

Cloud Federation: Characterisation and Conceptual Model

Marcio R. M. Assis and Luiz F. Bittencourt
Institute of Computing - University of Campinas
Av. Albert Einstein, 1251
Campinas, São Paulo, Brazil
Email: {marcio.assis, bit}@ic.unicamp.br

Rafael Tolosana-Calasan
Departamento de Informatica e Ingenieria de Sistemas
Universidad de Zaragoza
c/ Maria de Luna, 1 50018, Zaragoza, Spain
Email: rafaelt@unizar.es

Abstract—This position paper brings to discussion a next step in the evolution of Cloud Computing as the organisation of multiple clouds. We summarise concepts related to the set of clouds called Inter-Cloud and its main elements (hybrid clouds, multi-clouds, etc.). We explore in more detail the Cloud Federation, which has stood as a well-behaved and voluntary organisation of clouds, focusing on its key features, advantages over other types of cloud organisation and existing models. Finally, we present a conceptual model of a Cloud Federation organised in layers, created based on the relationship of the main segments identified in this kind of organisation.

I. INTRODUCTION

Cloud computing as a paradigm aims at providing on-demand elasticity at different levels. However, in reality, the elasticity of a single cloud deployment is limited by the amount of available, which may also limit peak demands. As a consequence, a single cloud deployment may not have the ability to fulfill all the characteristics stated in the cloud computing paradigm [1]. The most noticeable limitation is that smaller providers are not well equipped for peaks in demand for resources [2] [3] [4] [5] [6]. Solutions such as maintaining idle resources to fulfil this demand can act as a palliative, but this can also lead to waste of resources and energy, consequently increasing costs and potentially decreasing resources lifetime. To mitigate this limitation, cloud providers started to organize themselves into more complex infrastructures such as the hybrid clouds, where public clouds are hired to fulfil requests when the private cloud cannot satisfactorily handle them [1], [7].

Following this, other approaches involving the organisation of multiple clouds emerged. Recent researches and implementations have appeared, but as an in-development concept, the multiple clouds organisation and aggregation yet needs characterisation and modelling, as well as standardisation efforts. In this work we aim to stimulate the discussion on architectural aspects of cloud federations, bringing as a contribution the identification and discussion of relevant work in the literature, and raising an initial point for discussion of cloud segments in a layered model for a cloud federation.

This paper is organised as follows. Section II presents the main *Inter-Cloud* characteristics, the large class where the cloud federation belongs to. Section III describes definitions, characteristics and models of the cloud federation. A conceptual model of federation is presented in Section IV, aiming to

raise a discussion about the organisation and management of multiple clouds.

II. INTER-CLOUDS

The idea of using federated Cloud resources in the scope of IaaS, across multiple, distributed Cloud infrastructures was discussed by Keahey et al. [8]. Grozev and Buyya [4] grouped into the inter-cloud class all multiple cloud organisations that have similar objectives (Figure 1). The inter-cloud contains solutions conceived to overcome the limitations presented by isolated clouds to fulfil all characteristics of the cloud computing paradigm, such as high availability and real elasticity. Grozev and Buyya utilise the same definition of inter-clouds presented in the Global Inter-Cloud Technology Forum [6], where:

A cloud model that, for the purpose of guaranteeing service quality, such as the performance and availability of each service, allows on-demand reassignment of resources and transfer of workload through a interworking of cloud systems of different cloud providers based on coordination of each consumers requirements for service quality with each providers SLA and use of standard interfaces.




Fig. 1. Inter-cloud elements.

According to Buyya et al. [4], multiple cloud organisations can be realised under a *voluntary* or *independent* model. In the

voluntary model, clouds participating explicitly know they are taking part in an inter-cloud and their behavior is ruled in terms of a contract that defines how the items so that the resources of each cloud are shared by the other clouds. In the *independent* model, clouds work in conjunction in an orchestrated manner provide by third part service (e.g. OPTIMIS¹), and the participating are not aware of the orchestration or the existence of the other clouds. In this approach we can mention the *hybrid clouds* [9] and the *multi-clouds* [4], [10].

Kurze et al. [11] describe a cloud federation classification that can be extended to inter-clouds. This classification defines how organisation expand their resources, namely: *horizontal*, *vertical*, and *hybrid*. In the horizontal class, the inter-cloud expands within the same type of service. In the vertical class, the multiple cloud organisation expands along a variety of services, as IaaS and PaaS, for instance. In the hybrid class, both vertical and horizontal expansion can take place.

