

Comparative Study of Eucalyptus, Open Stack and Nimbus

Lakshmi D Kurup, Chandni Chandawalla, Zalak Parekh, Kunjita Sampat

Abstract- Cloud computing is a Service Oriented Architecture which reduces information technology overhead for the end-user and provides great flexibility, reduced total cost of ownership, on-demand services and many other benefits. Hence it delivers all IT related capabilities as services rather than product. Services on cloud are divided into three broad categories: Software as a Service, Infrastructure as a Service & Platform as a Service. Same as services cloud is also classified as Private Cloud, Public Cloud & Hybrid Cloud. Private cloud is gaining popularity, not only among large organizations but also small and medium enterprises. To deploy public or private cloud there are many open source software platforms available such as Eucalyptus, Nimbus, OpenStack, Open Nebula, Cloud Stack and Amazon Web Services. In this paper, we provide a comparative study of three open source cloud management platforms: Eucalyptus, OpenStack and Nimbus. We believe that the comparison presented in this paper would benefit enterprises as well as research institutes in selecting best open source platforms to meet their technology demands.

Keywords: - OpenStack, Cloud computing, Eucalyptus, Nimbus, Private Cloud, Public Cloud, Hybrid Cloud.

I. INTRODUCTION

As per the NIST definition, "cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction." There are basically below mentioned three kinds of services provided by cloud computing platforms:

- i. IaaS (Infrastructure as a Service): Infrastructure as a service (IaaS) is a standardized, highly automated offering, where compute resources, complemented by storage and networking capabilities are owned and hosted by a service provider and offered to customers, on-demand. Customers are able to self-provision this infrastructure, using a Web-based graphical user interface that serves as an IT operations management console for the overall environment. API access to the infrastructure may also be offered as an option.

- ii. PaaS (Platform as a Service): A platform as a service (PaaS) offering, usually depicted in all-cloud diagrams between the SaaS layer above it and the IaaS layer below, is a broad collection of application infrastructure (middleware) services (including application platform, integration, business process management and database services). However, the hype surrounding the PaaS concept is focused mainly on application PaaS (aPaaS) as the representative of the whole category.
- iii. SaaS (Software as a Service): Gartner defines software as a service (SaaS) as software that is owned, delivered and managed remotely by one or more providers. The provider delivers software based on one set of common code and data definitions that is consumed in a one-to-many model by all contracted customers at any time on a pay-for-use basis or as a subscription based on use metrics.

IaaS services are further categorized as: Private Cloud — the cloud infrastructure has been deployed, and is maintained and operated for a specific organization. Public Cloud — the cloud infrastructure is available to the public on a commercial basis by a cloud service provider. Hybrid Cloud — the cloud infrastructure consists of a number of clouds of any type, but the clouds have the ability through their interfaces to allow data and/or applications to be moved from one cloud to another. The main goal of this paper is to present a comparison among the three main open source software for IaaS cloud implementation that are: Eucalyptus, OpenStack and Nimbus.

A. Eucalyptus

EUCALYPTUS stands for Elastic Utility Computing Architecture for Linking Your Program To Useful System. It is open source software that was developed by University of California-Santa Barbara for Cloud Computing to implement Infrastructure as a Service. In early 2008, it become the first open source software which is compatible with Amazon Web Service API for deploying On-premise private cloud. Amazon Web Service (AWS) is one of the major players for providing Infrastructure As A Service. [2, 3] They have two popular services Elastic Compute Cloud (EC2) and Simple Storage Service (S3). Eucalyptus provide an EC2-compatible cloud Computing Platform and S3-compatible Cloud Storage thus its services are available through EC2/S3 compatible APIs [4]. Eucalyptus can leverage (drag) a heterogeneous collection of virtualization technologies within a single cloud, to incorporate resources that have already been virtualized without modifying

Manuscript Received on December 2014.

Lakshmi D Kurup, Asst. Prof., Dwarkadas J Sanghvi College of Engineering, Mumbai, India.

Chandni Chandawalla, Bachelor of Engineering, Dwarkadas J Sanghvi College of Engineering, Mumbai, India.

Zalak Parekh, Bachelor of Engineering, Dwarkadas J Sanghvi College of Engineering, Mumbai, India.

