

Risk-based Dynamic Access Control for a Highly Scalable Cloud Federation

Daniel Ricardo dos Santos, Carla Merkle Westphall, Carlos Becker Westphall

Networks and Management Laboratory
Federal University of Santa Catarina
Florianópolis, Brazil
{danielrs, carlamw, westphal}@inf.ufsc.br

Abstract— Cloud Computing is already a successful paradigm for distributed computing and is still growing in popularity. However, many problems still linger in the application of this model and some new ideas are emerging to help leverage its features even further. One of these ideas is the cloud federation, which is a way of aggregating different clouds to enable the sharing of resources and increase scalability and availability. One of the great challenges in the deployment of cloud federations is Identity and Access Management. This issue is usually solved by the creation of identity federations, but this approach is not optimal. In this paper, we propose an access control system for a highly scalable cloud federation. The presented system is dynamic and risk-based, allowing the use of cloud federations without the need of identity federations. We also present results of a prototype implementation and show that it is scalable and flexible enough to meet the requirements of this highly dynamic and heterogeneous environment.

Keywords- cloud computing; access control; risk; cloud federation

I. INTRODUCTION

Cloud computing is a model for enabling on-demand network access to a shared pool of computing resources [1]. It is widely adopted and provides advantages for customers and service providers.

As cloud computing grows in popularity, new ideas and models are developed to exploit even further its full capacity, increasing efficiency and scalability. One of these ideas is the deployment of cloud federations [2, 3]. A cloud federation is an association among different Cloud Service Providers (CSPs) with the goal of sharing data and resources [4].

However, to make such a scenario feasible it is necessary to develop authentication and authorization models for largely distributed, dynamic and heterogeneous environments.

This problem is usually treated by the deployment of identity federations. An identity federation is a model of identity management where identity providers and service providers share users' identities inside a circle of trust [32].

This solution, nevertheless, is not optimal, since identity federations present problems such as the necessity of attribute and trust agreements, interoperability issues and, in practice, show limited scalability [5]. This paper shows that it is possible to provide authorization in cloud federations without the need for an identity federation.

The difference between cloud federations and identity federations is that cloud federations are built to share resources and identity federations are built to share users and identity information.

In this paper, we propose to use a risk-based dynamic access control to enable authorization in a cloud federation without the necessity, but allowing the possibility, of using identity federations.

The rest of the paper is organized as follows: Section II presents the related work; Section III discusses the concept of cloud federations; Section IV analyses dynamic access control; Section V presents our proposal; Section VI shows some results and Section VII is the conclusion.

II. RELATED WORK

There are two main kinds of work which are related to this paper: those which study cloud federations and authorization in these scenarios and those which propose dynamic access control models.

CLEVER Clouds [6, 7, 8] is a “horizontal federation” model, built on top of a component called Cross Cloud Federation Manager (CCFM), responsible for the discovery of clouds in the federation, finding the best match for resource requests and handling authentication. Based on this architecture, there is the proposal of using federated identity management with a third party identity provider to handle authentication and authorization [9].

The Contrail project [10] is a framework for the construction of cloud federations. It is built upon a set of core components: the Virtual Execution Platform (VEP), the XtreamFS and the Cloud Federation. Contrail is a big project funded by the European Union and is under active development. It also uses federated identity management and provides support for eXtensible Access Control Markup Language (XACML) authorization and the Usage Control (UCON) access control model.

A basic blueprint for the Intercloud is presented in [11] and [12]. In those papers the aggregate of clouds is envisioned based on an architecture comprised of an Intercloud Root, responsible for naming and trust; Intercloud Gateways, responsible for enabling communication between protocols and standards; and finally the clouds. These papers propose that trust be managed by the Intercloud Root, in a configuration that is similar to an identity federation.

Some challenges for access control in highly distributed environments are presented in [13], which compares the Attribute-based Access Control (ABAC), UCON and Risk-adaptive Access Control (RAAdC) models.

The idea of using risk-based access control in cloud computing is presented in [14], where the authors claim this model is adequate to solve the problems presented by multi-tenancy and also that a dynamic environment requires a dynamic access control model. The paper presents a scenario where RAdAC is used to enforce access control among the tenants of a cloud, considering the risks of illegal access to tenants' data by other tenants or by administrators. The paper, however, shows only an overview of the proposal and lacks validation.

