

Open Source Cloud Technologies

Salman A. Baset
sabaset@us.ibm.com
IBM T. J. Watson Research Center

ABSTRACT

Open source cloud technologies such as OpenStack, CloudStack, OpenNebula, Eucalyptus, OpenShift, and Cloud Foundry have gained significant momentum in the last few years. For a researcher and practitioner, they present a unique opportunity to analyze, contribute, and innovate new services using these technologies. The first part of the tutorial will provide an overview of OpenStack and CloudStack, two open source infrastructure as a service (IaaS) cloud platforms. The second part of the tutorial will present a detailed analysis of different OpenStack components, namely, glance (image service), nova (compute service), keystone (security service), quantum (network service), and swift (object storage service). In particular, the tutorial will describe the scheduling and provisioning process in OpenStack, and how different configuration options lead to myriad provisioning performance. Further, the tutorial will describe how OpenStack has evolved over releases. Finally, the tutorial will describe weaknesses of OpenStack and highlight important areas where researchers can contribute.

Categories and Subject Descriptors

C.2.4 [Distributed Systems]: Distributed applications

General Terms

Documentation, Management, Measurement

Keywords

Cloud computing, OpenStack, IaaS

1. PART 1: OVERVIEW OF OPENSOURCE CLOUD TECHNOLOGIES

1.1 OpenStack

OpenStack [2] is an open source software that can be used to convert compute and storage hardware into an IaaS Cloud. It provides compute and storage services similar to Amazon EC2, S3, and EBS. OpenStack was initiated by Rackspace and NASA in July 2010, and is released under Apache 2.0 License. OpenStack is completely written in Python and supports KVM, Xen, and VMware hypervisors. Table 1 shows the evolution of OpenStack code base.

There are five main components of OpenStack, namely, nova (compute service), glance (image management service), keystone (identity service), quantum (network service), and swift (object storage service). From the management perspective, OpenStack comprises of a Cloud controller(s) and physical servers that are part of OpenStack's compute or storage cloud.

1.2 CloudStack

CloudStack [1] is an open source IaaS platform. Originally developed by cloud.com, it was acquired by Citrix in July 2011. The source code of CloudStack was released under Apache 2.0 license in April 2012. CloudStack supports KVM, Xen,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SOC'12, October 14-17, 2012, San Jose, CA USA

Copyright 2012 ACM 978-1-4503-1761-0/12/10 ...\$15.00.

Release	Nova	Swift	Glance	Keystone	Quantum
Austin Oct'10	16858	12979	*2302		
Bextar Feb'11	28272	16014	3652		
Cactus Apr'11	43408	16665	4648		
Diablo Sep'11	66070	15591	9779	12493	*17556
Essex Apr'12	88006	17646	15448	10999	*22014
Folsom Sep'12	134954	19114	20298	13939	42118

Table 1: OpenStack evolution in terms of lines of code (LOC) written in python excluding any testing code. An empty box indicates that the component was not part of OpenStack at the time of release, while an * indicates that the project was still in incubation. The LOC for nova folsom release also include cinder (31279).

Feature	OpenStack (Folsom)	CloudStack (Acton 3.0.x)
Source code	Python	Java, Python, Shell scripts
License	Apache 2.0	Apache 2.0
Installation	Available through Ubuntu (apt-get) and RHEL (yum) default repos	git
Hypervisors	KVM, Xen Server, VMware ESX	KVM, Xen Server, VMware ESX, Oracle OVM
Network as a service	Quantum, OpenFlow integration	SDN, OpenFlow integration
Network capabilities	VLAN, basic networking	VLAN, VPN, firewall, load balancer
Block storage	Local disk, iSCSI, NFS	Local disk, iSCSI, NFS, Fiber channel
Deployment architecture	Region (data center), hypervisor	Zones (data center), pods, cluster (rack), hypervisor
Rich organizational control	No	Yes
Multi-role support	Some	Rich
Console access	AJAX (JRE, ActiveX, Flash)	VNC only
Web interface	Basic	Rich

Table 2: Feature comparison of OpenStack and CloudStack

VMware, Oracle VM (OVM) hypervisors. It is mostly written in Java and uses native Java mechanisms for threads and message passing. Its latest release (Acton 3.0.x) has more than 1.2 million lines of java source code.

