

OpenID Authentication As A Service in OpenStack

Rasib Hassan Khan^{*‡}, Jukka Ylitalo[†] and Abu Shohel Ahmed[‡]

^{*}Aalto University, School of Science and Technology, Finland

[†]Ericsson Research, Finland

[‡]Royal Institute of Technology (KTH), School of ICT, Sweden

rkh@cc.hut.fi, jukka.ylitalo@ericsson.com, ahmed.shohel@ericsson.com

Abstract—The evolution of cloud computing is driving the next generation of internet services. OpenStack is one of the largest open-source cloud computing middleware development communities. Currently, OpenStack supports platform specific signatures and tokens for user authentication. In this paper, we aim to introduce a cloud platform independent, flexible, and decentralized authentication mechanism, using OpenID as an open-source authentication mechanism in OpenStack. OpenID allows a decentralized framework for user authentication. It has its own advantages for web services, which include improvements in usability and seamless Single-Sign-On experience for the users. This paper presents the OpenID-Authentication-as-a-Service APIs in OpenStack for front-end GUI servers, and performs the authentication in the back-end at a single Policy Decision Point (PDP). Our implementation allows users to use their OpenID Identifiers from standard OpenID providers and log into the Dashboard/Django-Nova graphical interface of OpenStack.

Keywords—Authentication; EC2API; OpenID; OpenStack; OS-API; Security;

I. INTRODUCTION

Cloud computing is a new paradigm for utilization of scalable resources over the internet. It is a relatively new cyber-infrastructure, implying a service oriented architecture (SOA) for computing resources. Users access cloud services over a simple front-end interface to utilize the virtualized resources. Ian Foster *et al.* in [1] have defined it as:

"A large-scale distributed computing paradigm that is driven by economies of scale, in which a pool of abstracted, virtualized, dynamically-scalable, managed computing power, storage, platforms, and services are delivered on demand to external customers over the Internet."

The SOA in clouds is usually defined in a hierarchical structure. The layers of cloud computing services in SOA can be described as: Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS).

IaaS providers, such as Amazon AWS¹, provide virtual CPUs, storage facilities, memory, etc. according to user requests. PaaS acts as an abstraction between the physical resources and the service. PaaS providers, such as Google App Engine², supply a software platform and the application programming interfaces (APIs), where users execute their

software components. SaaS provider, such as Salesforce.com³, provide end users with integrated services from the providers, comprising of hardware, development platforms, and applications.

The Pay-Per-Use model for cloud infra-structures has introduced wide interest among users to utilize such services. Major cloud service providers such as Amazon AWS, Rackspace⁴, Salesforce, etc. have driven development of multiple open-source cloud platforms. The most prominent among the open source cloud projects are OpenStack⁵, CloudStack⁶, Eucalyptus⁷, and OpenNebula⁸. The open-source cloud platforms provide the ability to deploy private IaaS clouds. Many open-source cloud platforms have compatible application programming interfaces (APIs) with public clouds such as Amazon AWS EC2APIs [2], which improves the flexibility and usability of the private clouds.

However, the cloud solutions available today have little flexibility in their authentication system. All of the above mentioned platforms allow user authentication, based on proprietary mechanisms, which include tokens, signatures, etc. With the recent shift in identity solutions, from being organization centric to user centric, these platforms have no provision for open authentication mechanisms, such as OpenID [3, 4]. Overcoming the existing limitations and lack of provisions, this paper presents the design for OpenID-Authentication-as-a-Service APIs in OpenStack, including the implementation of a prototype for the proposed architecture.

We chose OpenStack for our research on cloud platforms, and its architecture is discussed in section II. In section III, we provide a detailed discussion on the shortcomings of cloud platforms, specifically OpenStack. Section IV includes a brief description of the OpenID authentication mechanism, followed by section V, where we present our innovative design for implementing OpenID authentication with APIs in OpenStack. The design is applicable till to the Cactus release, which was the stable version at the time of the ongoing research. Finally, the implemented prototype for the proposed design is discussed in section VI of the paper.

