

Apache CloudStack: Open Source Infrastructure as a Service Cloud Computing Platform

Rakesh Kumar¹, Kanishk Jain², Hitesh Maharwal³, Neha Jain⁴, Anjali Dadhich⁵

^{1,2,3,4,5}Department of Information Technology & JECRC, Jaipur, India

Abstract— Cloud computing is changing the way IT is delivered in enterprises around the world, as well as the whole world’s leading open source cloud computing platform. Apache CloudStack, helps you implement a cloud computing service in your enterprise and it is a turnkey open source cloud management as well as flexible cloud orchestration platform for deploying reliable and also scalable public and private Infrastructure-as-a-Service (IaaS) in cloud computing environments. Apache CloudStack was mainly designed to deploy and manage a large number of virtual machines (VMs) and the software-side infrastructure required to support it. Apache CloudStack is a top-level project of the Apache Software Foundation (ASF) which provides a management server and agents for hypervisor hosts so that you can run an Infrastructure-as-a-Service (IaaS) cloud. In this research paper, we are discussing introduction, benefits, background, features, use cases and purposes of Apache CloudStack. Further discussing deployment architecture, components, API, pros and cons of Apache CloudStack and lastly discuss about some new features supported by latest release of Apache CloudStack 4.4. This paper show importance of Apache CloudStack as a Cloud provider and give the best solution for service providers, web content providers, SaaS providers, and enterprises.

Keywords— *Apache CloudStack; Cloud Computing; Cloud infrastructure; Hypervisor; IaaS; JIRA; Management Server; Virtual Machines (VMs)*

I. INTRODUCTION

Apache CloudStack [1], [16] is an open source cloud computing software, which is used to build private, public and hybrid Infrastructure as a Service (IaaS) clouds by pooling computing resources. Apache CloudStack [1], [17] manages computing, networking as well as storage resources. Apache CloudStack uses multiple hypervisors such KVM, vSphere and XenServer/XCP for virtualization purpose. Apache CloudStack [1], [21] also supports the Amazon Web Services (AWS) API, apart from its own APIs.

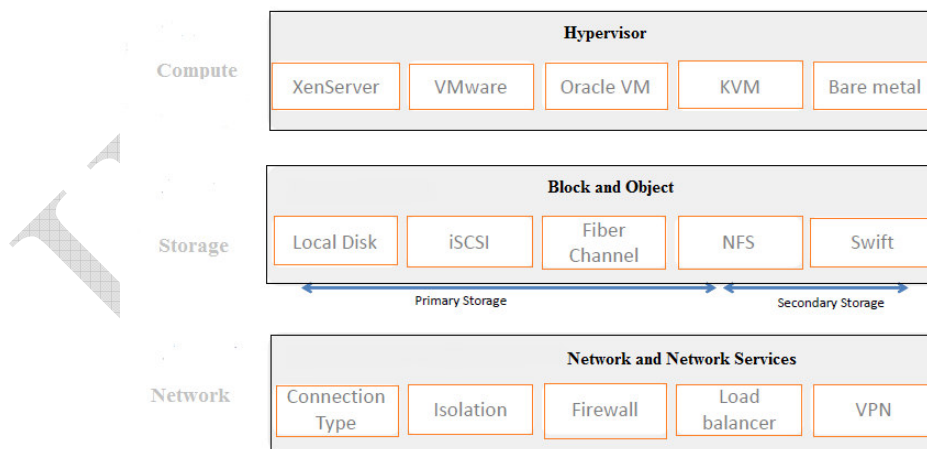


Fig. 1 Open Flexible Platform of Apache CloudStack

Apache CloudStack [8], [27] is a turnkey solution that contains the entire stack of features most organizations want with an Infrastructure-as-a-Service (IaaS) cloud such as Network-as-a-Service (NaaS), compute orchestration, user and account management, a full and open

native API, resource accounting, and a first-class User Interface (UI). Apache CloudStack [9], [26] also provides an API that's compatible with AWS EC2 and S3 for organizations which help to deploy hybrid clouds and some benefits of Apache CloudStack [7] such as:

- It mainly focus on solving business problems
- It automates deliver services quickly and helping to reduce IT operations costs
- It allows IT to provide standardized workloads, which ensure consistency
- It enhances visibility into resource allocation
- It increases the server/admin ratio and delivers benefits of scale, if deployed globally

II. APACHE CLOUDSTACK BACKGROUND

CloudStack [10], [30] was developed by Cloud.com, and Cloud.com released CloudStack as free software under the GNU General Public License, version 3 (GPLv3) in May 2010. Citrix purchased [10] Cloud.com in July 2011. In April 2012, Citrix donated CloudStack to the Apache Software Foundation (ASF). Apache CloudStack [10] became a Top-Level Project (TLP) of ASF On March 20, 2013. Some notable users of Apache CloudStack such as Amdocs, Autodesk, NTT, Zynga, China Telecom, Dell, Walt Disney, Huawei, Orange and Tata.