Arias et al. [15] proposed a set of metrics, organized in a taxonomy, to be used in the establishment of identity federations in the cloud and to handle access requests. The authors claim that the federated identity management model is hindered by the underlying trust models that must be pre-established, and that the use of risk metrics can mitigate this problem.

The main difference between our approach and the related work is the use of a risk-based access control model to enable the deployment of cloud federations without the need for identity federations. This proposal is detailed in Section V, and a deeper comparison to the related work can be found in the conclusions.

III. CLOUD FEDERATIONS

The cloud computing paradigm has reached a relative success due to its well-known advantages in scalability and cost reduction, but to enable its full potential we must step forward towards cloud federations [16].

As seen in Section II, there are several proposals of architectures for cloud federations in the literature, but they all share a common goal of aggregating different clouds through standard protocols, enabling their interaction and the sharing of resources available in each one. Cloud federation comprises services from different providers aggregated in a single pool supporting resource migration, resource redundancy and combination of complementary resources or services [4].

The main benefits of this new approach are an increase in scalability, availability and interoperability. It also helps in reducing costs of single providers, since the workload may be shared among the members of the federation.

Thinking even further, there are already proposals for an Intercloud, which is a global aggregate of clouds, such as the Internet is a global aggregate of networks [11, 17].

The establishment of cloud federations presents challenges such as the definition of standard protocols and the migration of virtual resources among diverse providers, but the focus of this work is in the security aspects of the federations, especially Identity and Access Management.

Cloud security is a challenge, since providing availability, integrity and confidentiality for a huge number of users and resources in an Internet-accessible environment is not easy. Cloud federations tend to increase concerns because of the increase in the number of users and resources, the use of different protocols and the exchange of sensitive data among providers. Issues such as governance, auditing and risk management are being actively researched for

clouds [18]; these studies must be extended to understand the influences of a cloud federation.

Compliance to regulations and the definition and fulfillment of Security Service Level Agreements (SecSLAs) are also indispensable.

One of the most important issues in the establishment and running of a cloud federation is Identity and Access Management (IAM) [28].

When in a single cloud, it is possible to use traditional IAM procedures and authorization models to handle access control because all of the users and resources are within the same security domain. When resources and subjects are scaled to a federation of clouds, nevertheless, there is the concern with the fact that subjects may come from a different security domain than the resource to which they are requesting access.

To implement authorization using models such as Role-Based Access Control (RBAC) or Attribute-Based Access Control (ABAC), the cloud must use information provided by a system about a user. This information may be, for instance, the user's identity or attributes of this identity, such as name, organizational role and date of birth.

For a cloud to trust the identity or attribute information of a user that comes from another cloud, both clouds must share some agreement of trust. That is why this process is commonly mediated by an identity federation. With Federated Identity Management (FIM), every participant of the federation is expected to agree that the information received by another participant is correct, in what is called a Circle of Trust (CoT).

A problem with this approach is the fact that this agreement requires previous negotiation, which may be an extensive process and hinder dynamic collaboration. Dynamic collaboration is achieved when entities which have a need to collaborate can instantly form a federation, without the need for a previous trust agreement.

Another problem faced by identity federations is the extensive number of protocols and standards, which actually reduces interoperability. Federations tend to get bigger and bigger and users may participate in different federations. All of those facts combined lead to, in practice, a limited scalability of identity federations, reducing their effectiveness in real world.

But even with the formation of identity federations and the possibility of using ABAC, there are challenges to be considered. The static policies which are predefined to be used in traditional access control models cannot comprehend every possible access situation, because in the cloud this is an ever changing process, with users and resources being deployed and deleted all the time. Static models, thus, lack the flexibility necessary to support exceptional situations, which are common in military and medical applications, among others [19] and important for collaboration and information sharing [20]. Examples of these exceptional access requests are given in the next section and are abundant in the literature.

IV. DYNAMIC ACCESS CONTROL

Identity and access management encompasses several processes related to the identification, authentication, authorization and accountability of users in computer systems [21]. Authorization or access control is the process through which a system guarantees that access requests are validated using well-established rules. These rules are known as policies and the way through which the policies are enforced together with the mechanisms used in this enforcement is known as an access control model.