According to the resulting draft from GICT Forum [6], the inter-clouds must have the ability to offer end-to-end quality guarantees for services hosted in their environments, as well as availability guarantees for resources/services and convenience in the cooperation among clouds. The same document describes a set of functional requirements (Table I) which must be present in the inter-clouds.

TABLE I. FUNCTIONAL REQUIREMENTS FOR INTER-CLOUD ORGANISATION.

Functional Requirements
Support for service level agreements - SLAs
Environment and services monitoring
Provisioning according to clients and environment requirements
Resource discovery
Security for services and for the organisation
Resources management in the organisation
Services configuration
Authentication among different networks
Inter-cloud network quality assurance
Data exchange between clouds

III. CLOUD FEDERATION

Cloud federation studies are relatively new in the cloud computing scope. Because of that, some concepts are yet to be clearly established. In one of the pioneering works in this area [4], the authors grouped a series of concepts related to inter-cloud. In [2], the authors mention that a *cloud federation must present at least three characteristics to be effectively recognised as a federation*: (i) capacity to dynamically expand or resize resources to fulfil incoming demand; (ii) be able to operate as a part of a market directed to resource lending; and (iii) be able to deliver reliable services with interesting costs, complying with established quality of service requirements.

Govil et al. [3] points out that this organisation emerged from the alliance of cloud service providers, which joined efforts to offer more resources to their clients and mitigate problems related to lack of SLA compliance. From the authors standpoint, clouds organised as federations provide some advantages: (i) *performance guarantees* – by lending resources, it is possible to maintain the necessary levels of performance

to the rendered services; (ii) *availability guarantees* – location diversity for data and services allows migration of services out from, for example, disaster-prone areas, maintaining higher availability to the client; (iii) *convenience* – federation provides convenience to the client in relation to contracted services, allowing a unified view from services of different providers; and (iv) *dynamic workload distribution* – geographic distribution allows to scatter load according to the client location.

Manno et al. [12] describe the cloud federation as a geographically dispersed community, where many heterogeneous and autonomous clouds cooperate, sharing resources to reach a common objective defined in the federation contract. This contract defines economic and technical aspects of the federation – policies, restrictions, and penalties – that should be followed. The authors state that each cloud is seen as an autonomous domain and can leave the community at any time.

Bringing together the cloud federation definitions stated above by different authors, this work follows an abridged definition that explores the fact that no constraint is imposed to an organisation to contribute to a federation with more than a single private cloud, and emphasising the geographical distribution of the clouds in the federation. Following this, we define: *A cloud federation is an inter-cloud organisation with a volunteer characteristic. It must present a significant geographical dispersion, a well-defined commercialisation system, and must be composed in terms of a Federative Contract that regulates a set of autonomous and heterogeneous clouds. It must be able to provide effective resource scalability, guaranteeing service performance, to realize the dynamic distribution of participating resources, and to respect end-to-end SLAs established with its clients.*

A. Interoperability

Interoperability, the ability for systems to interact each other, is a crucial aspect for forming federation of clouds. It involves the development of interaction protocols and interfaces that must be known in advance for all the interacting parties. Cloud providers offer their functionality to their customers, and to other providers (for inter-cloud purposes) through *application programming interfaces (APIs)*. Current situation in practice is that as there is still little consensus on open and standard APIs for fostering cloud interoperability, each cloud provider has developed its own proprietary API reflecting different interpretations and architectural designs [13]. In consequence, one possible solution (not really flexible) for dealing with such a situation would be to build mediator components among all the interacting providers in the federation.

A better approach is to take the lessons learnt in interoperability from Service Oriented Computing (SOC), and Grid computing based on flexible and fully interoperable architectural components. Such components are exposed as services following the SOC principles, and offer open and standard APIs. The communication among services is based on the remote procedure call (e.g. by means of Web services technologies such as REST). The interfaces should also be platform and programming language independent [13]. This approach requires a common understanding of terms, notions, and architectural relationships involved in a cloud platform.