³SalesForce CRM & Cloud Computing, www.salesforce.com

⁴Rackspace US, <http://www.rackspace.com>

⁵OpenStack, <http://www.openstack.org/>

⁶CloudStack, <http://cloud.com>

⁷Eucalyptus, <http://www.eucalyptus.com>

⁸OpenNebula, <http://opennebula.org>

¹Amazon Web Services (AWS), <http://aws.amazon.com>

²Google App Engine, <https://code.google.com/appengine>

II. CLOUD COMPUTING WITH OPENSTACK NOVA

Nova, the cloud computing middleware fabric controller from OpenStack, is a widely utilized open source project with many contributors. It originated as a project at NASA Ames Research Laboratory and started as open-source software in early 2010. Till the time of this research, OpenStack has released the following versions: Austin (October 2010), Bexar (February 2011), and Cactus (April 2011). The upcoming stable release of OpenStack, Diablo, is scheduled in late September 2011.

OpenStack manages computing resources: CPU, memory, disk space, and network bandwidth. The middleware application uses an hypervisor running in the back-end to allow the creation of virtual machines (VMs). These VMs emulate physical computers, and each have a CPU, memory, disk, and network resources. The actual physical resources for the creation of the VMs are provided by virtual hosts. OpenStack supports virtualization with KVM, UML, XEN, and HyperV, using the QEMU emulator. In the implementation, the *libvirt* [5] C/C++ library is used to communicate with the hypervisor from the middleware layer.

A. Architecture Overview

The components in the OpenStack architecture are: *Cloud Controller*, *API Server*, *Auth Manager*, *Nova-Manager*, *Scheduler*, *Object Store*, *Volume Controller*, *Network Controller*, and *Compute Controller*.

The Cloud Controller is the central component which represents the global state, and interacts with the other components. The Cloud Controller interacts with the API Server and the Auth Manager with internal method calls, with the Object Store over HTTP, and with the Scheduler, Network Controller and the Volume Controller, together, over the RabbitMQ [6] server using Advanced Message Queuing Protocol [7].

The API Server is an HTTP server which provides two sets of APIs to interact with the Cloud Controller: the Amazon EC2 APIs and the OpenStack OS APIs.

The Auth Manager provides authentication and authorization services for OpenStack, which can interact with the Nova-Manager client using local method calls. Nova-Manager is an admin tool to communicate with the Auth Manager to directly interact with the OpenStack database.

The Scheduler is responsible for selecting the most suitable Compute Controller to host an instance, and the Compute Controller subsequently provides the compute server resources, according to the commands from the Scheduler. The Object Store component is responsible for storage services. The Volume Controller provides permanent block-level storage for the compute servers, while the Network Controller handles the virtual networks for the VMs to interact with the public network respectively.

B. Authentication and Authorization Framework

The authentication of requests from a user, and the authorizing of resources for the request are handled by the

Auth Manager module. OpenStack uses a *Role Based Access Control* (RBAC) [8] mechanism to enforce policies.

When a user is created, an *Access Key* and a *Secret Key* are assigned to the user. They can be randomly generated or can be specified by the administrator during user creation. The credentials in the user database for each OpenStack user are shown in table I. These credentials are used in different ways to authenticate a user's incoming API requests.

TABLE I: User Credentials in OpenStack

Credential	Description
<i>id</i>	Unique identifier for each user
<i>name</i>	Usually, human readable <i>username</i> for a user
<i>access_key</i>	Unique, and can be randomly generated or specified during user creation
<i>secret_key</i>	Unique, and can be randomly generated or specified during user creation
<i>is_admin</i>	Set to "1" if an "admin" user is created, or "0" otherwise

C. Application Programming Interfaces

OpenStack Nova exposes two sets of APIs: the OpenStack API (OSAPI) and Amazon Elastic Compute Cloud API (EC2API). The OSAPI is the list of APIs being developed as OpenStack matures with time. On the other hand, the EC2 APIs are a list of comprehensive APIs, designed and defined by Amazon Web Services (AWS). In all cases, OpenStack relies on Representational State Transfer (REST) to handle the responses from the APIs.