Classical access control models are known to present problems in highly distributed and dynamic environments [13], especially scalability and flexibility limitations and the use of static policies [14]. Role-based models, for instance, lack granularity of control, because roles share their permissions with every user they are attributed to.

To enable more flexible access control decisions, which reflect current needs for information sharing and allow for a secure handling of exceptional requests, dynamic access control models were developed [22, 26, 29, 33].

In contrast with classical models, dynamic access control has the characteristic of using more than predefined policies to compute access decisions. These models are based on dynamic characteristics, which are assessed in “real time” as the subject requests access to a resource. Characteristics such as trust, context, history and risk are often used to reach decisions, and exactly which characteristics to use and how to measure them is discussed in several works [23, 24, 25, 26].

Risk-based access control models are often used as a “break-the-glass” mechanism, allowing for exceptional access requests to be handled by the system more effectively than simply granting full access [13, 30].

Exceptional requests and special access are sometimes necessary in medical and military applications, among others. A well-known example is in a healthcare facility where only doctors have access to patients’ histories, but in the case of an emergency, a nurse may need to access this information to save a patient’s life. If this kind of situation was not predicted in any policy, either the nurse won’t be able to perform his/her duty or the nurse may be given a doctor’s access, which may grant a broader access than the necessary in this case, allowing misuse. In either case, it represents a greater risk to the system than if a dynamic access control system were used and the access control needs were evaluated per request.

Granting special access in exceptional cases usually involves some form of monitoring by the system. It may be in the form of: obligations, which are post-conditions that a user must fulfill in order to keep his or her access right [13]; a reputation system, which logs users’ actions and assigns rewards and penalties to them [26]; or a market system, in which users have a limited amount of points that may be used to “buy” exceptional accesses [27].

Supporting this kind of access control involves an effective logging system for posterior audit and incident response.

There are several different approaches to risk-based access control, but they all share some common features. Fig. 1 presents an overview of a risk-based access control model.

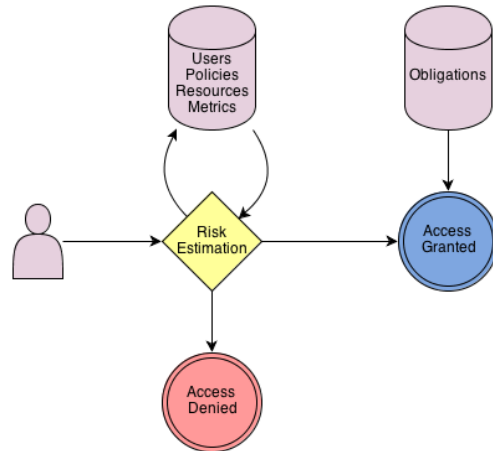


Figure 1. Risk-based access control overview

The figure is based on common points found in diverse models, and the main elements present are the subject, the resource and the risk estimation engine.

The subject tries to access a resource by issuing an access request, which is then processed by a risk estimation engine that uses all the information it deems necessary to come to a decision. Usually there is a risk threshold defined by the system administrators, and if the risk is lower than this threshold, access is granted. Other variations measure risk versus benefit of an access, and decide based on which one is greater [31].

V. PROPOSAL

In this paper, we propose that it is possible to provide a way to establish cloud federations without the need for identity federations, by using risk-based access control and relying on the authentication provided by each cloud. This can increase the scalability of this model and handle exceptional requests.

A. Cloud Federations

Fig. 2 presents an overview of the cloud federation architecture that we are considering. This architecture is based on the common points found in the main federation projects currently being developed, some of which were described in Section II.

The main application scenarios for such federations are medical, military and scientific collaborations, which require large storage and processing capabilities, as well as efficient information sharing. In this architecture we have the following components:

CloudProvider: this is the Cloud Service Provider (CSP) itself, who provides the infrastructure over which the virtual resources are allocated (they are represented by the clouds in the figure);

CloudManager: responsible for attaching a CloudProvider to the federation. It is composed of several services that deal with users, resources, policies, service-

level agreements, security and the CloudProvider. It is modular so that it can be attached to different cloud management software just by changing one of its services.