¹ <http://www.optimis-project.eu/>

B. Federation versus other Inter-Clouds

As previously related, there exist a variety of approaches to organize clouds to fulfil limitations presented by isolated clouds. In comparison to other approaches, cloud federation has a set of characteristics that makes it advantageous in some contexts. In a cloud federation, all participating clouds know in advance they are taking part in a larger entity. Through this knowledge, it is possible to participating clouds to know the current state of the available resources in the federation, mitigating possible losses and ameliorating application scheduling and resource allocation. Moreover, due to the voluntary characteristic of the organisation, it is possible to establish an internal contract who determines the rights and duties (penalties) of each participant. This contract allows minimum service guarantees and defines the expected behaviour within the federated environment.

Grozev and Buyya [4] highlighted that one of the benefits of the cloud federation is the possible dynamic coordination and distribution of workloads among the various clouds. Another benefit is the explicit geographic dispersion, which allows locality optimisation towards better quality of service and response time. The better resiliency furnished by the federation mitigates problems such as paralysation of services since it allows the delegation to another cloud. Besides this, lock-in problems can be avoided among federation participants, who agreed to follow a common objective and are guided by a federative contract. Thus, the client can migrate to another cloud with a better service condition without data storage or interface incompatibility.

From a legal standpoint, the geographical borders in a federation can also be attractive to sensitive data that is location-restricted due to regulations. The service level agreement can establish rules for data placement, and the federation can consider this when distributing the workload, while non-regulated data can travel freely inside the federation.

C. Federation Models

Some architectural models for cloud federation are appearing as the cloud paradigm matures. In this section we discuss four different approaches to model the cloud federation. Such models differ in business focus, internal components, services of interest, and target public (personal clients of service providers). Based on the discussion presented here, we propose a conceptual model in Section IV.

1) *Semantics based*: A theoretical federation model based on semantics and Infrastructure as a Service (IaaS) is proposed in [12]. The authors utilize ontology to provide interoperability between autonomous clouds in a resource sharing environment. The use of ontologies in this scenario is justified by the difficulties in offering interoperability, which is a critical aspect in a federated cloud. The difficulties in providing interoperability include different implementation schemes for the same type of entity or components on each cloud, where ontologies are utilised to understand and model such differences.

2) *Market-oriented*: The federation model oriented to computation services was proposed in [2]. It focuses on the commercialisation of infrastructure resources in a structure that resembles a services market. To support this model, four

components are utilised as the core of the cloud federation: (i) Clouds: where resources are located and the services will be offered to the clients. (ii) Application broker: interface responsible for intermediating operations between client and federation. (iii) Cloud coordinator: component located in each cloud and responsible for maintaining the integrity of the federation. (iv) Concentrator: acts as the market of resources and services.

3) *Reservoir*: The Reservoir [14] is a project led by IBM² and developed with the objective to provide a cloud federation environment that offers software as a service to providers. The Reservoir modelling is focused in loose coupling, and also in avoiding limitations shown by isolated clouds, such as: (i) difficulties small providers have in providing scalability; (ii) lack of interoperability; and (iii) lack of support to business service management (BSM³). The official documentation highlights four functional requirements: *rapid and automatic installation of applications and services, dynamic elasticity, automatic and continuous optimisation, and independency of virtualisation technologies*.

4) *Service-layers-oriented*: Villegas et. al [15] explore services (IaaS, Paas, and SaaS) and their relationships to model a federation. They isolate services into layers and determine that the relationship among clouds in the federation must be performed solely between identical layers. The authors expose the possibility of the clouds in a federation to be heterogeneous not only regarding their resources, but also regarding the layers of services. This related work puts in evidence the information flow and the parameter translation that occur between layers in the same cloud, as well as deals with the brokers that perform the scheduling in the environment.

IV. CONCEPTUAL MODEL

We have hitherto discussed previous work and definitions within the inter-clouds framework. Following the literature review and current understanding specifically regarding cloud federation, we aimed at a model to conceptually represent a layered architecture, illustrating interaction between layers and highlighting their importance in the federation. Therefore, in this section we describe a conceptual proposal for a cloud federation. With this, we have as objective the discussion of a general model to represent the cloud federation in current or yet-to-come models of logical connection among clouds, as for example in the emerging FOG computing paradigm [16]. The conceptual model is illustrated in Figure 2, with its components discussed in details in the remainder of this section.