REST [9] is an architecture for designing web applications. In typical implementations, REST relies on a stateless, client-server, cacheable communications protocol, usually over HTTPS. An example of a RESTful client-server interaction is shown in figure 1.

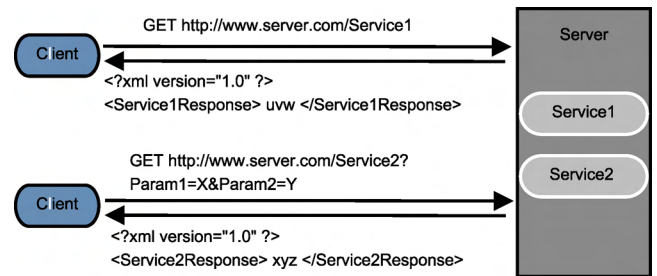


Fig. 1: RESTful Client-Server Interaction

III. PROBLEM AREA

Identity management in web services is experiencing a paradigm shift, from organization centric, to user centric authentication mechanisms. User centric identity allows both scalability, and flexibility in application to multiple service points over the Internet [10]. Additionally, these user centric frameworks aim to provide Single-Sign-On (SSO) mechanism for its users, and thus, provide a certain leverage to introduce login federation, which greatly improves the usability for any service architecture.

OpenStack services can be utilized via API tools, such as *euca-tools* EC2API client [11] and the *python-nova* OSAPI client [12]. However, a graphical interface provides better usability, especially for users without much knowledge of API tools and commands. Thus, the web GUI has become a widely deployed front-end for delivering cloud services, both to administrators and users. The Dashboard/Django-Nova framework provides a suitable GUI for users. However, there are certain limitations in what the front-end GUI for OpenStack make available in the context of authentication of users.

OpenStack encourages the use of its APIs (EC2API [2] or OSAPI [13]) for implementing front-end GUI services. OpenStack performs authentication, based on access and secret keys. The weakness in this implementation is that, the users are required to be authenticated in a separate authentication framework in the front-end, with a username/password pair (valid up to Cactus release). Once authenticated, the front-end GUI server then uses the Admin credentials to retrieve the user's credentials. This way, the backend OpenStack server never participates to the front-end user authentication. Basically, OpenStack does not support federated identity management properties that are available today in OpenID [3, 4], Shibboleth [14], SAML [15], etc [16].

In standard federated login architectures, the Policy Administration Point (PAP) and the Policy Decision Point (PDP) are logically located at a single point in the architecture. There can be multiple Policy Enforcement Points (PEPs), which communicate with the PDP [17, 18]. However, as the front-end GUI server becomes a separate security domain in OpenStack, the front-end needs to maintain separate user credentials, due to an absence of a federated login architecture. The absence of a centralized authentication architecture causes the problem of multiple PAPs and PDPs. For the above mentioned reasons, we see that there should be complete trust (single security domain) between front-end GUI and back-end cloud platform. This also means that the front-end GUI can not simply be a dumb web server, but has to be more tightly coupled to the back-end.

IV. AUTHENTICATION WITH OPENID

OpenID is a well known open source authentication mechanism. It provides decentralized user centric identity management for web services, and allows seamless SSO authentication. The current version of OpenID is 2.0 [3, 4]. Previously, OpenID 1.0 only supported stateful authentication. However, OpenID 2.0 supports both stateful and stateless OpenID authentication. The sequence of a stateless OpenID authentication is shown in figure 2.