Kunjita Sampat, Bachelor of Engineering, Dwarkadas J Sanghvi College of Engineering, Mumbai, India.

their configuration. Eucalyptus has five high-level components: Cloud Controller (CLC) -it is the entry point into the private cloud for end user, project managers developers and administrator. It also help in manage virtualized resources. Walrus: It implements bucket-based storage, which is available inside and outside the cloud system. Cluster Controller (CC):It executes on a machine that has network connectivity to the machines that are running on Node Controller and Cloud Controller .It manages the Virtual Machine (VMs) Network. All Node Controllers associated with a single CC must be in the same subnet.; Walrus is the storage system, which allow user to store data, organized as bucket and object, it is also used to create, delete, and list buckets. Storage Controller (SC) provides block-level network storage including support for Amazon Elastic Block Storage (EBS) semantics; and Node Controller (NC) is installed in each compute node to control Virtual Machine activities, including the execution, inspection, and termination of VM instances.

B. OpenStack

In July 2010, OpenStack was announced and the initial contributors of it are NASA and Rackspace. It is the fastest growing free open source software. Rackspace contributed their "Cloud Files" platform (code) to power the Object Storage part of the OpenStack, while NASA contributed their "Nebula" platform (code) to power the Compute part. OpenStack is a collection of open source software project that developers and cloud computing technologist can use to setup and run their cloud compute and storage infrastructure. Its services are available through Amazon EC2/S3 compatible APIs and hence the client tools written for AWS can also be used with OpenStack. [15, 16]It consists of three core software projects: OpenStack Compute Infrastructure also called Nova; OpenStack Object Storage Infrastructure also called Swift and OpenStack Image Service Infrastructure also called Glance. Nova is the main part of Infrastructure as a service and it also is the computing Fabric controller for the OpenStack cloud .Enterprises/Organization can use Nova to host and manage their cloud computing systems. Nova manages all the activities that are needed to support life cycle of instances within the open stack. Swift offers a distributed, consistent virtual object container in which lots of data can be store and from which data can be retrieved. It is capable of storing large number of object distributed across nodes. Glance is a lookup and retrieval system for virtual machine images.

C. Nimbus

Nimbus is an open-source toolkit focused on providing Infrastructure-as-a-Service (IaaS) capabilities to the scientific community. To achieve this, the focus is on three goals:

Enable providers of resources to build private or community IaaS clouds: The Nimbus Workspace Service provides an implementation of a compute cloud allowing users to lease computational resources by deploying virtual machines (VMs) on those resources. A complementary tool, Cumulus, provides an implementation of a quota-based storage cloud. Cumulus is designed for scalability and allows providers to configure multiple storage cloud implementations. Enable users to use IaaS clouds: An example of a tool in this area is the Nimbus Context Broker, which creates a common configuration and security context across resources

provisioned from potentially multiple clouds. Nimbus also offers scaling tools allowing users to automatically scale across multiple distributed providers. These tools are called "sky computing tools" as they often operate in a multi-cloud environment combining private and public cloud capabilities. Enable developers to extend, experiment and customize IaaS: To achieve this goal we provide a high-quality, highly configurable and extensible open source implementation. For example, the Workspace Service can be configured to support different virtualization implementations (Xen or KVM), resource management options (including schedulers such as PBS), interfaces (including compatibility with Amazon EC2), and many other options. Nimbus components are regularly acceptance tested which allows developers to easily extend them. Over time, the project has attracted a community of open source contributors and committers.

Combining those tools and capabilities in different ways allows users to rapidly develop custom community-specific solutions.

II. COMPARATIVE STUDY

The comparison between Eucalyptus, OpenStack and Nimbus is based on the study of the architecture of respective open source platform. Each Open source software provides Infrastructure as a Service to deliver virtualization environment, as a service to user.

A. Origin and Community Support

Eucalyptus: Eucalyptus was the result of research project of the University of California, Santa Barbara, Department of computer science. It has a powerful community that contributes to platform development and assists in finding and fixing [1].

OpenStack: OpenStack started in summer 2010 when RackSpace and NASA jointed its initial project "Cloud Files" and "Nebula" respectively. OpenStack is led by a foundation (integrated by 850 companies and 4500 individual members) and has a broad range of support from major tech industry players, ranging from HP, Dell, IBM, RackSpace, NASA , Cisco, NEC, AT&T ,Bull, EMC, Brocade and dozens of other companies. Its community integrated by around 7000 person over 87 countries.