CloudStack defines a deployment architecture, which divides the cloud into zones, pods, clusters, and hypervisors. A data center can have one or more zones. A pod is typically a rack of machines, where as cluster are a group of machines within a pod. CloudStack also enables rich organizational control. It introduces the concept of domains, which can have one or more accounts, which in turn have one or more users. Users within the same account can view the activities of other users, but they cannot view the activities of users in different accounts or domains.

Table 2 shows a brief comparison of OpenStack and CloudStack features and design choices.

2. PART 2: OPENSTACK ANALYSIS

The second part of the tutorial will provide a detailed analysis of different OpenStack (Essex and Folsom release) components, and interaction between these components to accomplish operations such as creating a server, maintaining users and quotas, message passing, network virtualization, configuration and storage options, monitoring options, weaknesses, and evolution of OpenStack over time. Next, we give a brief overview of these aspects.

Keystone [4] provides a single sign-on identity and access control mechanism across different components of OpenStack using a token-based authentication mechanism.

Glance [3] provides an image catalog from which VMs can be provisioned. It provides a registry meta-data and API service. User authentication for image access is performed through keystone.

Nova [5] provides an EC2 like compute service. It comprises of an API server (nova-api) which accepts requests for creating virtual machines, and their local or remote disks. Nova has a scheduler (nova-scheduler) which determines where to place virtual machines (VMs) or volumes. Each compute node (or hypervisor) runs a compute (nova-compute) service. The persistent state of the system is stored in a database. All nova components interact with each in a completely asynchronous

fashion using a message passing system which runs the advanced message queuing protocol (AMQP). Like glance, nova also authenticates incoming requests with keystone using keystone's REST API.

By default, VM disks are created on the same hypervisor on which the VM runs. Nova provides a nova-volume service which allows for exporting block storage for VMs over iSCSI. The block storage is referred to as a volume. Nova also supports live migration.

Quantum [6] provides the network virtualization service for achieving on-demand network topology creation, multi-tenancy, and isolation.

Swift [7] is a redundant storage system. It can store objects of different sizes which are replicated by the Swift software to other storage nodes. Swift does not require RAID or specialized storage hardware.

Provisioning Process of OpenStack heavily depends on the underlying configuration of the OpenStack system. For the out-of-the-box OpenStack configuration, an image is copied over the network from glance to the compute node. If the image is in qcow2 format, it is converted into raw format. Then a copy of the image is made to match the instance type. Finally, qcow2 disks are created from the copy of the image. During subsequent provisioning of the same VM type, there is no need to copy the image over the network. Consequently, subsequent provisions can be fast.

2.1 Limitations of OpenStack (Folsom)

OpenStack (Folsom release) does not have monitoring service, and neither does it automatically provide high availability of its components. It does not have rich organizational control. Some functions that can be performed using command-line tools cannot yet be performed by API. In Folsom release, the network manager has been centralized again with the introduction of Quantum which may present challenges when scaling the virtualized network service. The scheduler, although pluggable, is still relatively immature and will likely see future enhancements.

3. REFERENCES

- [1] CloudStack. <http://cloudstack.org/>, 2012. [Online; accessed September 2012].
- [2] OpenStack. <https://www.openstack.org>, 2012. [Online; accessed September 2012].
- [3] OpenStack Glance. <http://docs.openstack.org/developer/glance/>, 2012. [Online; accessed September 2012].
- [4] OpenStack Keystone. <http://docs.openstack.org/developer/keystone/>, 2012. [Online; accessed September 2012].
- [5] OpenStack Nova. <http://docs.openstack.org/developer/nova/>, 2012. [Online; accessed September 2012].
- [6] OpenStack Quantum. <http://wiki.openstack.org/Quantum>, 2012. [Online; accessed September 2012].
- [7] OpenStack Swift. <http://docs.openstack.org/developer/swift/>, 2012. [Online; accessed September 2012].