

IAAS PRIVATE CLOUD COMPARISON

OPENSTACK VS CLOUDSTACK

IONEL GORDIN

“Stefan cel Mare” University of Suceava, Romania

ionel@usv.ro

Keywords: cloud computing, OpenStack, CloudStack, IAAS

Abstract: Many organizations are now investing in cloud computing because they have realized it has the ability to promote rapid growth while at the same time reducing the speed and costs of application deployment. Enterprises no longer need to maintain computing resources that are used periodically and left idle most of the time. Choosing the best cloud solution for our organization can be challenging. A cloud solution requires high processing power and a very large storage space. Efficiency, easiness of maintenance and configuration are the key components when choosing our cloud solution. For open source Infrastructure as a service (IaaS) there are two key players: OpenStack and CloudStack. Both solutions are open source software platforms for IaaS that offer cloud orchestration architectures used to make the management of cloud computing easier and more efficient. This article compares each component of these platforms and offers clear solutions adapted for our infrastructure. Although there are many articles taking in consideration this thematic, the article takes a new approach by making a comparative analysis on for each similar component showing advantages and disadvantages over the other cloud solution.

1. INTRODUCTION

OPENSTACK

OpenStack is an open source IaaS platform for overseeing and making vast gatherings of virtual private servers in a distributed computing environment [1]. It was at first created by Rackspace and NASA in 2010. With more than 500 organizations embracing this stage, it is unquestionably a standout amongst the most famous cloud models out there. This cloud solution principle objective is to bolster interoperability between cloud administrations while empowering endeavors to make cloud administrations inside of their own server farms.

OpenStack APIs are compatible with Amazon EC2 and Amazon S3 and client applications that are composed for Amazon Web Services can be utilized with OpenStack with insignificant porting exertion.

The conceptual structure of Openstack can be seen on Fig. 1

The Openstack version of 2015 (codename Liberty) and their assigned code names are shown on Table I.

The system consists of several independent parts, named the OpenStack services.

All services authenticate through a common Identity service. Individual services interact with each other through public APIs, except where privileged administrator commands are necessary.

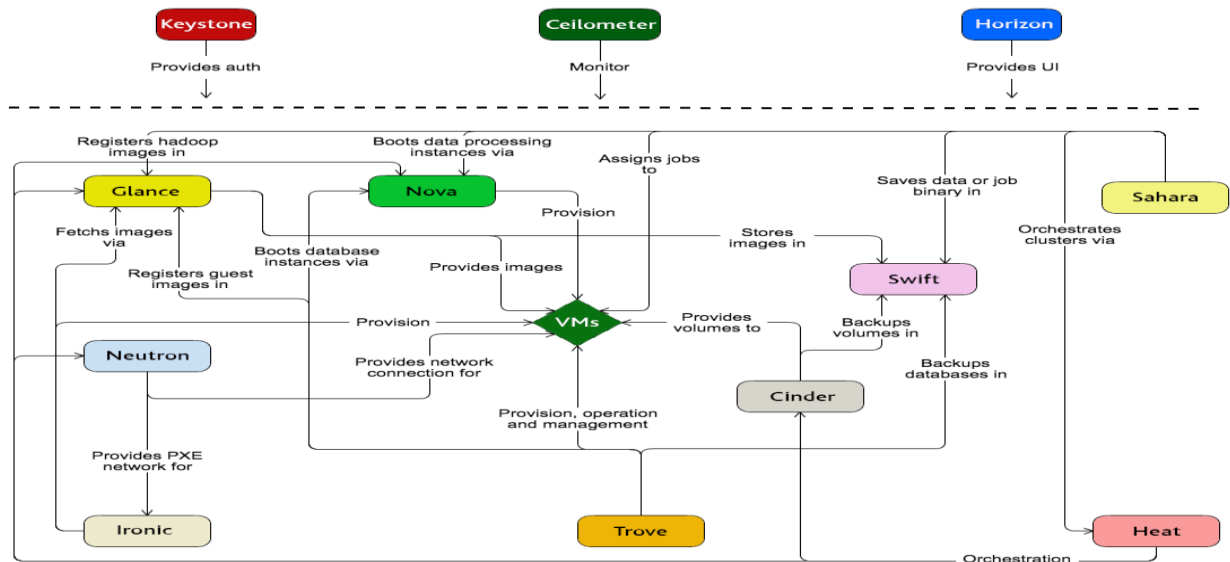


Fig. 1 OpenStack - Conceptual architecture [1]