Nimbus: Nimbus is a cloud computing infrastructure project currently directed by Argonne National Laboratory operated by the University of Chicago. The project went live in 2005 with the goal of creating an open source cloud computing infrastructure to allow scientists working on data-intensive research projects to access configurable virtual machines for computation purposes. Nimbus aims to allow building of private clouds, allow users to use clouds by leasing resources and allow developers to experiment with Nimbus.


B. System Architecture and Design

Eucalyptus

Eucalyptus was designed with the goal of its use to be in research setting. According to its designer [1, 6, 7, 8, 9], it is designed to be extensible and easy to modify and use with not require dedicated resources. Eucalyptus components were designed and implemented as web services and thus each service had a WSDL interface. Use of web services made it possible to utilize WS-

security to secure communications between components and to plug in components as needed. The Eucalyptus system is composed of composed of a 5 main components interacting together; Client, cloud controller, storage controller (Walrus), cluster controller and a node controller. These components are organized in hierarchical structure as illustrated in figure 2 and communicate through private and public networks as described in the networking section. Following is a description of each component according to [6, 7, 8, 9]:

1. Node Controller: Node controller is installed on each physical resource (node). It is responsible of hosting and administering virtual machine (VM) instances running on it. The node controller interface described through WSDL provides functionalities to manage its VMs such as *runInstance*, *terminateInstance*, *describeInstance*, *describeResource* and *startNetwork*. The first three functions interact with the node's underlying hypervisor to execute and govern such services. Eucalyptus supports both XEN and KVM hypervisors. The *describeResource* is a query that reports to the cluster controller its current capabilities such as number of cores, memory and disk capacity. The *startNetwork* sets up the node's virtual network.

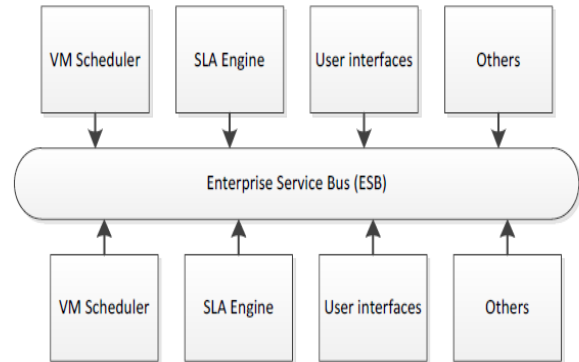


2. Cluster Controller: Cluster controller runs on one machine (usually the head node) per cluster and works as an intermediary between the cloud controller and node controller. A cluster controller manages a number of node controllers in its cluster and is connected to both public and private networks. A cluster controller collects state information about its node controllers, schedules VM creation requests and configures both public and private networks. Its WSDL is similar to that of a node controller but is specific to a number of instances rather than one. Its functions are *runInstances*, *terminateInstances*, *describeInstances* and *describeResource*. These functions interact with NC to execute. The cluster controller performs scheduling by running *describeResource* on each node controller and selecting the first free node.

3. Storage controller: Storage controller (Walrus) is a component that provides storage services for storing virtual machine images and user's data. Walrus runs also as a web service and has a WSDL interface compatible to that of S3.

4. Cloud Controller: Cloud controller is the user's entry point into the Eucalyptus system and only one instance is run on the system. It provides users with a way managing

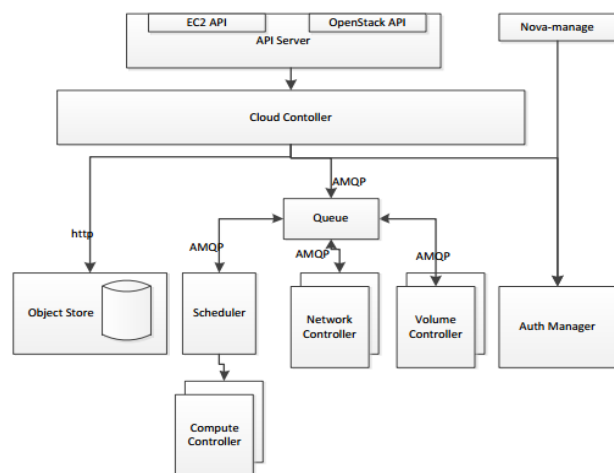
the system. Cloud controller processes and handles user's (regular users and administrators) requests, schedules VM instance creation, maintain system and user data and process service level agreements. The cloud controller is built using the Enterprise service bus providing decoupling from the services' implementation and the users requests which enabled the use of multiple interfaces and different implementations of such services.