The User-Agent (UA) first requests a page over HTTP from a web service point, and the web server returns the page to the UA. The user then submits his OpenID Identifier. The web server acts as a Relay Point (RP). It normalizes the Identifier, and performs the discovery process, using Yadis discovery protocol [19] (XRI Resolution protocol [20] was used in OpenID 1.0, but is avoided in OpenID 2.0). The RP

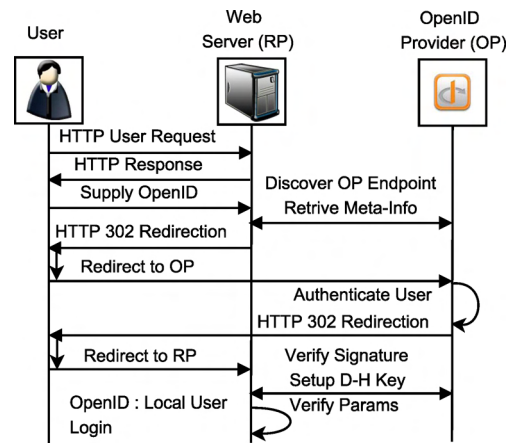


Fig. 2: Stateless OpenID Authentication Mechanism

receives the meta-information from the OpenID Provider (OP) to redirect the UA to the OP endpoint URL. The RP then sends an HTTP 302 Redirection response to the UA, with multiple parameters.

After the UA is redirected to the OP endpoint URL, the OP can use any method to authenticate the user (such as username/password, certificates, smart-cards, generic bootstrapping architecture based device authentication, etc). After authentication, the OP returns the UA back to the RP, and passes a long string in the HTTP GET request line, also called the assertion URL. The RP verifies the signature in the assertion URL, and sets up a key association using Diffie-Hellman (D-H) key exchange [21]. Then, the RP verifies the response for the specific `openid.identity`.

Once the parameters are all successfully verified, referred to as the "assertion", the RP links the `openid.claimed_id` with the identity of a local user in its own server and allows the user to login to the service point.

In a "stateful" OpenID authentication, the D-H key exchange [21] occurs in the initial discovery phase at the RP. The shared key is stored in a key database at the RP. The UA is then redirected to the OP. After authentication at the OP and when the user is redirected back to the RP, there is no D-H key exchange. Instead, the stored key from the previous step is retrieved from the database.

V. OPENID IN OPENSTACK

Implementing OpenID at the front-end GUI server as a simple RP is not the target. To apply OpenID authentication mechanism in OpenStack, we needed to combine a dual-PDP scenario into a single-silo formation. The following sections discuss the design considerations, and the solution architecture, followed by the usability and a detailed analysis of our solution. The design is applicable to OpenStack, till the Cactus release, which was the stable version at the time of the ongoing research.

A. Design Considerations

In an architecture to integrate OpenID with OpenStack, the front-end GUI should be a "dumb" server, only processing the

views for the user. There should not be any requirement for the GUI server to maintain any user credentials for authentication.

Furthermore, even though the views on the GUI are based on responses from the API server, the initial authentication on the front-end should be granted by the **Auth Manager** in the back-end server. However, as the HTTP User-Agent only interacts with the front-end server, the back-end OpenStack server should not have any direct communication with the User-Agent.

Therefore, the process of authentication of a user in the front-end should be realized as a service from the back-end server. Thus, we converged on a solution based upon OpenID-Authentication-as-a-Service for OpenStack. However, all interaction between the front-end and the back-end should be stateless, as required by the RESTful API server [9, 13].

Additionally, the design should meet all of the specifications of OpenID [3, 4], and ensure all security requirements for OpenID in all phases of interaction. This would allow interoperability between all OpenID providers, thus greatly improving the usability for the users.

Finally, the requirements at the front-end server to implement OpenID authentication in OpenStack should be simple, and secured. Also, we need to maintain a modular and distributed structure to comply with the current architecture and scalability of OpenStack.

B. Implementing OpenID as a Service

OpenID authentication at an RP involves two phases: (a) OP endpoint URL discovery and retrieving meta-information, and (b) Verifying an authentication assertion URL received from an OP. Therefore, we divided the OpenID-Authentication-as-a-Service operation in OpenStack into two phases, each invoked with a separate API. The functions of the two APIs have been defined as:

- **Authentication Request API:** This API is invoked in the initial phase by the front-end server. It executes an OpenID authentication request, and performs the first phase in the process.
- **Authentication Verification API:** This API is invoked in the second phase by the front-end server. It executes the authentication verification for the OpenID authentication assertion URL received from the OP.