This proposal was developed considering the challenges listed in [2], [12], [14], [15] and through the mapping of desired functionalities in a federation into components. This concept took into account the independency of technology, hence we do not discuss which technologies could be employed at each level to implement the identified components, since our focus in this paper is to discuss the conceptual model. In Section IV-A the components are organised in segments, and based on relations between components, Section IV-B presents a layered architecture for cloud federation.

²<http://www.ibm.com/>

³Business methodology to align information technology departments/policies to the business strategic objectives

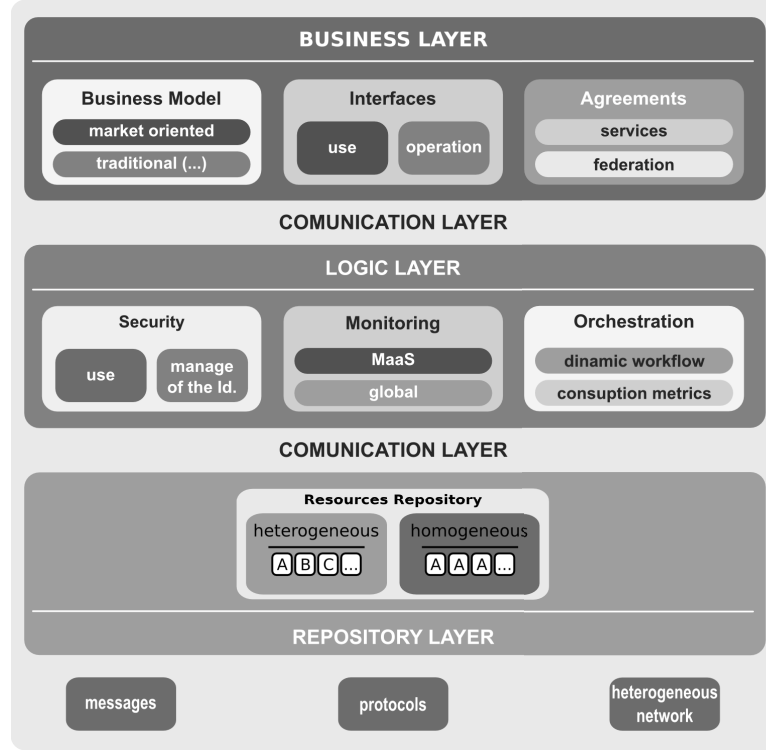


Fig. 2. Proposed 4-layer architecture and their respective components: *business*, *logical*, *repository* and *communication*.

A. Segments in a Federation

The proposed model utilises a layered representation. This model couples components with a certain level of affinity into segments, which are then grouped into layers. The cloud federation idealised in this paper contemplates eight segments: 1) *communication*, 2) *contracts*, 3) *interface*, 4) *commercial model*, 5) *monitoring*, 6) *orchestration*, 7) *resources* and 8) *security*. Hereon, we describe these segments and their components, emphasising their functions and main aspects to consider in future implementations.

1) **Communication:** Components and protocols for data traffic in a cloud federation are in this segment. Due to the distributed characteristic of the federation, data transfers are an important aspect of the architecture. Thus, this segment establishes relation with all the other components in the environment. The communication within the federation is limited to the borders of the clouds that compose the federation. Traffic management inside each cloud is not in the scope of this segment in the proposed architecture.

Relevant points to be considered: (i) Heterogeneous communication channels among federated clouds should perform with a certain level of predictability, avoiding high uncertainties present in public channels. (ii) Due to the geographic dispersion of federation elements, networks with varying characteristics (MTU, latency, etc) are present in the federation. (iii) Modelling the federated cloud network can help in avoiding concurrency and bottlenecks, as well as in routing data flows among clouds.

2) **Contracts:** Service contracts take a fundamental role in cloud environments. It is through these contracts that clients specify their expectations (e.g. QoS and QoE) and terms of use regulated by the providers that offer such a service (e.g., fees and prices). In a federation, besides the contracts between client and service providers, there must exist a *federative contract*, which regulates the functioning of the clouds as components of the federation, and how they should behave in relation to resource offering, monitoring, security, and so on. Moreover, the federative contract must establish penalties for clouds that do not obey rules agreed in the contract.

Relevant points to be considered: (i) The federative contract must state the norms and restrictions of the federation. (ii) Immaturity in contract representations limits the conversion of textual description into logical rules, which can compromise the automation of ensuring contract rules. (iii) Difficulties in expressing Quality of Service (QoS), Quality of Experience (QoE), and Quality of SEcurity (QoSE) are currently challenging issues to be addressed.

