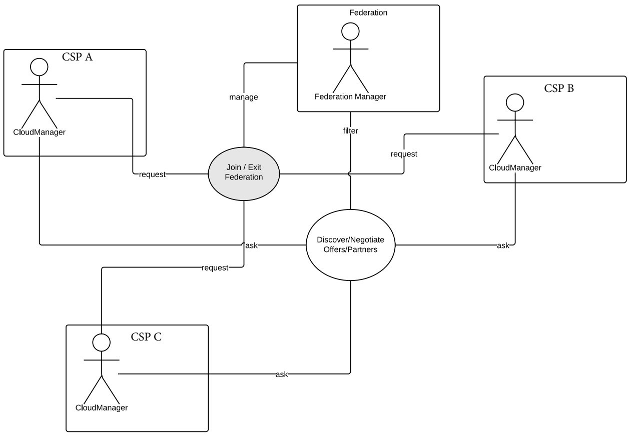

Figure 4: Cloud Federation model - Use Case

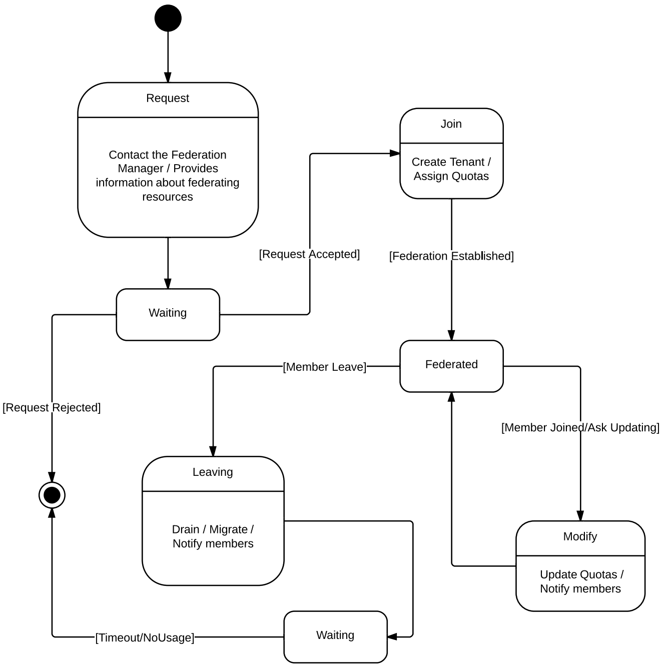
Request
Contact the Federation Manager / Provides information about federating resources
Join
Create Tenant / Assign Quotas
[Request Accepted]
[Federation Established]
Waiting
Federated
[Member Leave]
[Request Rejected]
[Member Joined/Ask Updating]
Leaving
Drain / Migrate / Notify members
Modify
Update Quotas / Notify members
[Timeout/NoUsage]
Waiting

Figure 5: Join and exit federation - State Diagram

the information to all members. Obviously, during the information update the federation can reject a member because it does not comply any more with the rules. A CSP can leave the federation, either for its or the federation manager decision. The federated CSP cannot leave immediately since some resources might be committed and still used, therefore the CSP enters in a leaving state. This state will terminate when all the resources are free or if the leaving period defined in the join agreement has expired, in this case the resource will be forcibly released and remaining data or services discarded. The federation manger notifies all the members about the current disconnection of the CSP. The members have to release the resources the leaving CSP supply with before expiring of a given timeout period.

The federation defines the technical aspects in order to access remote resources and maintains a list of CSPs providing resources with both qualitative and quantitative information. Nevertheless, in order to access member resources a new negotiation is requested between the two members, acting one as CSC and the other as CSP, with the supervision of the federation manager acting as CFA. Figure 6 shows the state diagram related with the discovery and negotiation of federated resources.

