

Storage Area Networks with OPEN STACK and Installation of Havana

Vamsi Krishna Pelluru

Abstract— *In Traditional Enterprise-class storage technology, many organizations now have a variety of storage needs with varying performance and price requirements. Open stack has support for both Object and Block storages with many deploying options based on the organizations needs.*

Index Terms—Open Stack, SAN, Storage, havana

I. INTRODUCTION

There is need to develop, understand and implement Storage area network in Open Stack, this paper enables the reader to Understand objects and Block storages, its advantages and how it can be installed with Havana.

II. OBJECT STORAGE CAPABILITIES

- ❖ Open Stack provides redundant, scalable object storage using clusters of standardized servers capable of storing peta bytes of data
- ❖ Object Storage is not a traditional file system, but rather a distributed storage system for static data such as virtual machine images, photo storage, email storage, backups and archives. Having no central "brain" or master point of control provides greater scalability, redundancy and durability.
- ❖ Objects and files are written to multiple disk drives spread throughout servers in the data center, with the Open Stack software responsible for ensuring data replication and integrity across the cluster.
- ❖ Storage clusters scale horizontally simply by adding new servers. Should a server or hard drive fail, Open Stack replicates its content from other active nodes to new locations in the cluster. Because Open Stack uses software logic to ensure data replication and distribution across different devices, inexpensive commodity hard drives and servers can be used in lieu of more expensive equipment.

III. BLOCK STORAGE CAPABILITIES

- ❖ Open Stack provides persistent block level storage devices for use with Open Stack compute instances.
- ❖ The block storage system manages the creation, attaching and detaching of the block devices to servers. Block storage volumes are fully integrated into Open Stack Compute and the Dashboard allowing for cloud users to manage their own storage needs.
- ❖ In addition to using simple Linux server storage, it has unified storage support for numerous storage platforms including Ceph, NetApp, Nexenta, SolidFire, and Zadara.

- ❖ Block storage is appropriate for performance sensitive scenarios such as database storage, expandable file systems, or providing a server with access to raw block level storage.
- ❖ Snapshot management provides powerful functionality for backing up data stored on block storage volumes. Snapshots can be restored or used to create a new block storage volume.

IV. USING SAN OBJECT AND BLOCK STORAGE CAN POTENTIALLY OFFER THE FOLLOWING BENEFITS

- ❑ Improvements to application availability: Storage is independent of applications and accessible through multiple data paths for better reliability, availability, and serviceability.
- ❑ Higher application performance: Storage processing is off-loaded from servers and moved onto a separate network.
- ❑ Centralized and consolidated storage: Simpler management, scalability, flexibility, and availability.
- ❑ Data transfer and vaulting to remote sites: Remote copy of data enabled for disaster protection and against malicious attacks.
- ❑ Centralized management: Single image of storage media simplifies management.

V. SAN- HOW BLOCK STORAGE WORKS WITH ISCSI PROTOCOL

- ❖ iSCSI is an emerging standard which defines the encapsulation of SCSI packets in TCP. These packets are then routed using IP. This technology allows block-level storage data to be transported over widely used IP networks, enabling end users to access the storage network from anywhere in the enterprise.
- ❖ Several iSCSI products are now available, including iSCSI HBAs and iSCSI-to-Fiber Channel storage routers. Storage routers, when utilized with a Fiber Channel switch, will enable block-level storage traffic to be accessed from a SAN over an IP network.
- ❖ The iSCSI concept can be expanded to create a storage network by using any network-connected device. iSCSI enables HBAs to provide I/O processing on order to reduce the load on the server CPU.

VI. OPEN STACK INSTALLATION PROCEDURE

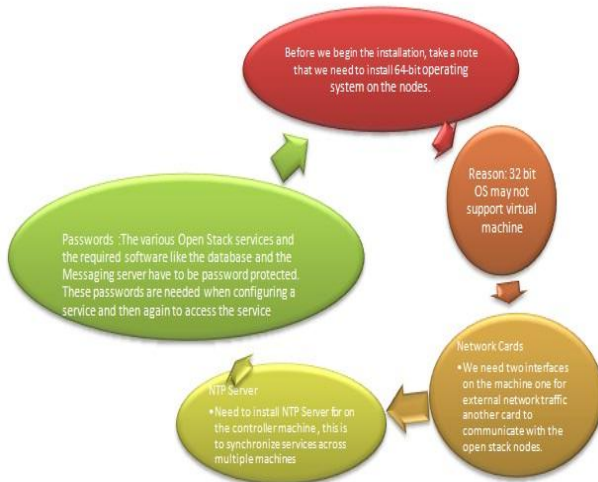
Open Stack® system consists of several key projects that you install separately but that work together depending on your cloud needs. These projects include Compute, Identity Service, Networking, Image Service, Block Storage Service, Object Storage, Telemetry, and Orchestration. You can install any of these projects separately and configure them standalone or as connected entities. This guide walks through an installation by using packages available through Ubuntu

Manuscript received March, 2014.