III. APACHE CLOUDSTACK KEY FEATURES AND USE CASES

Apache CloudStack [22] supports broad hypervisor such as XenServer, VMware, OracleVM, vSphere and KVM. Apache CloudStack [1] support scalable architecture, Multi-node installation as well as load balancing makes it highly available. Virtual networking ability is an important feature of Apache CloudStack, and also MySQL replication is useful for maintaining high availability.

Apache CloudStack [5], [7] can be used for multiple purposes such as:

- Enterprises can stage development and testing in a consistent way easing the development as well as publishing processes for applications
- Web content providers can offer scalable and elastic web infrastructure
- SaaS providers can support true multi-tenant software hosting
- Service Providers can attempt virtualized hosting in a cloud computing configuration

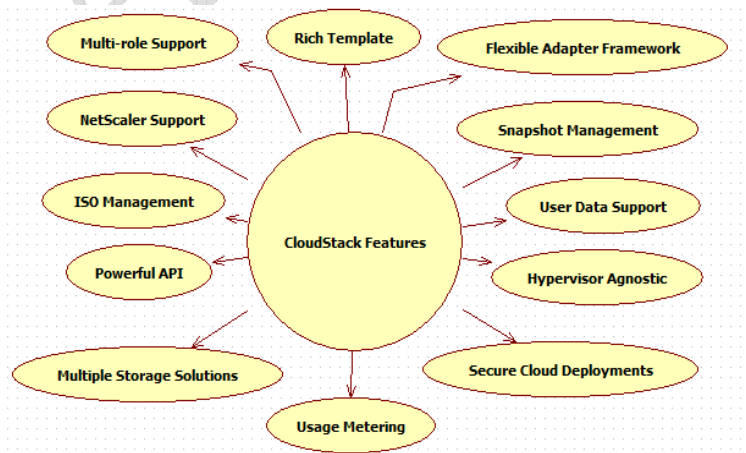


Fig. 2 Features of Apache CloudStack

Service providers, organizations [1] use Apache CloudStack to set up an on-demand as well as elastic Infrastructure as a Service (IaaS). Apache CloudStack [1], [29] can be used to set up an on premise private cloud behind the organization's firewall for internal purposes such as gaining better control over infrastructure. It can also be used to some purposes like AWS insourcing and traditional enterprise workloads.

IV. APACHE CLOUDSTACK DEPLOYMENT ARCHITECTURE

Apache CloudStack [1], [2], [15] installation consists the Management Server (Apache CloudStack software that manages cloud resources) and the cloud infrastructure. When you set up as well as manage a Apache CloudStack cloud, you provision resources such as storage, hosts, and IP addresses into the Management Server [20].




Fig. 3 Basic Deployment

A single machine [10] can act as both the Management Server and the hypervisor host when they using the KVM hypervisor, in smallest deployment. The management server controls cloud resources [1], as well as the administrator can manage and interact with the management server by using a UI or APIs.

V. APACHE CLOUDSTACK COMPONENTS

A host [1] is a computer which provides the computing resources like the CPU, storage, memory, networking, etc. to run the virtual machines (VMs) and also every host has a hypervisor installed to manage the virtual machines (VMs). Primary storage is coupled with a cluster which stores the disk volumes for all the virtual machines (VMs) running on hosts in that specific cluster.

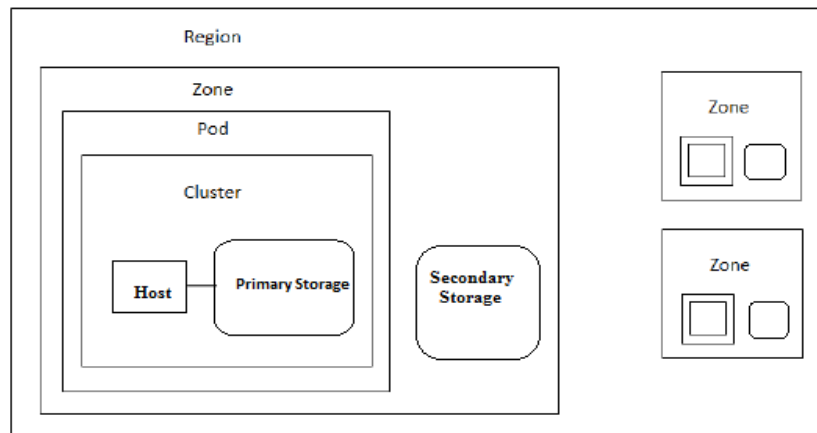


Fig. 4 A region with multiple zones

A cluster [3] contains of one or more hosts and one or more primary storage servers, and also every cluster can consist of one or more nodes. A pod [3], [6] is simply a dedicated rack of hardware, and also each pod can consist of one or more clusters. A host [24] is a single compute node which can consist of one or more virtual machines (VMs). A zone [3] simply corresponds to a single data centre and also it is permissible to have multiple zones in a data centre. Pod is contained within zones as well as every zone can contain one or more pods, and also zones can be public or private. Apache CloudStack [19] supports two types of networking basic networking, which is used for AWS style networking and advanced networking, which is used for more sophisticated network topologies.