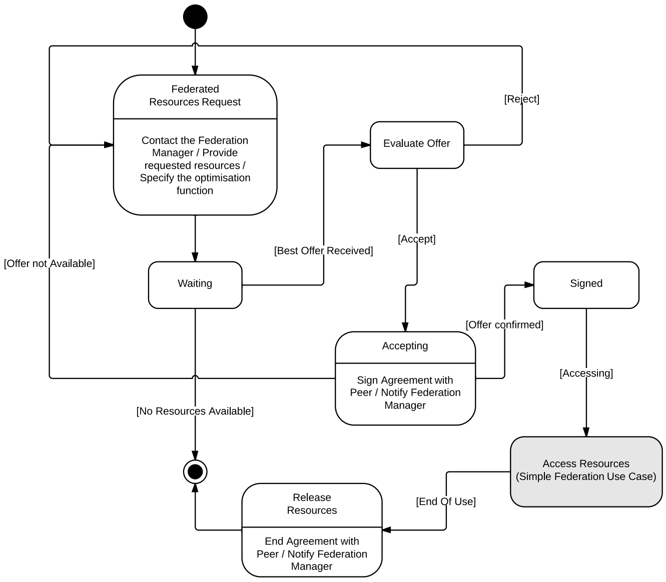
Federated Resources Request
Contact the Federation Manager / Provide requested resources / Specify the optimisation function
[Reject]
Evaluate Offer
[Accept]
[Best Offer Received]
[Offer not Available]
Waiting
Signed
[Offer confirmed]
Accepting
Sign Agreement with Peer / Notify Federation Manager
[Accessing]
[No Resources Available]
Access Resources (Simple Federation Use Case)
Release Resources
End Agreement with Peer / Notify Federation Manager
[End Of Use]

Figure 6: Discovery and negotiate federated resources - State Diagram

To access federated resources the cloud manager has to send a request to the federation manager including the list of requested resources (e.g. number of cores, storage or other) and specify an optimisation function used to pick the best fit among the possible results the CSPs supply with[1]. The optimisation function contains constraints related to the resources, like performance, location, reliability, etc. . . as well as parameters describing QoS / SLA constrains. The federation manager, upon reception of the request, queries the members able to provide the relevant resources based on the information published in the federation about current availability and prices. The optimisation function is then applied to select the best fit for the waiting CSC. This activity is performed automatically and unattended so it does not require any human interaction.

The cloud manager can reject the offer selected by the federation manager and then could send a new request with a different optimisation function. If the offer is accepted an agreement has to be established before the CSC can access the resources. The agreement is an XML document based on WSAgreement [12], which has to be signed by the two parties and the federation because it is responsible for all the relations among its members. Therefore, the federation manager is notified when the agreement takes place and is over, as well. This allows the federation manger to have

[1] The request is an XML document based on WSAgreement and include RDF elements for resource description.

full knowledge of the resources usage among the members and implements strategies for a better distribution and optimisation of workload in the federation.

After the agreement is signed the CSC can start deploying services on the resources of the CSP, upon user requests. The deploy and access to remote resources by the user is shown in Figure 7 and described below. Resources under the agreement are reserved to the CSC and cannot be used by the owner.

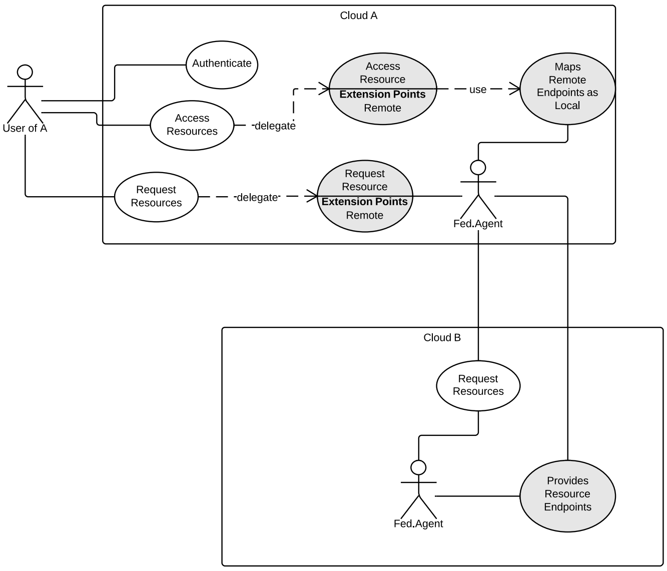
Cloud A
Authenticate
Access Resource Extension Points Remote
use
Maps Remote Endpoints as Local
User of A
Access Resources
delegate
Request Resources
delegate
Request Resource Extension Points Remote
Fed.Agent
Cloud B
Request Resources
Provides Resource Endpoints
Fed.Agent