FederationManager: responsible for coordinating the federation. It acts as a naming service and is also responsible for message passing.

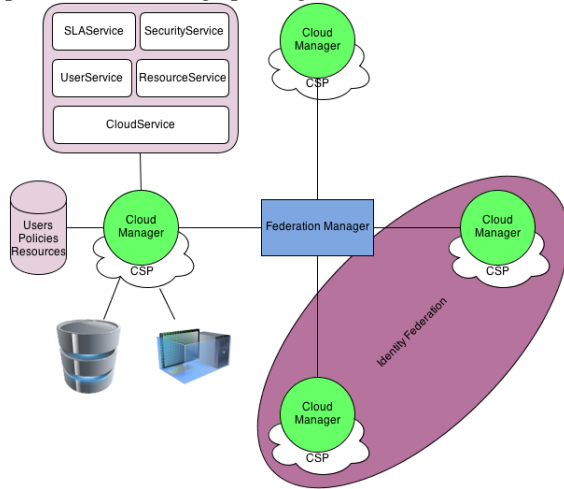


Figure 2. Overview of the federation

B. Access Control

As shown in Fig. 2, some of the participating clouds may form identity federations among themselves.

Under the point of view of a user there are two types of clouds in this architecture: a home cloud (the user’s original CSP) and foreign clouds (the other clouds in the federation). Users can deploy and access resources in both types of cloud, but access control behaves differently for each case.

When users deploy a resource in their home cloud they may choose if it will be available for users of foreign clouds. In any case the user must upload an XACML policy file together with the resource, which will be used for ABAC.

Users may also deploy resources in a foreign cloud and it will automatically be available to every user of the federation. Finally, users may access resources in their home cloud or shared resources in foreign clouds.

When a user tries to access a resource in their home cloud, this request is handled by a classical ABAC model. Based on user attributes and XACML policies the system grants or denies the requested access.

When a user tries to access a resource in a foreign cloud, the system first verifies if both clouds are in an identity federation, in which case the access will also be handled by ABAC, but if there is not an identity federation between them, the “break-the-glass” mechanism is activated and the risk-based access control Policy Decision Point (PDP) is called.

The PDP is located in the cloud handling the access request (foreign to the requester) and the metrics and parameters of risk estimation are defined by the administrators of this cloud and the users who own the resources.

These metrics are informed in an eXtensible Markup Language (XML) file, containing definitions of risk metrics

and how to measure and aggregate them, as well as a threshold level for granting access to the resource and possible obligations that users will have to follow. This file is known as a risk policy.

Each cloud provider must provide a set of basic metrics with their quantification rules. Those will be used to create a baseline risk policy for the provider. This guarantees that a cloud provider is able to maintain their minimal security requirements.

Each resource has its own risk policy, which must respect what is defined in the baseline policy, but may be extended to become more or less restrictive as the user desires. The XML file of the policy must be uploaded by users when they choose to deploy a shared resource. The system does not generate risk policies on the fly and all the risk policies must follow a predefined XML schema, so that different clouds can communicate.

If a user chooses to define a risk metric that is not available in the server, he/she must provide a way for the CSP to quantify this risk. This is done by defining a Web Service that will be called by the PDP upon the evaluation of the access request. The PDP will forward the access request to the Web Service, which will have to parse it, process it and return a numeric value representing the associated risk for the metric being evaluated.

To handle the access request for a given resource all of the metrics are valued, based on the rules defined by the CSP and the Web Services defined by the user. The chosen aggregation engine is used to reach a final risk value. This value is then compared to the defined threshold and, if lower, the subject is given special access.

Before granting access, however, the policy is analyzed in search of obligations that were defined by the user. Those obligations are stored in a system monitor, which will watch and log every user action once the access is granted.