VI. APACHE CLOUDSTACK API

The Apache CloudStack API [1], [4] supports three access roles such as root admin API, domain admin API and user API. The root admin API [4] can access all the features in addition to both virtual as well as physical resource management. The domain admin API can access the virtual resources which belong to the administrator's domain. User API can access the features that allow the management of the user's virtual machines (VMs), storage as well as network.

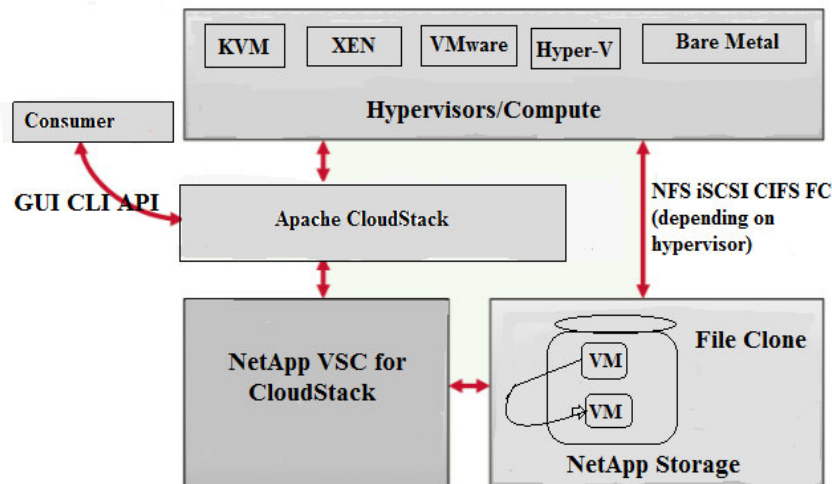


Fig. 5 Apache CloudStack

VII. APACHE CLOUDSTACK PROS AND CONS

Apache CloudStack [6], [28] abridgments complexities of dealing with compute, network, and storage. It aims at abstracting as well as unifying heterogeneous IT infrastructure into uniform services, through a common UI and user experience.

A. Apache CloudStack Pros

Apache CloudStack [18] gives featured solution designed to manage a cloud through an easy-to-use web interface, command line tools and a RESTful API. The RESTful API allows easy integration for other tools and automation frameworks. Some Apache CloudStack pros [6] such as:

- Easy to configure, use, and maintain
- Enterprise level support
- Unified management



- Allow plugin framework
- It is a mature product

B. Apache CloudStack Cons

To use the Apache CloudStack API [25], knowledge of Java, PHP, HTTP GET/POST, query strings, XML, JSON, URL of the CloudStack server, and also API key as well as secret key is necessary. Some Apache CloudStack cons [6] such as:

- Flexibility and modularity not as customizable
- It only supports Fiber Channel as primary storage through the hypervisor
- Backup and restore is not solved sufficiently and arguably

Apache CloudStack 4.4 [11] includes the some new features such as:

- Resize root disk
- Monitoring VR services
- Improvements of Hyper-V support
- Primary storage plug-in for root disks
- Added per primary storage overprovisioning
- Windowsfication of CloudStack Management Server
- Region level Guest networks as well as VPC deployment
- VMware Support for DRS (Distributed Resource Scheduler)
- Support Distributed routing and network ACL with OVS plug-in

Apache CloudStack [13], [23] is a software designed to deploy and also manage big networks of virtual machines (VMs), as well as highly scalable Infrastructure as a Service (IaaS) cloud computing platform. All new features and bugs in Apache CloudStack 4.4.0 [12] have been tracked in Jira, and also have a standard naming convention of “CLOUDSTACK-NNNN” where “NNNN” represents the issue number such as CLOUDSTACK-3968.

JIRA [14] is a commercial software product, developed since 2002 by Atlassian, Inc. JIRA is written in Java and it provides bug tracking, issue tracking, and project management functions.