5. Client: The client component provides the user with a way to access the Eucalyptus system (cloud controller). Eucalyptus provides two interfaces; One is a WSDL which is a SOAP client interface similar to AWS EC2 [3] interface called *euca2ools*. Another interface is a HTTP query based interface. The administrator is provided with an interface to manage users and VM images.

OpenStack

OpenStack compute was designed to be component based, highly available, fault tolerant, recoverable and provides API compatibility [18]. OpenStack compute consists of 8 main components as illustrated in figure 6. The components are designed to interact based on message based architecture. Communication between the cloud controller, scheduler, network controller and volume controller is performed through the AMQP (Advanced Message Queue Protocol) using asynchronous communication.



Following is a description of each component [17, 18]:

1. API Servers: User's front end providing interaction with the cloud controller allowing users holding all kinds of roles to access and manage storage,

network, images, projects, users and instances. Two interfaces are provided; one in EC2 compatible and another is openStack specific. Also, a nova-manage interface is provided for users holding admin roles for administration and maintenance of the cloud.

2. Cloud controller: maintains cloud state and handles interaction between services such as scheduler, volume and network controller through the use of a message queue called RabbitMQ[19].

3. Scheduler: Selects the available compute controller to host an instance.

4. Network controller: manages networking resources on each node builds virtual networks to be used by computer controllers. It provides the commands such as *Allocate Fixed IP Addresses, Configuring VLANs for projects and Configuring networks for compute nodes.*

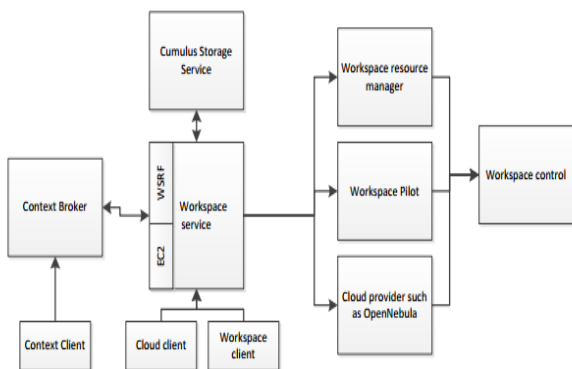
5. Volume controller: provides permanent storage for the compute controllers. It provides commands such as *Create Volumes, Delete Volumes and Establish Compute volumes.*

6. Auth Manager: Provides authorization and authentication services.

7. Compute controller: manages hosted VM instances and provides commands such as *Run instances, Terminate instances, Reboot instances, Attach volumes, Detach volumes and Get console output* through an API interface

III. NIMBUS

Nimbus offers a “cloudKit” that allows users to lease remote resources by allocating and configuring Virtual Workspace Service (VWS). The design of Nimbus which consists of a number of components based on the web service technology as illustrated in figure.



Following is a description of each component[8, 9, 23, 24]:

1. Workspace service: Allows clients to manage and administer VMs by providing to two interfaces; One interface is based on the web service resource framework (WSRF) and the other is based on EC2 WSDL. This service communicates with a workspace resource manager or a workspace pilot to manage instances.

2. Workspace resource manager: Implements VM instance creation on a site and management.

3. Workspace pilot: provides virtualization with significant changes to the site configurations.

4. Workspace control: implements VM instance management such as start, stop and pause VM. It also provides image management and sets up networks and provides IP assignment.

5. Context broker: allows clients coordinate large virtual cluster launches automatically and repeatedly.

6. Workspace client: a complex client that provides full access to the workspace service functionality.

7. Cloud client: a simpler client providing access to selected functionalities in the workspace service.

8. Storage service: Cumulus is a web service providing users with storage capabilities to store images and works in conjunction with GridFTP[26].

Feature Comparison

Feature \ IaaS	Eucalyptus	OpenStack	Nimbus
Virtual network	VDE	VDE & Open vSwitch	VDE
VM isolation	VLAN	VLAN	Not available
System Security	Firewall filters (Security groups) ,SSH & WS-Security	VPN access, SSH, security groups and CloudAudit	VPN access, SSH & use of globus certificate credentials.
User Security	User credentials provide through web interface	User credential using certificate authority for VPN access	Users x509 public key certificates which is provided to the cloud
DHCP	On the cluster controller	On the network controller	On the individual node
Hypervisor support	XEN, KVM & VMware	XEN, KVM, UML, QEMU & Microsoft hyper-V	XEN & KVM
Storage System	Warlus	Object store	Cumulus