Table 1 Openstack vs Apache Cloudstack [2]

No	Component (Code name)	Short description
1	Compute (Nova)	Fundamental IAAS part. Hypervisors upheld: Xen, KVM, Hyper-V, VMware
2	Image Service (Glance)	Discovery, enlistment, and delivery services for disk and server images
3	Object Storage (Swift)	Scalable redundant storage system
4	Dashboard (Horizon)	Graphical interface to access, provision, and automate cloud-based resource
5	Identity Service (Keystone)	Provides a central registry of users mapped to the OpenStack services they can access
6	Networking (Neutron)	Framework for overseeing systems and IP addresses
7	Block Storage (Cinder)	Persistent block-level storage system for use with OpenStack compute instances
8	Orchestration (Heat)	Manages multiple composite cloud applications using templates
9	Telemetry (Ceilometer)	Provides a Single Point Of Contact for charging frameworks, providing all details necessary to set up customer charging, over all present and future OpenStack components
10	Database (Trove)	Database-as-a-service provisioning relational and non-relational database engine
11	Elastic Map Reduce (Sahara)	Quick provision for Hadoop clusters by specifying necessary parameters
12	Bare Metal Provisioning (Ironic)	Provisions bare metal machines instead of virtual machines
13	Multiple Tenant CloudMessaging (Zaqar)	Multi-tenant cloud messaging service for Web developers
14	Sared File System Service (Mania)	API for managing shares in a vendor agnostic framework
15	DNSaaS (Designate)	DNS as a Service
16	Security API (Barbican)	REST API intended for the protected storage, provisioning and administration of secrets

APACHE CLOUDSTACK

CloudStack is quickly gaining momentum amongst several organizations. Initially developed by Cloud.com on 2010, CloudStack was purchased by Citrix then later on released into the Apache Incubator program. It is now governed by the Apache Software Foundation (ASF). Since the Apache transition, other vendors have also joined the effort by enhancing and adding more capabilities to the core software. On 2013, CloudStack graduated from Apache Incubator and became a Top-Level Project (TLP) of Apache Software Foundation. The first stable (maintenance) release after graduation is version 4.0.2. [3] [4]

Today CloudStack is a project built around a committee of developers and a VP/Chair. Even though the project has a smaller community than OpenStack, it is a project that is commercially supported its main contributor, Citrix.

Resources within the cloud are managed as follows:

- **Regions:** Grouping cloud resources into geographic regions and managing them with one or multiple servers.
- **Zones:** A zone is structured like a datacenter. A zone is composed from one or more pods and must include also a secondary data storage.
- **Pods:** Physically is the equivalent of a data rack that includes at least a layer-2 switch and one or multiple clusters.
- **Clusters:** Can be identified with one or more hosts that includes also primary storage.

- *Host*: A single server as part of a cluster. Usually the host is the hypervisor.
- *Primary Storage*: The storage resource used by a single cluster mainly for running the operating system. This type of storage usually is much faster than the other types of storage used because of higher RPM hard drives or SSD hard drives.
- *Secondary Storage*: A high capacity storage resource which stores usually disk images, OS ISO files, etc.

CloudStack structure is represented on Fig. 2.

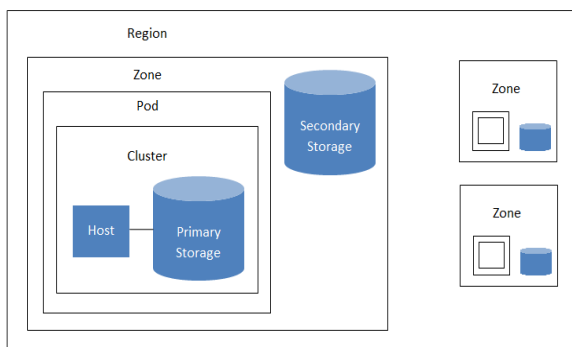


Fig. 2 Cloudstack cloud infrastructure overview [5]

hypervisors bridge. AWS (Amazon Web Services) uses also this type of networking.

2. Advanced: Layer-2 network isolated by using VLANs, and/or SDN technologies like as Nicira NVP.

The CloudStack conceptual architecture is represented on Fig. 3.