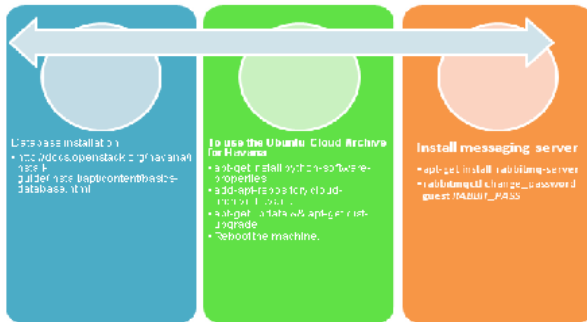
Vamsi Krishna Pelluru, System Validation team, Freescale semiconductor, Hyderabad, India.

12.04 (LTS). Explanations of configuration options and sample configuration files are included.

a. Basic requirements for installation procedures



b. Installation using Havana:



Open Stack Havana, an open-source Infrastructure-as-a-Service cloud, has added features that make it even better suited for enterprise use. Businesses have taken to Open Stack, the open-source, Infrastructure-as-a-Service (IaaS) cloud, like ducks to water. Major companies such as PayPal, Intel, and Wikimedia, the group behind Wikipedia, are using Open Stack in their clouds. Open Stack's Havana, the eighth release of this IaaS for public, private, and hybrid clouds, brings its enterprise customers several new improvements.

c. Configure Identity services

Configuring Identity services

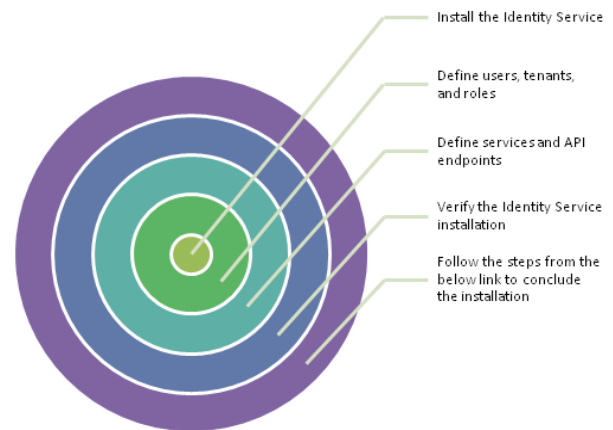
- Identity Service concepts
- Install the Identity Service
- Define users, tenants, and roles
- Define services and API endpoints
- Verify the Identity Service installation
- Use the link below for reference
- http://docs.openstack.org/havana/Install-guide/Install/apt/content/ch_keystone.html

d. configure Image services

- ❖ The Open Stack Image Service enables users to discover, register, and retrieve virtual machine images. Also known as the glance project, the Image Service offers a REST API that enables you to query virtual machine image metadata and retrieve an actual image. Virtual machine images made available through the Image Service can be stored in a variety of locations from simple file systems to object-storage systems like Open Stack Object Storage.
- ❑ Image Service overview
- ❑ Install the Image Service
- ❑ Verify the Image Service installation

e. configure compute services

- ❖ The Compute service is a cloud computing fabric controller, which is the main part of an IaaS system. Use it to host and manage cloud computing systems. The main modules are implemented in Python.
- ❖ Compute interacts with the Identity Service for authentication, Image Service for images, and the Dashboard for the user and administrative interface. Access to images is limited by project and by user; quotas are limited per project (for example, the number of instances). The Compute service scales horizontally on standard hardware, and downloads images to launch instances as required.



f. Open Stack Dashboard

Horizon is the canonical implementation of Open Stack's Dashboard, which provides a web based user interface to Open Stack services including Nova, Swift, Keystone, etc.



VII. CONCLUSION

Providing storage solutions with Open Stack is a Global requirement at this point of time. The Need for this solution is becoming more and more Critical with major enterprise level companies looking to move towards cloud solutions. Small to medium Companies cannot afford high end solutions, so Open Stack with SAN becomes an inevitable option.

ACKNOWLEDGMENT

The preferred spelling of the word “acknowledgment” in American English is without an “e” after the “g.” Use the singular heading even if you have many acknowledgments. Avoid expressions such as “One of us (S.B.A.) would like to thank” Instead, write “F. A. Author thanks” **Sponsor and financial support acknowledgments are placed in the unnumbered footnote on the first page.**

REFERENCES

- [1] <http://docs.openstack.org/havana/install-guide/install/apt/content/>
- [2] http://docs.openstack.org/havana/install-guide/install/apt/content/ch_kystone.html



Vamsi Krishna Pelluru Completed Masters in Computer Applications from Kakatiya University, working as Senior Quality Assurance Engineer in Freescale Semiconductor for the past Eight years, his expertise lies across network security, Open Stack and Storage Area Networks.