VIII. CONCLUSIONS

Of all the technology trends but, cloud computing is one of the more interesting and least controversial, in my opinion. Cloud computing is most revolutionary types of technology entering corporate information technology strategies today. The latest innovations in cloud computing technology are making our business applications even more mobile and also collaborative, much similar to popular consumer apps such as Facebook and Twitter. Cloud computing will continue to expand and come to dominate information transactions because it provides many advantages, allowing users to have easy, instant, and individualized access to tools and information they need wherever they are, locatable from any networked device. Large businesses are far less likely to put most of work in the cloud anytime soon because of control and security issues, so further research is going on related to most efficient solution of security issues and increase the performances in dynamic reconfiguration in an IaaS Cloud Computing.



REFERENCES

- [1] <http://www.opensourceforu.com/2013/12/cloudstack-overview-open-source-iaas-platform/>
- [2] https://cloudstack.apache.org/docs/enUS/Apache_CloudStack/4.0.2/html/API_Developers_Guide/deploymet-architecture-overview.html
- [3] http://support.rightscale.com/09-Clouds/CloudStack/CloudStack_Reference_Architecture#Overview
- [4] <http://cloudstack.apache.org/docs/api/>
- [5] <http://cloudstack.apache.org/software/features.html>
- [6] <https://communities.netapp.com/community/netapp-blogs/the-raised-floor/blog/2014/06/23/apache-cloudstack-citrix-cloudplatform-introduction-part-i-of-iii>
- [7] <http://cloudstack.apache.org/software.html>
- [8] <http://cloudstack.apache.org/>
- [9] <http://kangaroot.net/blog/cloudstack-open-source-cloud-computing>
- [10] http://en.wikipedia.org/wiki/Apache_CloudStack
- [11] <http://cloudstack-release-notes.readthedocs.org/en/latest/about.html>
- [12] http://cloudstack-release-notes.readthedocs.org/en/latest/known_issues.html
- [13] <https://issues.apache.org/jira/browse/CLOUDSTACK/?selectedTab=com.atlassian.jira.jira-projects-plugin:summary-panel>
- [14] <http://en.wikipedia.org/wiki/JIRA>
- [15] Apache CloudStack Architecture pdf; Alex Huang, Software Architect, Citrix Systems
- [16] Getting to Know Apache CloudStack pdf; Joe Brockmeier, PPMC Member Apache CloudStack, jzb@apache.org
- [17] Apache CloudStack Cloud Computing pdf; Leverage the power of CloudStack and learn to extend the CloudStack environment; Navin Sabharwal Ravi Shankar
- [18] Apache CloudStack 4.0.0-incubating CloudStack API Developer's Guide pdf
- [19] Apache CloudStack 4.0.0-incubating CloudStack Installation Guide pdf
- [20] Apache CloudStack 4.0.2 CloudStack Installation Guide pdf
- [21] Taxonomy, Classification & Implementation of open source cloud computing platforms; CS 692 R&D Project Report; Arpit Malani (10305901); May 16, 2012 pdf
- [22] Barrie Sosinsky, "Cloud Computing Bible", Published by Wiley Publishing, Inc.; 10475 Crosspoint Boulevard; Indianapolis, IN 46256 www.wiley.com
- [23] CloudStack Documentation Release 4.3.0; Apache CloudStack; May 28, 2014 pdf
- [24] CloudStack Advanced Installation Guide; For CloudStack Version 3.0.0 – 3.0.2; Revised August 16, 2012 1:33 AM Pacific pdf
- [25] CloudStack Installation Documentation; Release 4.3, Apache Software Foundation; June 30, 2014 pdf
- [26] CloudStack Overview; Written by: Chiradeep Vittal, Alex Huang @ Citrix Revised by: Gavin Lee, Zhennan Sun @ TCloud computing pdf
- [27] Comparing Open Source Private Cloud (IaaS) Platforms; Lance Albertson OSU Open Source Lab Associate Director of Operations; lance@osuosl.org / @ramereth pdf
- [28] Building Clouds on Apache CloudStack: An Introduction; Linux CloudOpen 23 Oct 2013; Giles Sirett CEO ShapeBlue, Giles.sirett@shapeblue.com; Twitter: @ShapeBlue pdf
- [29] Hireen Parmar, Tushar Champaneria "Comparative Study of Open Nebula, Eucalyptus, Open Stack and Cloud Stack"; International Journal of Advanced Research in Computer Science and Software Engineering; Volume 4, Issue 2, February 2014; ISSN: 2277 128X
- [30] Rakesh Kumar, Neha Gupta, Shilpi Charu, Kanishk Jain, Sunil Kumar Jangir; "Open Source Solution for Cloud Computing Platform Using OpenStack"; International Journal of Computer Science and Mobile Computing; Vol. 3, Issue. 5, May 2014, pg.89 – 98; ISSN 2320-088X