Fig. 3 shows an example of a risk policy file. In this example a metric for transport layer encryption will be quantified, along with other metrics. They will be aggregated based on a maximum value rule. If the final value is lower than 10, access will be granted.

```

<risk-ac>
  <resource id="1"/>
  <user id="2"/>
  <metric-set name="transport layer">
    <metric>
      <name>Transport Layer Encryption</name>
      <description>Quantifies the strength of the encryption scheme
        used in the access request</description>
      <quantification>https://example.com/quantify-tl-encryption
      </quantification>
    </metric>
  </metric-set>
  <aggregation-engine>maximum_value</aggregation-engine>
  <risk-threshold>10</risk-threshold>
</risk-ac>

```

Figure 3. XML risk policy example

Fig. 4 presents a step-by-step flow of the handling of an access request in a foreign cloud. In this figure, step 1 is the issuing of an access request from a user to a foreign shared

resource (since resources that are not shared are not visible to foreign users). The Policy Enforcement Point (PEP) receives this request and forwards it to a PDP (step 2). The PDP verifies if the user's home cloud and the foreign cloud participate in the same identity federation. If they do, then the PDP requests the XACML policies applicable to the resource (step 3a), the Policy Access Point (PAP) responds to this request (step 4a), and the PDP retrieves the necessary attributes from the Policy Information Point (PIP) in steps 5a and 6a.

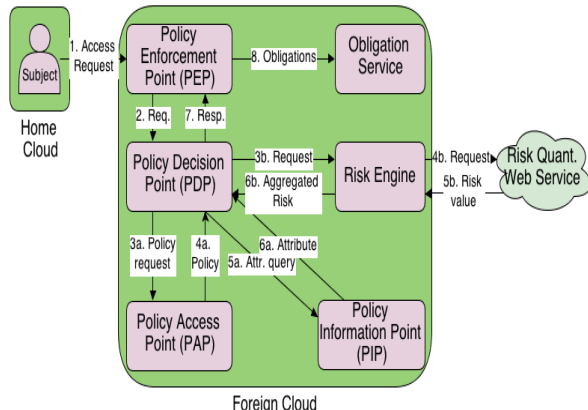


Figure 4. Access control step-by-step

Steps 3a to 6a represent a classical XACML access control decision. However, if the user's home cloud and the foreign cloud are not participants of the same identity federation, the PDP will forward the access request to a risk engine (step 3b). This risk engine will then parse the XML risk definition file associated with the resource and quantify the metrics defined. If the quantification rules are local, the predefined functions are called, if any of the rules are defined in a web service, then it is invoked, having the access request as a parameter (step 4b). The risk quantification web service performs its role and returns a risk value. After all of the metrics are valued, the risk engine applies an aggregation rule, which is always local. The aggregated risk is then returned to the PDP, which uses this value to decide upon the granting of the access request, once again based on what is defined in the XML file. After reaching a decision, the PDP returns it to the PEP, which applies the necessary obligations.

The dynamic nature of access control is present in the system because the access decision may vary according to contextual information evaluated by the metrics.

VI. RESULTS

To validate our proposal and measure some performance characteristics we implemented the key parts of the federation system and the whole access control system.

The implementation used the Python programming language, the zeromq library to handle message passing, MySQL for persistence, the ndg-xacml library for XACML evaluations and the web.py framework for the web services.

The infrastructure over which the federation was deployed was composed of two OpenNebula clouds running on a laptop with a 2,53GHz Core i5 processor and 4GB of RAM. All of the experiments were repeated 50 times to obtain the averages and all of the times shown here refer only to the execution of the access control decision function, ignoring message passing between the clouds. Table I shows four different cases of access request. Case A represents 10 requests handled by local XACML only; case B represents a risk decision that involves 10 risk quantification rules performed locally; case C uses 5 local rules and 5 external (web service) rules; and case D represents a risk policy with 10 external risk quantification rules.

TABLE I. COMPARISON OF DIFFERENT ACCESS REQUEST CASES

	min. (ms)	max. (ms)	average (ms)
A	1.057	9.372	1.46
B	1.824	15.564	4.574
C	1556.182	2813.56	1726.71
D	3247.563	10350.5	4220.6

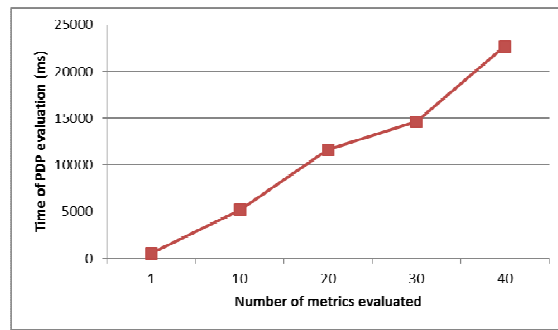


Figure 5. Performance with a varying number of external metrics

It is possible to see that the use of local risk quantification rules has no significant impact on performance, while the use of web services does affect performance, as expected, because of the HTTP invocations that must be performed for each metric.