As shown in figure 3, the user requests for an OpenID based authentication to the front-end GUI server. The GUI server then invokes the initial authentication request API on the API Server in the back end. The back-end server responds to the request with all the necessary OpenID parameters required for the redirection. The GUI server parses the information, and sends a HTTP 302 Redirection to the UA.

The UA redirects to the OP, where it is authenticated. Upon successful authentication, the OP sends the UA to the front-end GUI server.

At this point, the front-end invokes the second authentication verification API on the back-end API server. The back-end communicates with the OP, and completes all the processes for verification. Once verified, the authentication is granted by

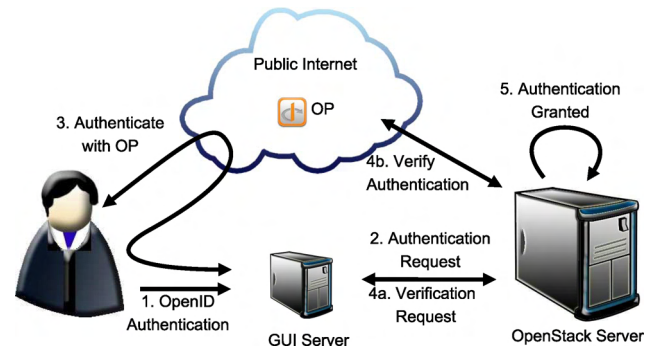


Fig. 3: OpenID Authentication in OpenStack

Auth Manager. Based on the authenticated user information, the front-end then allows the user to log into the managerial interface.

C. Design Analysis

HTTPS can be used to secure interaction with the UA. Additionally, it is required to protect the integrity of the assertion messages relayed from the OP through the UA, and includes nonce checking, and signature verification.

All RESTful requests to the API Server include signatures with a pre-shared secret between the GUI server and OpenStack. Thus, unless the front-end server is vulnerable to a compromise by an attacker, the connection to the back-end can be considered as an integrity protected channel. Using SSL between the front-end and the API Server is a common practice for confidentiality in RESTful services. Security technologies such as IPSec [22] are intended for network level host-to-host security, rather than application-to-application security, and hence is not a recommended security solution for the RESTful API Server. However, HTTPS support in the OpenStack API Server has not been implemented yet, and remains as a future task.

The verification of the assertion URL by OpenStack server and the OP occurs in the back-end. However, the back-end communication with the OP cannot be considered hidden from an attacker. An attacker can sniff packets from the network to intercept the communication between OpenStack and the OP. Hence, this communication takes place over an encrypted channel with the D-H shared key between OpenStack and the OP.

However, there is still scope for an attacker to manipulate the information. Based on security issues of OpenID, the solution can be vulnerable to Discovery Tampering, Adversary Relay Proxy, and DoS attacks.

Session management between the UA and the GUI front-end is another area where the security should be improved. Services on OpenStack are RESTful services, and no session information is stored, while the front-end is a session-based service point for the UA. It is contradictory with the design principles of RESTful services to maintain such session based security. Therefore, OpenStack needs to trust the front-end GUI to manage the user session.

D. Usability of OpenID in OpenStack

A lot of discussions on the usability of OpenID have occurred so far. OpenID is the most widely used open standard for authentication, with many OP providers, such as Google⁹, Yahoo¹⁰, MyOpenID¹¹, and LiveJournal¹².

Integrating OpenID in OpenStack will provide a decentralized user centric authentication delegation for using OpenStack services, where the users will have control over his or her own identity management and authentication. Thus, usability will improve, as OpenID aims for a single user versus multiple service points applicability, and allows users to have a seamless SSO experience. This would allow providers to introduce a federated login architecture, and also reduce the IT maintenance cost by management of user credentials at 3rd party OPs.

