

OpenStack Deployment on Multi-Controllers Using DevStack and Integration of Keystone with Centralized LDAP Server

¹Sahana S J, ²Dr.M.N.Jayaram

¹Post Graduate Student, ² Professor, EC Department, SJCE, Mysore, India

Abstract: Cloud computing has proven to be a successful distributed computing model as demonstrated by its wide spread industrial adoption. OpenStack is free open source Cloud computing software originally released by Rackspace and NASA, which strives to close the gap in the lack of a comprehensive Cloud platform with a fast pace of development and innovation, and supported by both an active community of people and large companies. In today's Security Mechanism Authentication plays a vital role. In this paper, we go through OpenStack deployment on Multi-Node/Controller and integration of OpenStack project called Keystone with Centralized LDAP Server. Keystone is responsible for providing a common Authentication/Authorization for all OpenStack services.

Keywords: Keystone, Authentication, Centralized LDAP, Multi- Controller, Dashboard.

I. INTRODUCTION

In today's world, Internet has been a driving force for various technologies that have been developed. One of the most used technology is cloud computing. Cloud computing using Internet-enabled services to operate the application software. OpenStack began in 2010 as a joint project of RACKSPACE hosting and NASA. It provides both large and small organizations an alternative to closed cloud environments, reducing the risk of lock-in associated with proprietary platforms. The openStack community collaborates around a six-month, time-based release cycle with frequent development milestones. OpenStack technology consists of a series of interrelated projects that control pools of "processing", "storage", and "networking" resources throughout a data centre which users manage through a web-based dashboard, command-line tools, or a Restful API. OpenStack allows the users to deploy Virtual Machines (VMs) and other instances which handle different tasks for managing a cloud environment on the fly. It makes horizontal scaling easy, which means that tasks which benefit from running concurrently can easily serve more or less users on the fly by just spinning up more instances. Cloud based networking and applications are presently one of the Cutting-edge technology and due to its potential use in cloud computing. Multi-node and Multi-region cloud has requirement for multi-vendor OpenStack distribution, multi-OpenStack instance, multi- OpenStack version co-existence Multi-vendor: anti-vendor lock in business policy. The open source OpenStack project provides an Infrastructure as a Service (IAAS) layer for building public and private clouds. Corporations, service providers, value-added resellers, small and mid-sized businesses, researchers, and global data centres all use OpenStack to deploy large-scale private or public clouds. OpenStack have their users stored in an existing centralized authentication service. This is typically an LDAP server, or Active Directory. Keystone has an LDAP driver for the identity backend to allow it to use LDAP for authentication and storage of users and groups.

II. RELATED WORK

Generally there are three major ways to deploy an OpenStack cloud includes Manual deployment procedure, DevStack and Pack Stack. DevStack has evolved to support a large number of configuration options and alternative platforms and support services. DevStack is a opinionated script/tool that was initially developed to speed the deployment of OpenStack

for development purposes, hence “Dev” Stack. The Multi-node installation setup runs different OpenStack services on different nodes. A basic installation requires the three nodes they are Controller-node, Network-node and Compute-node.

- 1) Controller Node: It runs control services, such as message queue, database and API services for the Identity Service (Keystone).
- 2) Network Node: It runs networking services and is responsible for virtual networking needed for people to create private or public networks, and uplink their virtual machines into external networks.
- 3) Compute Node: It run the virtual machine instances in OpenStack.

Authentication plays a vital role in today’s world. Every multiuser service needs some mechanism to manage who can access the application and which actions each person can perform. A private cloud is no exception and OpenStack has streamlined these functions into a separate project called Keystone. Authentication plugins are used and implemented should be generic enough to cover completely customized authentication solutions in cloud environment with the OpenStack Identity Service (Keystone). In OpenStack cloud services, authentication is a process which positively verifies the user identity by validating a set of credentials supplied by the user. These credentials are initially a user name and password or a user name and API key, and authentication token will be issued to the user in response to these credentials. Many enterprise applications use LDAP as the foundation for user authentication. The LDAP- Lightweight Directory Access Protocol is a client/server protocol for accessing and managing directory information specifically X.500-based directory services. Keystone is the only one in OpenStack to responsible for creation of users, their roles, and to assign which project(s) they belong to. OpenStack have their users stored in an existing centralized LDAP server, or Active Directory (which offers LDAP capabilities). The main challenging part here is to achieve integration of Keystone with an existing identity store such as an LDAP server.