Fig. 5 shows the growth in time spent reaching an access decision as we increase the number of metrics which call web services in a risk policy file.

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we proposed a risk-based dynamic access control system to enable cloud federations without the need, but allowing the possibility, of identity federations. By eliminating the need for identity federations our proposal eases the use of cloud federations, since it doesn't depend on the establishment of agreements and circles of trust, also enhancing scalability, by avoiding the formation of "identity islands" [5].

The main contributions of this paper are the definition of a risk-based access control system for cloud federations and the proposed use of risk policies in the form of XML files to allow the use of different risk metrics and quantification methods that are not necessarily predefined.

The proposal is flexible enough to handle the needs of a cloud federation and the performance evaluations indicate

that it is scalable and that the risk estimation process is not a big hindrance in the process, especially if the quantification is performed locally.

In comparison to the related work we first have to clarify that we have not implemented a whole cloud federation system such as [6, 7, 8, 10], since it is a huge task and not our focus. We have, however, described and implemented a simple federation model that is sufficient for our access control research and we can highlight that our proposal is the only that uses risk-based access control. Also, we still allow the use of identity federations, but offer a choice of establishing the cloud federation without the need for Federated Identity Management.

Compared to the works that deal with risk-based access control in cloud [14, 15], our approach has the advantage of allowing the resource owner to choose different risk quantification and aggregation engines through a risk policy definition file, also the cloud that hosts the resource can define a baseline risk policy, to ensure its minimum security requirements are met.

As future work we foresee the possibility of enlarging the federation used in our experiments and deploying it to real use. Also, we want to explore further the use of the risk policies with different risk metrics and quantification methods.