Standard OpenID implementation also includes using the Provider Authentication Policy Extension (PAPE), to allow a flexible authentication framework. PAPE allows an RP to specify different requirements to be implemented at an OP during authentication. Thus, OpenID could utilize PAPE to implement a requirement based security in OpenStack.

Additionally, users will have flexibility in the authentication mechanism, as OPs allow different authentication mechanisms for their users. Most OPs support username/password, and client certificate based authentication for users. Apart from that, Leicher et al. in [23] describe a trusted computing environment using OpenID, A. S. Ahmed in [24] presents a 3GPP standard authentication mechanism for smart phones, and Watanabe *et al.* in [25] illustrate a cellular subscriber ID and OpenID federated authentication architecture. Ericsson Labs also provides an Identity Management service, which uses the 3GPP standard Generic Bootstrapping Architecture [26, 27] based device authentication services with OpenID.

VI. PROTOTYPE IMPLEMENTATION

We initially started working with the Bexar release. After a successful implementation with Bexar, we then integrated our solution with Cactus, the third release, which came out in April 2011.

In our design, we introduced an additional module, the Nova-OpenID Controller, in the OpenStack architecture. The added module is responsible for all operations related to OpenID authentication in the back-end, and has an internal HTTP interface with the Cloud Controller, and a public interface to interact with the OP on the public internet. The implementation also included extension of the Nova-Admin tool. The admin can use the added functionality in Nova-Admin to add/modify OpenID information for existing OpenStack users. The architecture of OpenStack, including the added modules for the prototype is shown in figure 4.

We designed the two APIs according to the specification of EC2 APIs on OpenStack API server. Furthermore,

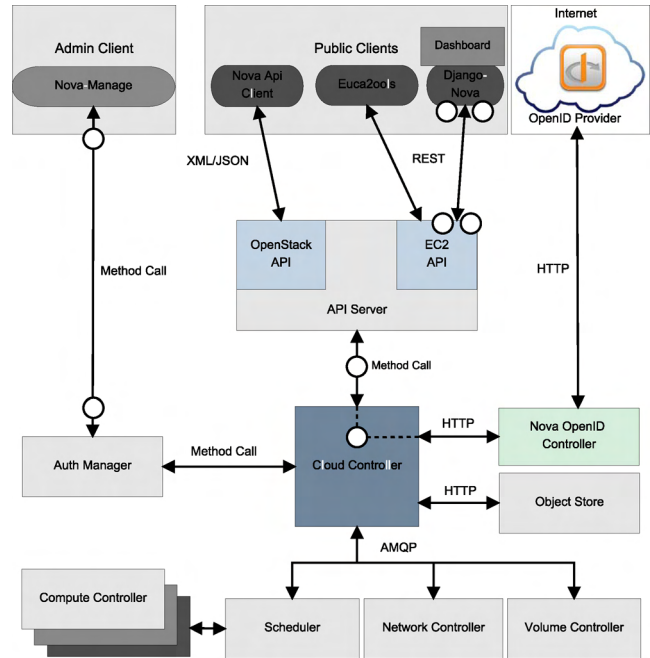


Fig. 4: Prototype Architecture for OpenID in OpenStack

we implemented the invocation of the APIs from the Dashboard/Django-Nova web GUI for OpenStack. In our current implementation, we have a one-to-one mapping of existing OpenStack users to the number of enabled OpenID Identifiers a user can use. The implementation was tested successfully against standard OpenID providers on the Internet.

As shown in figure 5, the user provides the OpenID Identifier on the GUI, and subsequently, the front-end invokes the *OpenidAuthReq* API. The response contains the required parameters for redirecting the UA to the OP. After the user authentication is completed, the UA returns to the front-end, and invokes the *OpenidAuthVerify* API. The back-end then verifies the assertion URL, links an existing OpenStack user to the verified OpenID Identifier, and returns a success or a failure. The front-end GUI then uses the information to allow or deny login to the user.