3) **Interfaces:** Interfaces are components that promote the interaction between two or more entities in the system. In cloud federations, at least two types of interfaces must be present: *external* and *internal*. The external interface exists between the federation and the outside world, namely the *user interface*. This interface is an extension of classical models representing the interaction portal and/or APIs offered to cloud clients. The internal interface, namely *operations interface*, are inherent to inter-cloud organisations. These interfaces are situated in borders between clouds, thus being responsible for interoperability among federation components.

Relevant points to be considered: (i) Interfaces of the federated cloud must receive, interpret, and execute restrictions defined in contracts. (ii) Interfaces should be able to exchange information regarding the environment: security, monitoring, availability, etc. (iii) Standards should be defined or adapted to enable interoperability between external and internal interfaces of different clouds. (iv) Interfaces should also be platform and programming language independent, requiring a common understanding of terms, notions, and architectural relationships involved in a Cloud platform.

4) *Commercial Model*: The commercial model establishes how resources in the cloud federation are to be negotiated. In an organisation with multiple clouds such as the federation, it is possible to create new commercial models as well as to adapt existing models from other contexts. In the federation model presented in Section III-C2, a market-oriented model is utilised [17], where a centralised entity acts as an intermediate point in the search/offering of resources in the federation as in a marketplace. Moreover, it is possible to implement an autonomous commercial model without interference of the federation, where each associated cloud is free to commercialise its resources according to its convenience, such as in Reservoir (see Section III-C3). In this last model, the federation is utilised solely as a way to improve resource availability.

Relevant points to be considered: (i) Cloud federations offer flexibility in choosing the commercial model to be utilised (model directed to resources market, for instance). (ii) Clients look for a better user experience in these environments (real elasticity). Thus, it is necessary to employ a commercial model that offers guarantees in regard to the client expectations. (iii) Who effectively realizes the commercial model can be external (provided by a third party) or implemented within the federation.

5) *Monitoring*: There are two types of monitoring, differing in their scope: one related to the federation itself (global); and one offered as a service to the clients [18]. As autonomous players in their private domain, the clouds in the federation can implement different monitoring systems. In face of that, the integration of such independent monitoring systems poses a challenge in the federation. The federative contract can establish interfaces for monitoring services and/or monitoring systems supported in the federation.

Relevant points to be considered: (i) Timely monitoring information is crucial to maintain efficient operation in the federated services. (ii) Monitoring as a Service (MaaS) can be offered to the cloud clients. (iii) A global monitoring system must be able to interact with different monitoring mechanisms of each cloud provider participating in the federation. (iv) Monitoring system security can be sensitive, since monitoring can access resources information at different levels.

6) *Orchestration*: Components that perform tasks related to orchestration in the federation are in this segment. The main component in this segment is responsible for performing the mapping of resources to client applications and the workload scheduling within the federation. Also, components that maintain the integrity of the federation are in this layer, verifying the compliance of the members to the federative contract and applying penalties when necessary.

Relevant points to be considered: (i) Dynamic workflows

can be complex to enact in the federation if resources availability is volatile. (ii) Application deployment in the federation must take into account many factors, such as data locality, data transfers, and the topology and architecture of the cloud at different levels. (iii) Metrics, such as reputations, can help in decision making regarding applications scheduling and virtual machine placement.

7) *Resources*: In the proposed model, the main focus is on IaaS offering, thus computational resources are the assets in the federation. Such resources are made available through Virtual Machines (VMs) that slice the physical resources to promote resource consolidation, where VMs have virtual processors, virtual network interfaces, segmented volatile/non-volatile memory space, and so on. In this proposal, clouds participating in the federation are repository alike, where some clouds can offer a variety of heterogeneous resources, while others can offer homogeneous resources. Both types of repository must be able to coexist in the federation.

Relevant points to be considered: (i) An IaaS federation is likely to make available heterogeneous computing assets. However, a cloud participating in the federation can be specialised in a type of asset, such as storage. (ii) As autonomous components, the participating clouds must define in the federative contract how these resources are offered to the federation. (iii) The resources offer and consumption from/by repositories must be accounted through metrics and priorities. Also, penalties can be defined according to these metrics.