Figure 7: Request and access federated resources - Use Case

During normal operation, shown in Figure 7, when a cloud user, or any CSC, requires new resources the cloud manager discriminate whether these will be provided as internal resources or taken from the federation. In the former scenario resources are managed by the CSP as usual. In case of federated resources the cloud manager will become a CSC of the federation and will start the negotiation procedure described above. These operation are internally managed by a federation agent inside the cloud. Upon agreement establishment, the required resources are committed in the remote CSP and the relevant endpoints sent to the federation agent who will activate a mapping service to generate local endpoints for the users. The mapping is requested to masquerade the real location of the services. As a result, the user can access the services transparently as resources managed by the cloud itself, hiding to the user the real owner of the resources and their location, which is an important aspect of the federation.

Finally, agreements could be defined in advance, before users request new resources. Moreover, users might release requested resources before the actual expiration of the corresponding agreement, thus leaving them unused within the owning cloud. Hence, internal policies of federated resources usage should be defined and pursued by each member.

5. Cloud federation for research institutes

Many research institutes, like *INFN* (the Italian National Institute for Nuclear Physics) where this work has been developed, have moved to cloud based solutions to support both their internal services (e.g. mail, web, etc...) and the computation needed to run experiments. In fact, computation and storage play a crucial role in the modern scientific investigation and cloud platform has simplified their access for the scientists removing the physical limitation, cloud services are accessed from everywhere by everyone.

INFN manages some sites dislocated in different Italian cities and each site has its computing facilities for the local experiments and, at the same time, provides them to bigger experiments like ALICE, CMS and other[2].

Many experiments require a well defined amount of resources during its lifetime whereas other have some peak requests only in specific moment of the experiment. This makes difficult for a site to define the amount of physical resources to deploy to implement its cloud avoiding waste of moneys and/or long waiting time to access the resources.

At INFN the sites are following two different path: integration and federation. The integration aimed at joining small sites in a unique cloud so the sites can use resources from others. Although this help to mitigate the shortage of resources during peak requests introduce problems in term of management. The administration will not be anymore at site level and this will reduce the agility to support local experiments. Additionally, if peak requests increase resources can still be saturated with the result of an increased waiting time.

The other approach followed by INFN is the federation as described in this paper. This should allow the creation of a federation among INFN sites and between these and external entities, either public or private. With this model should be easier to collaborate in order to run an experiment because the cloud of the partners can be federated and provide the resources needed. Nevertheless, including commercial cloud provider in the federation can create a big pool of resource to use on demand avoiding the problem of overplay the computing facilities.

6. Conclusion and Future Work

In this paper we have presented the idea of cloud cooperation among operators based on federation agreements. The challenge we want to address with the federation is to overcome all the problems raising in merging clouds with heterogeneous administration domains. Therefore, we introduced a high level model of cloud federation able to provide the scalability and flexibility needed by small clouds. The added-value of this work is in providing a high-level model not related to a specific technology which aims at federating different cloud infrastructures.

For the future we are looking at a concrete implementation useful for testing the goodness of our model, but also for providing new features and solving real problems that may occur in cloud federation accomplishments.

References

[1] S. Azodolmolky, P. Wieder, and R. Yahyapour. Cloud computing networking: challenges and opportunities for innovations. *Communications Magazine, IEEE*, 51(7):54–62, July 2013.