IV. CONCLUSION

In this paper, we took a closer look into three common infrastructure as a service (IaaS) framework implementations; Eucalyptus, OpenStack and Nimbus. For each framework, we illustrated its origin, architecture and design as well as provided a feature comparison. OpenStack is the best choice for public cloud providers due to its scalability. Eucalyptus is a possible choice for private cloud providers and researchers due to its extensibility and quick installation. Nimbus is also a good choice for researchers and scientific applications since it has an efficient job scheduler. Although Nimbus started in 2005, it was lacking in detailed documentation and publications. As a future work we plan to expand our review on the security of Infrastructure as a service implementations.



REFERENCES

- [1] Eucalyptus home page, <http://open.eucalyptus.com/>, march, 2010
- [2] Amazon S3 API, <http://docs.amazonwebservices.com/AmazonS3/latest/dev/>
- [3] Amazon EC2 API, <http://docs.amazonwebservices.com/AWSEC2/2009-04-04/DeveloperGuide/>
- [4] Ubuntu cloud computing home page, <http://www.ubuntu.com/business/cloud/overview>
- [5] Eucalyptus Enterprise home page, <http://www.eucalyptus.com/>
- [6] D. Nurmi et al., "The Eucalyptus Open-Source Cloud- Computing System," Cloud Computing and Applications 2008 (CCA 08), 2008,.
- [7] Nurmi, d., wolski, r., grzegorzcyk, c., obertelli, g., soman, s., youseff, l., and zagorodnov, d. Eucalyptus: a technical report on an elastic utility computing archietcture linking your programs to useful systems . Tech. Rep. 2008-10, university of california, santa barbara, october 2008
- [8] P. T. Endo, G. E. Goncalves, J. Kelner, and D. Sadok. A Survey on Open-source Cloud Computing Solutions. Brazilian Symposium on Computer Networks and Distributed Systems, May 2010.
- [9] Peter Sempolinski, Douglas Thain, "A Comparison and Critique of Eucalyptus, OpenNebula and Nimbus," cloudcom, pp.417-426, 2010 IEEE Second International Conference on Cloud Computing Technology and Science, 2010
- [10] M. Chang, J. He, and E. Castro-Leon. Service orientation in the computing infrastructure. In SOSE '06: Proceedings of the Second IEEE International Symposium on Service-Oriented System Engineering, pages 27–33, Washington, DC, USA, 2006. IEEE Computer Society.
- [11] Eucalyptus network configuration, http://open.eucalyptus.com/wiki/EucalyptusNetworkConfiguration_v2.0
- [12] Virtual Distributed Ethernet, <http://vde.sourceforge.net>
- [13] RackSpace cloud hosting home page, <http://www.rackspace.com/cloud/>
- [14] State of the Cloud – January 2011, <http://www.jackofallclouds.com/2011/01/state-of-the-cloud-january-201/>
- [15] OpenStack Tutorial, IEEE CloudCom, 2010 <http://salsahpc.indiana.edu/CloudCom2010/slides/PDF/tutorials/OpenStackTutorialIEEECloudCom.pdf>
- [16] OpenStack compute home page, <http://openstack.org/projects/compute/>
- [17] Open Stack Administration Manual home page, <http://docs.openstack.org/openstack-10compute/admin/content/>
- [18] Nova's Documentation, <http://nova.openstack.org/>
- [19] Rabbit MQ home page, <http://www.rabbitmq.com/>
- [20] Open Virtual Switch home page, <http://openvswitch.org/>
- [21] CloudAudit home page, <http://www.cloudaudit.org/>
- [22] The Globus alliance home page, <http://globus.org/>
- [23] Keahey, K. "Nimbus: Open Source Infrastructure-as-a-Service Cloud Computing Software", Workshop on adapting applications and computing services to multi-core and virtualization, CERN, Switzerland, 2009
- [24] Nimbus home page, <http://www.nimbusproject.org/docs/2.2/faq.html>
- [25] K. Keahey, T. Freeman. "Science Clouds: Early Experiences in Cloud computing for Scientific Applications," Cloud Computing and Its Applications 2008 (CCA-08), Chicago, IL. October 2008
- [26] W. Allcock, J. Bester, J. Bresnahan, A. Chervenak, L. Liming, and S. Tuecke, "GridFTP: Protocol Extension to FTP for the Grid," Grid Forum Internet-Draft, March 2001