8) *Security*: The cloud federation is a distributed environment with potentially a large number of interconnected components. Since each cloud is autonomous, their security policies and internal architecture can be unknown to the other participants. This lack of knowledge can bring security concerns and issues to the cloud federation. Because of this, security in a federated cloud is a critical point to maintain the creditability of the federation. Therefore, many security policies or tools may be necessary to achieve an acceptable level of security in the federation.

Relevant points to be considered: (i) Each cloud in the federation has its own authentication scheme, thus besides the clients access to the federation, access permissions must be managed to allow the federation to work properly [5]. (ii) It is necessary to map both the external and internal borders of the federation so the federation itself knows its acting scope. (iii) The cloud federation must have mechanisms to avoid both external and internal attacks to its resources.

B. Layered Architecture

Although a handful of work related to multi-clouds have been done, which were reviewed in this paper, a full architectural view in terms of layers have not been portrayed yet. Using the eight segments presented above, we present an architecture composed of 4 layers: *business*, *logical*, *repository* and *communication*. The first layer comprises segments related to business aspects of the federation, namely contracts and business model. In the logical layer, the technical aspects of the federation are grouped, namely orchestration, security, and monitoring. The third layer (repository) contains the assets that are offered in the federation. This layered architecture is illustrated in Figure 2.

In the same way as the service layers [15] discussed in Section III-C4, the proposed architecture focuses on a loose coupling. On the other hand, the presented architecture better details the components at each layer, yet allowing flexible implementations of the federation model. For instance, to change the commercialisation model of the federation to a market-driven model (Section III-C1), it is sufficient to change components in the appropriate layer. Also, with layers and the loose coupling, it is possible to vertically extend the service classes offered by the federation. For example, it is possible to insert a PaaS on a federation that provides infrastructure as a service without restructuring the federation. Furthermore, ontologies can be added to the resources layer (Section III-C1). With this, implementation technologies for resource allocation can be abstracted, and compatibility with other types of (newly developed or installed) resources can be achieved.

V. CONCLUSION

As the cloud computing paradigm matures, cloud organisations within more complex infrastructures comes into discussion. These organisations, called inter-clouds, present themselves as a solution to provide some characteristics of the cloud paradigm that are not achieved by isolated clouds. One of the main factors to be considered is the real elasticity that cannot be fulfilled by small- to mid-sized clouds, since they may not have the necessary infrastructure to deal with peak demand. Another desirable characteristic that is not present in many clouds is the geographical dispersion, so the cloud can be close to clients everywhere to provide low latency and response times. In the inter-cloud organisation framework, the cloud federation stands out by having its behaviour governed by a federative contract and by its volunteer characteristic, where all clouds know they are taking part in a federation.

This work presented the state-of-the-art related to cloud federation, where we exploited the main definitions and challenges to the effective implementation of this kind of infrastructure. Moreover, we proposed a conceptual model to be discussed and refined as a federation architecture, initially offered as infrastructure as a service, where the organisational segments necessary to the federation are arranged. These segments are posed in a way that their relationships allow an architectural view of the federation in four layers: *business*, *logical*, *repository* and *communication*. With this proposal, we aim to raise the discussion on cloud federation models and identify relevant problems to be tackled by the community.

As future work, we aim to detail each component of the architecture and identify technologies to be utilised at each layer. Also, we are working in the identification and definition of reputation metrics to serve as a basis of the federative contract and regulate the provision of resources by the clouds in the federation.

ACKNOWLEDGMENT

This work was supported by CAPES, and partially supported by the Spanish Ministry of Economy under the program “Programa de I+D+i Estatal de Investigación, Desarrollo e innovación Orientada a los Retos de la Sociedad”, project identifier TIN2013-40809-R.