Features and functionality [5]:

- Works with hosts running XenServer/XCP, KVM, Hyper-V, and/or VMware ESXi with vSphere
- Provides a friendly Web-based user interface for managing the cloud
- Provides API's for controlling the cloud
- Provides an Amazon S3/EC2 compatible API
- Manages storage for instances running on the hypervisors (primary storage) as well as templates, snapshots, and ISO images (secondary storage)

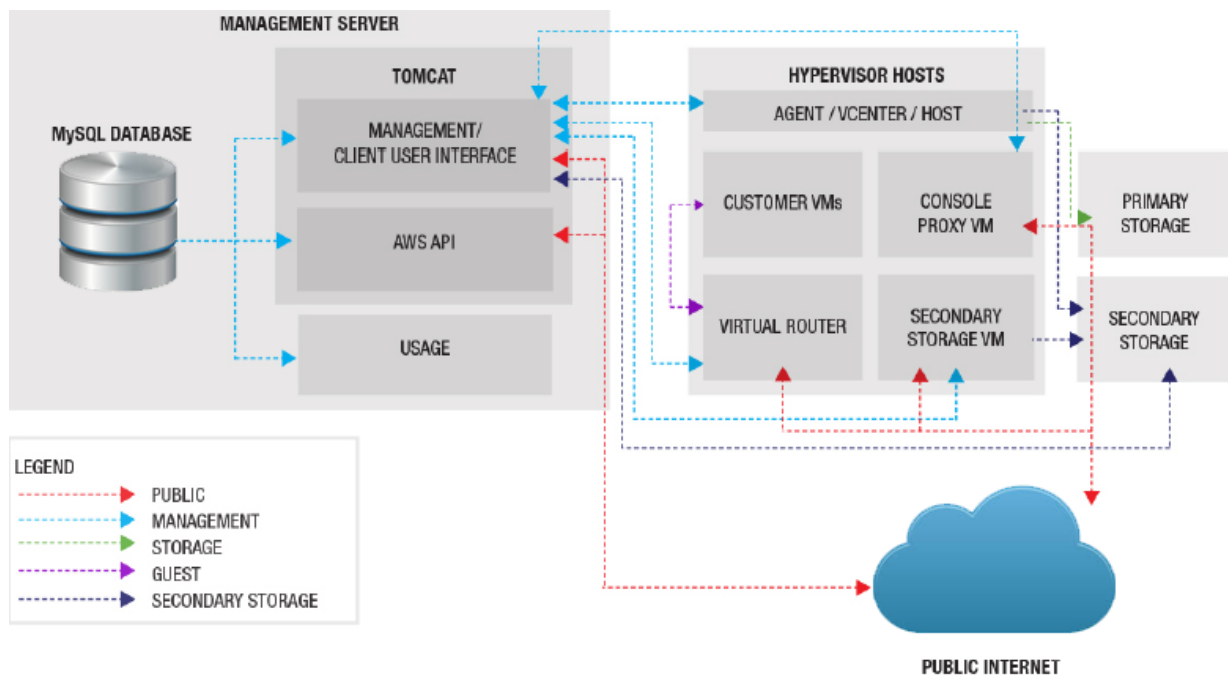


Fig. 3 Cloudstack cloud platform - Conceptual architecture

Cloudstack includes multiple types of networking that fall into two groups:

1. Basic: Single layer-2 network. Here the guest isolation is made at layer-3 using the

- Orchestrates network services from the data link layer (L2) to some application layer (L7) services, such as DHCP, NAT, firewall, VPN, and so on
- Accounting of network, compute, and storage resources
- Multi-tenancy/account separation
- User management

2. OPENSTACK ARCHITECTURE

COMPUTE

OpenStack's compute layer is managed by Nova project. Nova manage and automate pools of computer resources and it can work with many available virtualization technologies, respectively bare metal and high-performance computing (HPC) configurations.

The architecture supports XenServer/KVM, VMware, LXC and Hyper-V.

BLOCK STORAGE

Block storage is managed by Cinder project and it is having a separate management process and corresponding API. Cinder manages the provisioning of volumes and communication of the volume location to the compute project. It can integrate

It can integrate directly with many back-end systems and supports as well several protocol implementations. It is recommended to set up one Cinder management server per storage back-end. The most recent OpenStack variants permits arrangement of simultaneous back-ends through one management server.

Based upon utilized variant of OpenStack, and the wanted networking model, networking administration is done through either the Nova system framework or the Neutron project. This part of the framework of the system routes traffic through the networking management nodes. The Nova networking project uses a simplified system of networking functions while Neutron offers more advanced capabilities and includes also plugins to support external applications. Neutron is supporting mainly KVM and is having constrained backing for different hypervisors.