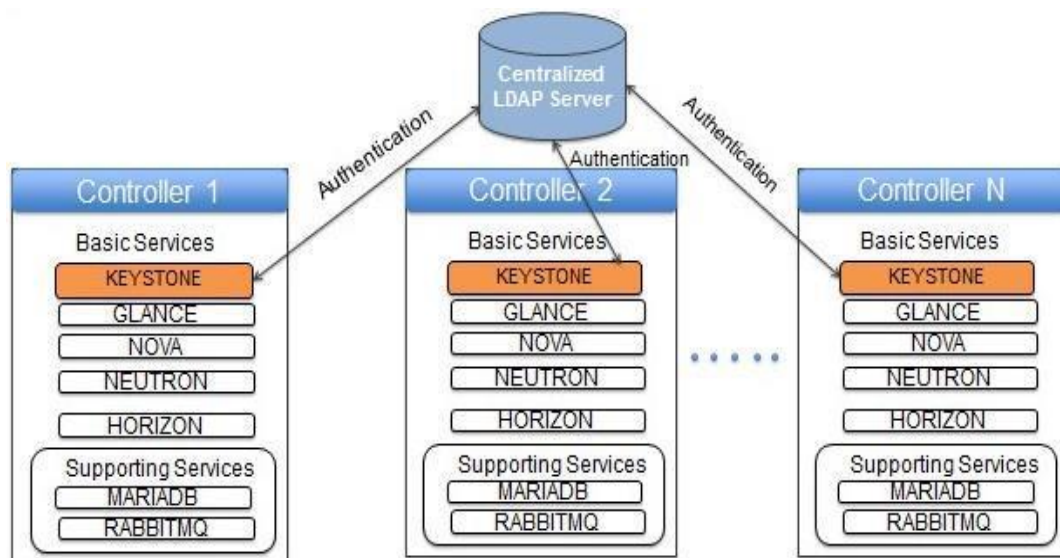


Fig.1: OpenStack deployed on Multi-Controllers with Centralized LDAP Topology

Keystone has an LDAP driver for the identity backend to allow it to use LDAP for authentication and storage of users and groups. Integrating an identity back end with Centralized LDAP, by enabling the LDAP identity driver in the “keystone.conf” file by replacing “driver = sql” with “driver = ldap”.

III. USER IDENTITY VALIDATION VIA KESTONE WITH OPENSTACK DASHBOARD

Horizon is a OpenStack’s web-based self-service portal and it sits on-top of all of the other OpenStack components via API interaction. The Fig.2 depicts the OpenStack cloud is deployed on a Node/Controller using DevStack. When DevStack “./stack.sh” script runs successfully, it will be provided with the specific IP address, username and password that can be used to access the Horizon (dashboard). According to our project OpenStack DevStack is successfully deployed on Multi-Nodes/Controllers and the OpenStack server is in “UP” state.

```
Horizon is now available at http://192.168.122.196/
Keystone is serving at http://192.168.122.196:5000/v2.0/
Examples on using novaclient command line is in exercise.sh
The default users are: admin and demo
The password: onecloud
```

Fig.2: DevStack-OpenStack deployment

The Fig.3 depicts the dashboard (horizon) it is accessed through the provided OpenStack Server IP and login credentials.



Fig.3: OpenStack Login Screen

a) Creating new user on LDAP server through keystone via CLI or GUI:

There are some commands to follow up to create new user, tenant and roles in OpenStack. Once the new OpenStack user is created and stored in database, we can check the newly created Openstack User by issuing the “*keystone user-list*” command on CLI. The Fig.4 depicts newly created OpenStack user with name “onecloud”.

id	name	enabled	email
c7d3ad78f3f841e5be190b01ce12f5b2	admin		
436ca551705f480ebf30422d71cb986d	cinder		
427c6741e2194b2db79edbceel1dc4768	glance		
d9d5595776fc4840b95264639906105d	neutron		
6ac2f7a16ff249c3bfa5b3d4b9b95037	nova		
fb0f19a89f1a4beab1b4246e4e3d7edd	onecloud		