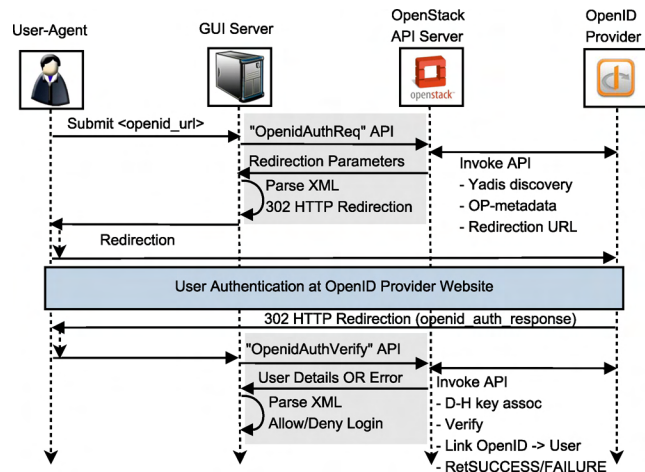


Fig. 5: Signalling Sequence for OpenID Authentication in OpenStack

⁹Google OpenID Services, <http://code.google.com/apis/identitytoolkit/>

¹⁰Yahoo OpenID Services, <http://openid.yahoo.com>

¹¹MyOpenID OpenID Services, <https://www.myopenid.com>

¹²LiveJournal OpenID Services, <http://www.livejournal.com>

Diablo, the fourth release of OpenStack is scheduled to be released in late September 2011. However, beginning with the Diablo release, the authentication framework design is supposed to utilize a new architecture. It will integrate the *KeyStone* project [28] as the authentication module, the design of which was still evolving at the time of writing this paper.

VII. CONCLUSION AND FUTURE WORKS

To improve usability, OpenStack proposes to utilize its API functions to provide a GUI for its users. However, the current architecture of OpenStack lacks specific GUI based services. For initial authentication, the front-end is required to incorporate a separate username/password validation. This introduces a dual PDP scenario, which is not a recommended practice for web services.

In this paper, we introduced a flexible decentralized authentication service for the front-end, using OpenID as an open-source authentication platform. In our design, OpenID authentication in the front-end is used as a service from the back-end OpenStack server. As a result, we were able to shift the dual points of decision making and perform the authentication at a single PDP in the back-end. The design was successfully implemented on OpenStack, and utilized the OpenID-Authentication-as-a-Service APIs from the Dashboard/Django-Nova GUI.

The research performed during this work revealed further possibilities. The first objective would be to introduce greater flexibility in the choice of authentication mechanisms for the user. To provide a generic solution for authentication, we aim to design a common Authentication-as-a-Service API in OpenStack. Additionally, we also aim to introduce open platforms for authorization delegation within OpenStack.