REFERENCES

- [1] P. Mell and T. Grance, "The NIST Definition of Cloud Computing", 2011. Available at: <http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf> [Retrieved : May, 2013]
- [2] E. Carlini, M. Coppola, P. Dazzi, L. Ricci, and G. Righetti, "Cloud Federations in Contrail", Proc. Euro-Par 2011: Parallel Processing Workshops, 2012, pp. 159-168
- [3] B. Rochwerger et al., "The reservoir model and architecture for open federated cloud computing", IBM J. Res. Dev., vol. 53, no. 4, July 2009, pp. 535-545
- [4] T. Kurze, M. Klems, D. Bermbach, A. Lenk, S. Tai, and M. Kunze, "Cloud Federation", The Second International Conference on Cloud Computing, GRIDs, and Virtualization, September 2011, pp. 32-38
- [5] K. Lampropoulos and S. Denazis, "Identity management directions in future internet", IEEE Communications Magazine, vol. 49, no. 12, December 2012, pp. 74-83
- [6] A. Celesti, F. Tusa, M. Villari, and A. Puliafito, "How to Enhance Cloud Architectures to Enable Cross-Federation", Proc. 3rd IEEE International Conference on Cloud Computing, July 2010, pp.337-345
- [7] A. Celesti, F. Tusa, M. Villari, and A. Puliafito, "Security and Cloud Computing: InterCloud Identity Management Infrastructure", 19th IEEE WETICE, June 2010, pp. 263-265
- [8] A. Celesti, F. Tusa, M. Villari, and A. Puliafito, "Three-Phase Cross-Cloud Federation Model: The Cloud SSO Authentication", 2nd AFIN, July 2010, pp. 94-101
- [9] A. Celesti, F. Tusa, M. Villari, and A. Puliafito, "Federation Establishment Between CLEVER Clouds Through a SAML SSO Authentication Profile", International Journal on Advances in Internet Technology, vol. 4, no. 12, 2011, pp.14-27
- [10] M. Coppola et al., "The Contrail approach to cloud federations", Proc. ISGC'12, 2012
- [11] D. Bernstein, E. Ludvigson, K. Sankar, S. Diamond, and M. Morrow, "Blueprint for the Intercloud - Protocols and Formats for Cloud Computing Interoperability", Proc. 4th ICIW, May 2009, pp. 328 - 336
- [12] D. Bernstein and D. Vij, "Intercloud Directory and Exchange Protocol Detail Using XMPP and RDF", Proc. 6th SERVICES, July 2010, pp. 431-438
- [13] V. Suhendra, "A Survey on Access Control Deployment", Proc. FGIT-SecTech, 2011, pp. 11-20
- [14] D. Fall, G. Blanc, T. Okuda, Y. Kadobayashi, and S. Yamaguchi, "Toward Quantified Risk-Adaptive Access Control for Multi-tenant Cloud Computing", Proc. 6th JWIS, October 2011
- [15] P. Arias-Cabarcos, F. Almenáez-Mendoza, A. Marín-López, D. Díaz-Sánchez, and R. Sánchez-Guerrero, "A Metric-Based Approach to Assess Risk for "On Cloud" Federated Identity Management", Journal of Network and Systems Management, vol. 20, no. 4, 2012, pp. 513-533
- [16] P. Harsh, Y. Jegou, R. Cascella, and C. Morin, "Contrail virtual execution platform challenges in being part of a cloud federation", Proc. 4th ServiceWave, 2011, pp.50-61
- [17] R. Buyya, R. Ranjan, and R. Calheiros, "InterCloud: Utility-Oriented Federation of Cloud Computing Environments for Scaling of Application Services", Algorithms and Architectures for Parallel Processing, vol. 6081, Springer, 2010, pp.13-31
- [18] D. Catteddu and G. Hogben, "Cloud Computing: benefits, risks and recommendations for information security", Technical Report. European Network and Information Security Agency, 2009
- [19] B. Farroha and D. Farroha, "Challenges of operationalizing dynamic system access control: Transitioning from ABAC to RADAC", Proc. SysCon, March 2012, pp. 1-7
- [20] F. Salim, J. F. Reid, and E. Dawson, "Towards authorisation models for secure information sharing : a survey and research agenda." The ISC International Journal of Information Security, vol. 2, 2010
- [21] B. McQuaide, "Identity and Access Management", Information Systems Control Journal, vol. 4, ISACA, 2003
- [22] JASON Program Office, "Horizontal Integration: Broader Access Models for Realizing Information Dominance", Technical Report. MITRE Corporation, December 2004
- [23] D. W. Britton and I. A. Brown, "A Security Risk Measurement for the RADAC Model.", Naval Postgraduate School Thesis, 2007
- [24] Q. Wang and H. Jin, "Quantified risk-adaptive access control for patient privacy protection in health information systems", Proc. 6th ASIACCs, 2011, pp. 406-410
- [25] Y. Li, H. Sun, Z. Chen, J. Ren, and H. Luo, "Using Trust and Risk in Access Control for Grid Environment", Proc. SECTECH, 2008, pp. 13-16
- [26] R. A. Shaikh, K. Adi, and L. Logrippo, "Dynamic risk-based decision methods for access control systems", Computers & Security, Elsevier, vol. 31, no. 4, 2012, pp. 447-464
- [27] I. Molloy, P. Cheng, and P. Rohatgi, "Trading in risk: using markets to improve access control", Proc. NSPW, 2008, pp. 107-125
- [28] D. N. Sriram, "Federated Identity Management in Intercloud", Der Technischen Universität München Thesis, January 2013
- [29] R. Lepro, "Cardea: Dynamic Access Control in Distributed Systems", Technical Report, NASA, November 2003
- [30] A. Ferreira et al., "How to Securely Break into RBAC: The BTG-RBAC Model", Proc. ACSAC '09, 2009, pp. 23-31
- [31] L. Zhang, A. Brodsky, and S. Jajodia, "Toward information sharing: benefit and risk access control (BARAC)", Proc. 7th IEEE POLICY, 2006, pp. 9-53
- [32] H. Lee, I. Jeun, and H. Jung, "Criteria for evaluating the privacy protection level of identity management services", Proc. SECURWARE, 2009, p. 155-160
- [33] J. Tigli, S. Lavirotte, G. Rey, V. Hourdin, and M. Riveill, "Context-aware Authorization in Highly Dynamic Environments", International Journal of Computer Science Issues, vol. 4, no. 1, 2009, pp. 24-35