Fig.4: Created new OpenStack User

To verify that the existing domain entry is present in LDAP Server with ldapsearch operation through CLI by using “*ldapsearch -x -b 'dc=openstack,dc=org' '(sn=onecloud)’*” command on CLI. The Fig.5 depicts domain entry stored in LDAP Server.

```
# extended LDIF
# LDAPv3
# base <dc=openstack,dc=org> with scope subtree
# filter: (sn=onecloud)
# requesting: ALL
#
# fb0f19a89f1a4beab1b4246e4e3d7edd, Users, openstack.org
dn: cn=fb0f19a89f1a4beab1b4246e4e3d7edd,ou=Users,dc=openstack,dc=org
objectClass: person
objectClass: inetOrgPerson
cn: fb0f19a89f1a4beab1b4246e4e3d7edd
sn: onecloud
userPassword:: b25lY2xvdWQxMjM=
# search result
search: 2
result: 0 Success
# numResponses: 2
# numEntries: 1
```

Fig.5: Newly created OpenStack User(onecloud) successfully stored in LDAP Server

b) Validate Keystone against Centralized LDAP/AD:

To validate OpenStack user against Centralized LDAP through Horizon by login with newly created User name (onecloud) and Password. The Fig.6 depicts OpenStack login screen with newly created username and password.

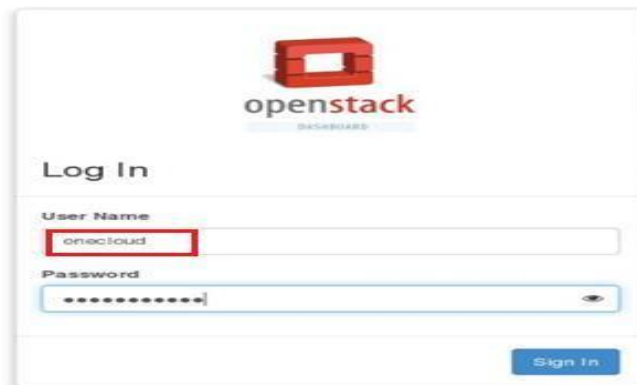
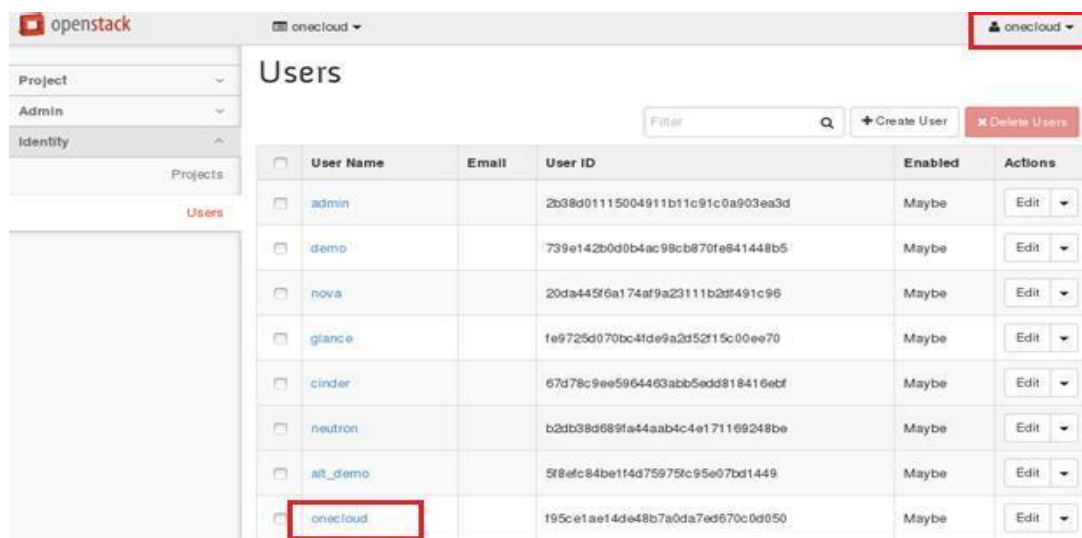


Fig.6: OpenStack Login Screen with new username and password

Through GUI/Dashboard, In Identity section, check users tab to see the newly added user in the database. The Fig.7 depicts Keystone user list through Dashboard.