OBJECT STORAGE

OpenStack utilizes object storage for resources that aren't required for typical VM operation. Swift project is the administrator for object storage. Glance project is in charge of template management. Each of these systems have their own administration process and APIs. The Swift sub-project it is regularly utilized outside of a full OpenStack deployment to provide replicated object storage for third party cloud providers (ex. Amazon S3)

USER INTERFACE

The OpenStack interface, part of the Horizon project, consolidates a few APIs from different sub-projects into a single web interface. It is another administration process running independently from the already mentioned systems, and must be configured to communicate with each of other frameworks. This interface permits users to have a friendly environment that allows access to all of cloud resources.

SYSTEM USAGE

The Ceilometer sub-project is in charge of gathering utilization data and for making an interpretation of it into consumable information. It is a new project that is not completely incorporated into the other sub-projects. Some sub-projects have been coordinated with Ceilometer and offers information through a common event bus, in this manner eliminating the requirement for intermittent data surveying.

AUTHENTICATION/AUTHORIZATION

This component is part of the Keystone sub-project and is in charge of identity verification for the entire cloud environment. It is also configured as a separate set of management process and APIs. Keystone uses for authentication RBAC (Role Based Access Control) permission scheme.

3. CLOUDSTACK ARCHITECTURE

COMPUTE

The compute segment of the general design is the Infrastructure-as-a-Service (IaaS) foundation [6]. The compute layer's sole intention is to give a running virtual machine. Each compute node is managed through the center administration

server and supports numerous hypervisors and virtualization platforms, including XenServer, vSphere, KVM and Hyper-V. It influences the local abilities of the hypervisor and is persistently adding more integration for all hypervisors. CloudStack has added recently support for VMware vMotion and vStorage Motion with the goal of handling enterprise workloads.

BLOCK STORAGE

The management server keeps communication with the hypervisor in order to provision volume resources and utilizes the protocols supported by it to have access to all of its resources. The latest cloud versions permit native storage area network (SAN) communication from the management server to facilitate provisioning, snapshots and thin provisioning. This allows more advanced capabilities of the back-end SAN to be leveraged. The block storage administration is configured so the management server is not involved in the block level operations. The block storage it is only used to manage volumes.

CLOUD NETWORKING

Networking is essential to provide best performance for cloud environment. Through CloudStack, networking is organized utilizing a mix of the center administration and a virtual machine (VM) application. This is deployed by the administration server automatically and services individual customer networks. There are two essential methods of network administration: Basic (AWS style) layer-3 and Advanced (VLAN) layer-2 isolation.

CloudStack networking likewise features custom network service offerings. This permits cloud managers to adapt the cloud networking to meet customer needs including external load balancing and firewall integration and also custom CIDR, VPN connection. A VM application provides all customers enhanced network functionality without extra administration servers to configure and maintain. This configuration gives adaptability in making multi-level isolated networks to support n-level enterprise applications and site-to-site VPN to existing networks and resources.

SECONDARY STORAGE



Secondary storage was intended to store templates, snapshots and ISOs. It is coordinated through a mix of the core management server and a VM application deployed afterwards analyzed by an external process and converted to the desired format.

AUTHENTICATION/AUTHORIZATION

CloudStack system features authentication and authorization in the core management server. Additionally, there is available the option to integrate with LDAP and Active Directory as part of the core management server. This type of architecture layer includes several levels of access and nesting to create a hierarchy of resource pools.

Each component of OpenStack and CloudStack have been analyzed and compared in detail. The results are synthesized on Table 2.