[2] ALICE and CMS are some of the experiments related with the Large Hadron Collider developed at CERN

[2] Antonio Celesti, Francesco Tusa, Massimo Villari, and Antonio Puliafito. Three-phase cross-cloud federation model: The cloud sso authentication. In *Proceedings of the 2010 Second International Conference on Advances in Future Internet*, AFIN '10, pages 94–101, Washington, DC, USA, 2010. IEEE Computer Society.

[3] DavidW. Chadwick, Kristy Siu, Craig Lee, Yann Fouillat, and Damien Germonville. Adding federated identity management to openstack. *Journal of Grid Computing*, 12(1):3–27, 2014.

[4] Yuri Demchenko, Canh Ngo, Cees de Laat, Marc X. Makkes, and Rudolf J. Strijkers. Intercloud architecture framework for heterogeneous multi-provider cloud based infrastructure services provisioning. *IJNGC*, 4(2), 2013.

[5] European Grid Infrastructure. Egi federated cloud. https://www.egi.eu/infrastructure/cloud/.

[6] FI-WARE. Open APIs for Open Minds. http://www.fi-ware.org, 2014.

[7] FI-XIFI. Joining The Federation Scenario Exploiting FI-Ware framework. http://wiki.fi-xifi.eu/Public:Joining_the_Federation_scenario, 2014.

[8] I. Goiri, J. Guitart, and J. Torres. Characterizing cloud federation for enhancing providers' profit. In *Cloud Computing (CLOUD), 2010 IEEE 3rd International Conference on*, pages 123 –130, july 2010.

[9] IEEE. P2302 - the ieee standards association. http://standards.ieee.org/develop/project/2302.html, 2014.

[10] A Lopez Garcia, E. Fernandez-del Castillo, and M. Puel. Identity federation with voms in cloud infrastructures. In *Cloud Computing Technology and Science (CloudCom), 2013 IEEE 5th International Conference on*, volume 1, pages 42–48, Dec 2013.

[11] Open Grid Forum. An Open Community Leading Cloud. http://occi-wg.org/.

[12] Open Grid Forum. Web Services Agreement Specification (WS-Agreement). https://www.ogf.org/ogf/doku.php/documents/documents, 2007 (update 2011). GFD-R.192 (Obsoletes GFD.107).

[13] The open source, open standards cloud, innovative, open source cloud computing software for building reliable cloud infrastructure. http://openstack.org/ jan 2014.

[14] Openstack inter cloud resource federation. https://wiki.openstack.org/wiki/Inter-Cloud-Resource-Federation, 2014.

[15] Benny Rochwerger, David Breitgand, Amir Epstein, David Hadas, Irit Loy, Kenneth Nagin, Johan Tordsson, Carmelo Ragusa, Massimo Villari, Stuart Clayman, Eliezer Levy, Alessandro Maraschini, Philippe Massonet, Henar Munoz, and Giovanni Toffetti. Reservoir - when one cloud is not enough. *Computer*, 44:44–51, 2011.

[16] Dinkar Sitaram, H.L. Phalachandra, Anush Vishwanath, Pramod Ramesh, Meghana Prashanth, Akshay G Joshi, Anoop R Desai, Harikrishna Prabhu C R, Prafulla, Shwetha R, and Yashaswini A. Keystone federated security. In *Internet Technology and Secured Transactions (ICITST), 2013 8th International Conference for*, pages 659–664, Dec 2013.

[17] F. Tusa, A. Celesti, M. Villari, and A. Puliafito. How to enhance cloud architectures to enable cross-federation. In *Proceedings of IEEE CLOUD '10*, pages 337–345. IEEE, July 2010.

[18] G. Vernik, A Shulman-Peleg, S. Dippl, C. Formisano, M.C. Jaeger, E.K. Kolodner, and M. Villari. Data on-boarding in federated storage clouds. In *Cloud Computing (CLOUD), 2013 IEEE Sixth International Conference on*, pages 244–251, June 2013.