REFERENCES

- [1] P. Mell and T. Grance, “The NIST definition of cloud computing,” National Institute of Standards and Technology, Information Technology Laboratory, Tech. Rep., jul 2009.
- [2] R. Buyya, R. Ranjan, and R. N. Calheiros, “Intercloud: Utility-oriented federation of cloud computing environments for scaling of application services,” in *Proceedings of the 10th International Conference on Algorithms and Architectures for Parallel Processing - Volume Part I*, ser. ICA3PP’10. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 13–31.
- [3] V. K. Chaurasiya, K. Srinivasan, K. Thyagarajan, S. B. Govil, and S. Das, “An approach to identify the optimal cloud in cloud federation,” *Institute of Advanced Engineering and Science*, vol. Vol 1, No 1 (2012), 2012.
- [4] N. Grozev and R. Buyya, “Inter-Cloud Architectures and Application Brokering: Taxonomy and Survey,” *Software: Practice and Experience*, 2012.
- [5] A. Celesti, F. Tusa, M. Villari, and A. Puliafito, “Security and cloud computing: Intercloud identity management infrastructure,” in *Proceedings of the 2010 19th IEEE International Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises*, ser. WETICE ’10. Washington, DC, USA: IEEE Computer Society, 2010, pp. 263–265.
- [6] G. I.-C. T. Forum, “Use cases and functional requirements for inter-cloud computing,” Global Inter-Cloud Technology Forum, Tech. Rep., August 2010.
- [7] L. F. Bittencourt, E. R. M. Madeira, and N. L. S. Da Fonseca, “Scheduling in hybrid clouds,” *IEEE Communications Magazine*, vol. 50, no. 9, pp. 42–47, 2012.
- [8] K. Keahey, M. Tsugawa, A. Matsunaga, and J. A. Fortes, “Sky computing,” *Internet Computing, IEEE*, vol. 13, no. 5, pp. 43–51, 2009.
- [9] L. F. Bittencourt and E. R. M. Madeira, “Hcoc: a cost optimization algorithm for workflow scheduling in hybrid clouds,” *J. Internet Services and Applications*, vol. 2, no. 3, pp. 207–227, 2011.
- [10] A. J. Ferrer, F. Hernándeiz, J. Tordsson, E. Elmroth, A. Ali-Eldin, C. Zsigri, R. Sirvent, J. Guitart, R. M. Badia, K. Djemame, W. Ziegler, T. Dimitrakos, S. K. Nair, G. Kousiouris, K. Konstanteli, T. Varvarigou, B. Hudzia, A. Kipp, S. Wesner, M. Corrales, N. Forgó, T. Sharif, and C. Sheridan, “Optimis: A holistic approach to cloud service provisioning,” *Future Gener. Comput. Syst.*, vol. 28, no. 1, pp. 66–77, JAN 2012.
- [11] T. Kurze, M. Klems, D. Bernbach, A. Lenk, S. Tai, and M. Kunze, “Cloud federation,” in *Proceedings of the 2nd International Conference on Cloud Computing, GRIDS, and Virtualization (CLOUD COMPUTING 2011)*. IARIA, sep 2011, Inproceedings.
- [12] G. Manno, W. W. Smari, and L. Spalazzi, “Fcfa: A semantic-based federated cloud framework architecture,” in *HPCS*, 2012, pp. 42–52.
- [13] D. Petcu, C. Craciun, and M. Rak, “Towards a cross platform cloud api,” in *International Conference on Cloud computing and Services Science (CLOSER)*, 2011.
- [14] B. Rochwerger, D. Breitgand, E. Levy, A. Galis, K. Nagin, I. M. Llorente, R. Montero, Y. Wolfsthal, E. Elmroth, J. Cáceres, M. Ben-Yehuda, W. Emmerich, and F. Galán, “The reservoir model and architecture for open federated cloud computing,” *IBM J. Res. Dev.*, vol. 53, no. 4, pp. 535–545, Jul 2009.
- [15] D. Villegas, N. Bobroff, I. Rodero, J. Delgado, Y. Liu, A. Devarakonda, L. Fong, S. Masoud Sadjadi, and M. Parashar, “Cloud federation in a layered service model,” *Journal of Computer and System Sciences*, vol. 78, no. 5, pp. 1330–1344, sep 2012.
- [16] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, “Fog computing and its role in the internet of things,” in *Proceedings of the first edition of the MCC workshop on Mobile cloud computing*. ACM, 2012, pp. 13–16.
- [17] R. Buyya, C. S. Yeo, and S. Venugopal, “Market-oriented cloud computing: Vision, hype, and reality of delivering computing as the 5th utility,” in *Proceedings of the 2009 9th IEEE/ACM International Symposium on Cluster Computing and the Grid*, ser. CCGRID ’09. Washington, DC, USA: IEEE Computer Society, 2009.
- [18] Y. Al-Hazmi, K. Campowsky, and T. Magedanz, “A monitoring system for federated clouds,” in *CLOUDNET*. IEEE, 2012, pp. 68–74.