Table 2 OpenStack vs Apache Cloudstack

		
Compute	Sub-project: Nova <ul style="list-style-type: none"> • It is functioning as a separate management process and API • Supports mainly XenServer/KVM and with limited functionality for VMware and Hyper-V 	<ul style="list-style-type: none"> • Managed through the core management server • Supports the following hypervisors: XenServer, VMware, KVM and Hyper-V
Block storage	Sub-project: Cinder <ul style="list-style-type: none"> • It is functioning as a separate management process and API to organize the provisioning of volumes • Communicates volume location to the compute project 	<ul style="list-style-type: none"> • Management server communicates with the hypervisor • Latest versions permits local provisioning, snapshots and thin provisioning
Cloud Networking	<ul style="list-style-type: none"> • Sub-project: Neutron • Provides "networking as a service" between interface devices (e.g. vNICs) managed by other Openstack services (e.g. Nova). 	<ul style="list-style-type: none"> • Coordinated using a blend of central administration and virtual machine appliances • Deployed by default and manages individual customer networks
Templates, Snapshots and ISOs	<ul style="list-style-type: none"> • Sub-project: Swift • Offers cloud storage software that permits storing and retrieving data with an API. It's built for scale, availability and concurrency across the entire data set. Swift is ideal for storing unstructured data that can grow without bound. • Template management is provided through Glance, another sub-project 	<ul style="list-style-type: none"> • Known as Secondary Storage in CloudStack, it was intended to store layouts, snapshots and ISOs • Provides a bridge between end users and the storage area
User Interface	<ul style="list-style-type: none"> • Sub-project: Horizon • Combines several APIs from different subprojects into a single web user interface 	<ul style="list-style-type: none"> • Part of the core management server application • AJAX base web UI • Handles requests of the system for administrators and end users
System Usage	<ul style="list-style-type: none"> • Sub-project: Ceilometer • Responsible for gathering usage event data and for making an interpretation of it into usable information 	<ul style="list-style-type: none"> • Core management server gathers events related to resources • Usage and Events Such as beginning, halting or changing a VM
Authentication/ Authorization	<ul style="list-style-type: none"> • Sub-project: Keystone • Identity service for authentication (authN) and high-level authorization (authZ). It currently supports token-based authN and user-service authorization. 	<ul style="list-style-type: none"> • Integrates with LDAP and Active Directory • Includes several levels of access that permits making a hierarchy of resources pools

4. CONCLUSION

CloudStack was designed as a singular system that operates in a cohesive manner. All components are working together and operate from a single unified core with a central database. Cloudstack includes distributed management servers through the various VM appliances. These servers are completely automated by the core management. The cloud administrator doesn't have to separately administer these components. The advantage of having a single unified core is the ability to administer everything with the help of a single API.

OpenStack was composed from the earliest starting point to be a cooperation of to some degree separate projects, which affected the general architecture and design throughout the system. This gives alternatives to the user to decide which component they want to utilize or

incorporate, additionally puts more obligation on the cloud supplier to settle on a certain design. There have additionally been occurrences of spin off projects inside of the OpenStack project. The Cinder project is a good illustration of a spin off, as it was initially part of the Nova volume project, however was split off to decrease the general size of the Nova project and to dedicate resources to its development.

Cloudstack installation and configuration is relatively straightforward, taking into account the complexity of the used infrastructure, whereas OpenStack is cumbersome and disjointed to deploy and configure. On the opposite side the

CloudStack management server works just right after the setup process.

While CloudStack is a monolithic system, its adaptability and versatility can be utilized to meet a wide arrangement of prerequisites. OpenStack has a more modular nature, which

results in the setup and configuration being more complex. CloudStack provides the functionality to deploy environments rapidly; with easy support for common protocols, it is possible to stand up a practical cloud in just few hours versus days Openstack that is more complex to install and setup.

Despite the fact that the CloudStack extend right now has less brand visibility in the cloud scene, Citrix's commitment offers credibility to the product and its ongoing development for service providers. Although, the installation of OpenStack is complicated, it offers a more stable platform than CloudStack and adding advanced features that offers is a more attractive adopted by large corporations. [7] [8].

5. REFERENCES

[1] "OpenStack," [Online]. Available: <https://www.openstack.org/>. [Accessed 29 03 2016].

- [2] "OpenStack," [Online]. Available: <https://en.wikipedia.org/wiki/OpenStack>. [Accessed 29 03 2016].
- [3] Salam Ismaeel, "Open Source Cloud Management Platforms: A Review," in 2015 IEEE 2nd International Conference on Cyber Security and Cloud Computing, New York, NY, 2015.
- [4] "Apache CloudStack," [Online]. Available: https://en.wikipedia.org/wiki/Apache_CloudStack. [Accessed 29 03 2016].
- [5] "Apache CloudStack - Open source cloud computing," [Online]. Available: <https://cloudstack.apache.org/>. [Accessed 29 03 2016].
- [6] B. Arshdeep and M. Vijay, "Cloud Computing: A Hands-On Approach," CreateSpace Independent Publishing Platform, 2013, pp. 86-88.
- [7] "A Comparative Study of OpenStack and CloudStack," in 2015 Fifth International Conference on Advances in Computing and Communications, Kochi, India, 2015.
- [8] Aaron Paradowski, "Benchmarking the Performance of OpenStack and CloudStack," in 2014 IEEE 17th International Symposium on Object/Component-Oriented Real-Time Distributed Computing, Reno, NV,