<input type="checkbox"/>	User Name	Email	User ID	Enabled	Actions
<input type="checkbox"/>	admin		2b38d01115004911b11c91c0a903ea3d	Maybe	Edit
<input type="checkbox"/>	demo		739e142b0d0b4ac98cb870fe841448b5	Maybe	Edit
<input type="checkbox"/>	nova		20da445f6a174af9a23111b2df491c96	Maybe	Edit
<input type="checkbox"/>	glance		fe9725d070bc4fde9a2d5215c00ee70	Maybe	Edit
<input type="checkbox"/>	cinder		67d78c9ee5964463abb5edd818416ebf	Maybe	Edit
<input type="checkbox"/>	neutron		b2db38d689fa44aab4c4e171169248be	Maybe	Edit
<input type="checkbox"/>	alt_demo		5f8efc84be1f4d75975c95e07bd1449	Maybe	Edit
<input type="checkbox"/>	onecloud		195ce1ae14de48b7a0da7ed670c0d050	Maybe	Edit

Fig.7: Keystone users list

IV. ADVANTAGES OF CENTRALIZED LDAP IN CLOUD

- i. Centralized LDAP supports high availability and redundancy
- ii. Centralized password security policies in one authority
- iii. Centralized identity and passwords across both UNIX and Windows
- iv. Simplified creation and deletion of users

V. LIMITATIONS OF CENTRALIZED LDAP

Since Centralized LDAP Server act as global service and serving large amount of user management data to “N” number of OpenStack servers(Controllers) on different regions. Most high availability systems will fail in the event of single or multiple independent (non-consequential) failures. A crucial aspect of high availability is the elimination of single points of failure (SPOFs). In this case, most systems will protect data over maintaining availability.

VI. SOLUTION TO OVERCOME OF LIMITATIONS

Load balancing is a staple solution in virtually every data center. A load balancer is a device that acts as a reverse proxy and distributes network or application traffic across a number of servers. Load balancing aims to optimize resource use, maximize throughput, minimize response time, and avoid overload of any single resource. Using multiple components with load balancing instead of a single component may increase reliability and availability through redundancy.

Failover support: When a major failure occurs in the Centralized LDAP Directory Server specified by the primary URL and the server no longer responds to network requests, LDAP clients attempt to connect to the secondary LDAP Directory Server specified by the secondary URL. The Fig.8 depicts virtual load balancing in cloud environment.




Fig.8: Virtual Load Balancing

VII. CONCLUSION

Virtualization can exist without the cloud, but cloud computing cannot exist without virtualization. LDAP is rapidly becoming a defacto standard for remote authentication and authorization of users. Centralized LDAP is a solution to access centrally stored information over network. This centrally stored information is organized in a directory that follows X.500 standard. The main advantage of this approach is that the information can be grouped into containers and clients can access these containers whenever needed. In this paper, Centralized LDAP approach enables to achieve High availability and redundancy.

REFERENCES

- [1] "Engineering Enterprise Applications To Ensure The Highest Level Of Availability And Fault Tolerance In The Cloud", by brian jimerson, published june 2012.
- [2] Stratus Technologies, "Server Virtualization and Cloud Computing: Four Hidden Impacts on Uptime and Availability," A White Paper by Stratus Technologies, June 2013.
- [3] Oracle, "Architectural Strategies for Cloud Computing," An Oracle White Paper in Enterprise Architecture, August 2009.
- [4] Juniper Networks, "Implementation Identity Federation in a Hybrid Cloud Computing Environment Solution Guide," October 2009.
- [5] "A User Identity Management Protocol for Cloud Computing Paradigm", by Lawrence Kehinde ,Vol.4 No.3, March 2011.
- [6] Cloud Standards Customer Council, "Security for Cloud Computing 10 Steps to Ensure Success", June 2012.
- [7] Energy Efficiency for Data Center and Cloud Computing: A Literature Review, Volume 3, Issue 4, October 2013.
- [8] Private Virtual Infrastructure: A Model for Trustworthy Utility Cloud Computing UMBC Computer Science Technical Report Number TR-CS-10-04.
- [9] The Security Division of EMC, "The Role of Security in Trustworthy Cloud Computing", RSA white paper, April 2009.