REFERENCES

- [1] Foster, I. and Yong Zhao and Raicu, I. and Lu, S., "Cloud Computing and Grid Computing 360-Degree Compared," in *Grid Computing Environments Workshop, 2008. GCE '08*, November 2008, pp. 1 – 10.
- [2] "Amazon AWS EC2 API Reference, <http://docs.amazonweb-services.com/awsec2/latest/apireference/>, last accessed 30th April 2011."
- [3] "OpenID Foundation, <http://openid.net>, last accessed 14th June 2011."
- [4] D. Recordon and D. Reed, "Openid 2.0: a platform for user-centric identity management," in *Proceedings of the second ACM workshop on Digital identity management*, ser. DIM '06. New York, NY, USA: ACM, 2006, pp. 11–16. [Online]. Available: <http://doi.acm.org/10.1145/1179529.1179532>
- [5] "Libvirt Online Working Group, The Virtualization API, <http://libvirt.org>, last accessed 25th April 2011."
- [6] RabbitMQ AMQP Server Working Group, "<http://www.rabbitmq.com/specification.html>, last accessed 30th June, 2011."
- [7] Advanced Message Queuing Protocol (AMQP) Working Group, "<http://www.amqp.org/>, last accessed 15th June, 2011."
- [8] R. Sandhu, E. Coyne, H. Feinstein, and C. Youman, "Role-based access control models," *Computer*, vol. 29, no. 2, pp. 38 –47, feb 1996.
- [9] R. T. Fielding, "Architectural styles and the design of network-based software architectures," Ph.D. dissertation, University of California, Irvine, 2000.
- [10] Jamie Bodley-Scott, "IDM09, Access or Identity, http://www.opengroup.org/jericho/idm2009_jbs.pdf."
- [11] "Euca-Tools, Eucalyptus Community, <http://open.eucalyptus.com/wiki/toolsecosystem>, last accessed 10th May 2011."
- [12] "Python-Openid 2.2.5, <http://pypi.python.org/pypi/python-openid>, last accessed 5th May 2011."
- [13] Rackspace US, Inc., "Openstack compute developer guide api 1.0, api v1.0, january 2011."
- [14] M. Erdos and S. Cantor, "Shibboleth architecture protocols and profiles, <http://shibboleth.internet2.edu/shibboleth-documents.html>."
- [15] R. Philpott, E. Maler, N. Ragouzis, J. Hughes, P. Madsen, and T. Scavo, "OASIS Open 2008, Security Assertion Markup Language (SAML) V2.0 Technical Overview, Committee Draft 02, <http://docs.oasis-open.org/security/saml/post2.0/sstc-saml-tech-overview-2.0.html>," March 2008.
- [16] J. Rosenberg and D. Remy, *Securing Web Services with WS-Security: Demystifying WS-Security, WS-Policy, SAML, XML Signature, and XML Encryption*. Pearson Higher Education, 2004.
- [17] S. Almulla and C. Y. Yeun, "Cloud computing security management," in *Engineering Systems Management and Its Applications (ICESMA), 2010 Second International Conference on*, 30 2010-april 1 2010, pp. 1 –7.
- [18] D. Gollmann, "Computer security," *Wiley Interdisciplinary Reviews: Computational Statistics*, vol. 2, no. 5, pp. 544–554, 2010. [Online]. Available: <http://dx.doi.org/10.1002/wics.106>
- [19] "Yadis 1.0, The Identity and Accountability Foundation for Web 2.0, <http://yadis.org>, last accessed 5th June 2011."
- [20] G. Wachob, D. Reed, L. Chasen, W. Tan, and S. Churchill, "Extensible resource identifier (xri) resolution v2.0, <http://docs.oasis-open.org/xri/2.0/specs/xri-resolution-v2.0.pdf>."
- [21] W. Diffie and M. Hellman, "New directions in cryptography," in *IEEE Transactions on Information Theory*, vol. 22, no. 6, nov 1976, pp. 644 – 654.
- [22] N. Doraswamy and D. Harkins, *IPSEC: The New Security Standard for the Internet, Intranets, and Virtual Private Networks*. Prentice Hall, 1999.
- [23] A. Leicher, A. Schmidt, Y. Shah, and C. Inhyok, "Trusted computing enhanced openid," in *Internet Technology and Secured Transactions (ICITST), 2010 International Conference for*, nov. 2010, pp. 1 –8.
- [24] A. S. Ahmed, "A user friendly and secure openid solution for smart phone platforms," Master's thesis, Faculty of Information and Natural Sciences, School of Science and Technology, Aalto University, June 2010.
- [25] R. Watanabe and T. Tanaka, "Federated authentication mechanism using cellular phone - collaboration with openid," in *Information Technology: New Generations, 2009. ITNG '09. Sixth International Conference on*, april 2009, pp. 435 –442.
- [26] 3rd Generation Partnership Project. 3GPP TR 33.924, "Identity management and 3gpp security interworking; identity management and generic authentication architecture (gaa) interworking, (release 9). <http://www.3gpp.org/ftp/specs/html-info/33924.htm>, 2009."
- [27] Ericsson Labs Identity Management Framework, "<https://labs.ericsson.com/developer-community/blog/identity-management-framework-now-available-download>, last accessed 15th May 2011."
- [28] OpenStack Keystone Project, "<http://wiki.openstack.org/openstack-authn>, last accessed 25th May, 2011."