A Survey on Commercial and Open Source Cloud Monitoring

Hemanshu A Patel, Arvind D Meniya

***Abstract*— Cloud Monitoring plays a crucial role in providing application guarantees like performance, availability, and security. We can understand cloud computing as technologies that rely on the Internet to satisfy the computing needs of users, who do not generally own the physical infrastructure. All services are often provided by a third party with several common business applications online. Users can choose the services they want and access them from a web browser, while the software and data are stored on third party's company.**

Cloud Monitoring is an Integral part of maintenance. Requirement for a monitoring solution for cloud are totally different from legacy and virtualized Monitoring Environment .There are many third party solutions are available for cloud monitoring. But there is lack of standard model which covers all required parameter needed to be covered in solution so that an exhaustive report can be produced for service provider. This paper is intend to provide a brief introduction of cloud computing with cloud monitoring fundamental and its requirement including EUCALYPTUS an open source software framework and other related framework needed to implement cloud solution.

***Index Terms*— AmzoneEC2, Eucalyptus, GoGrid, Microsoft Azure, Monitoring,OpenNebula, RackSpace**

I. INTRODUCTION

Cloud computing can be defined as style of computing in which dynamically scalable and often virtualized resources are provided as a service over the Internet. Users need not have knowledge of, expertise in, or control over the relevant technology infrastructure in the cloud that supports them Cloud computing is evolutionary in nature where IT is moving from being deployed on independent, dedicated set of infrastructure to shared infrastructure, this shift brings a need for a revolutionary and complex thinking when it comes to monitoring and managing the cloud infrastructure and applications. This is essential to keep the application management as simple as it used to be on dedicated infrastructure.

We are considering, a cloud is a group of machines configured in such a way that an end-user can request any number of virtual machines (VMs) of a desired configuration. The cloud will spawn these VMs somewhere on the physical machines that it owns. The word "cloud" in this context is meant to convey the semi-ethereal nature of these VMs. The end-user neither knows nor cares where exactly these VMs are physically located or the configuration of the underlying hardware, so long as they can access their bank of properly configured VMs. This kind of setup is ideal for applications where a specific hardware configuration is needed or users only occasionally need the high compute capacity. However, commercial cloud services charge, by the hour, for CPU time. In some settings, such as a large organization with many users, it might be more cost effective for the organization to purchase hardware to create its own private cloud. This is where open-source cloud frameworks such as Eucalyptus, OpenNebula and Nimbus enter the picture. These software products are designed to allow an organization to set up a private group of machines as their own cloud. In this work, we analyze these three open-source cloud frameworks and other commercial frameworks. We selected these three because they represent three different points of interest in the design space of this particular type of open-source cloud.

there are three types of the cloud.

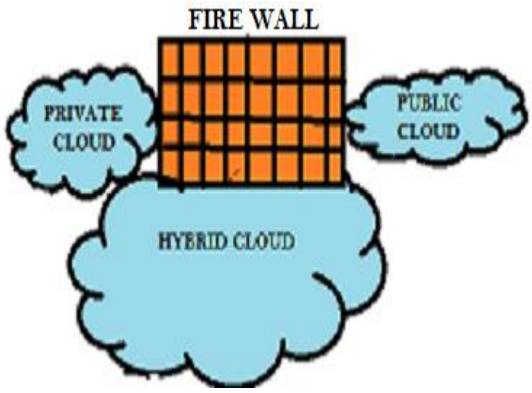

Public Cloud

This is based on standard cloud computing model, in which service provider offers its applications, storage, resources to the general public .It may be free or as pay-per-use bases depends on service provider.

The benefits of the public cloud services are as follows

- It provides easy and inexpensive set up, due to all the hardware and application maintenance costs are taken by the provider.
- It provides the scalability
- You have to pay only for what you need ,so no wastage of resources.

Private Cloud

It is Internal cloud for particular organization, which is set up according to organization or corporate need and provides its services to limited number of users.

Manuscript Received on January 2013.

Mr.Hemanshu A Patel, Department of Information technology, Gujarat Technological University ,Shantilal Shah Engg. College, Bhavnagar, Gujarat, India.

Prof. Arvind D Meniya,Department of Information technology, Gujarat Technological University ,Shantilal Shah Engg. College, Bhavnagar, Gujarat, India.

The specific organization's or company's employee can only get access and it will be accessible only within organization's premises and by authenticating each and every user, it is not open to all.

Hybrid Cloud

This types of cloud are combination of both public as well as private cloud. Most of the commercial use is influenced by this type of cloud, in which some resources are handled and provided by the organization and some are from external.

II. NEED FOR MONITORING IN CLOUD COMPUTING

Both provider and clients are the beneficiaries of monitoring. Cloud providers have to monitor the current status of allocated resources in order to handle future requests from their users efficiently and to keep an eye on malicious activity of users by identifying anomalous usage behavior. Monitoring is also beneficial to the end-users since it helps them to analyze their resource requirements, and ensure that they get the requested amount of resources they are paying for. Also, it enables them to know when to request for more resources, when to relinquish any underutilized resources, and what proportion of various physical resources are appropriate for the kind of applications they are running.

Consider an example of a web server. The usage pattern of a web server depends on various factors. One such factor is the time of the day. For example, a server hosting a banking website is likely to have more hits during day-time when most of the transactions take place instantaneously rather than at night-time. Similarly, a web server hosting news is likely to have more hits on the occurrence of some unusual event like a Tsunami. Web servers need to maintain sufficient resources so as to provide uninterrupted service to end users even during peak usage in that cases cloud monitoring plays an important role.

III. MONITORING IN COMMERCIAL CLOUD

Amazon EC2

Amazon EC2 provides a service called CloudWatch that allows monitoring of other Amazon services like EC2, Elastic Load Balancing and Amazon's Relational Database Service. The monitoring information provided to a Cloud client by the CloudWatch service is strictly related to the Cloud client's virtual platform. The CloudWatch service collects the values of different measurement types from its targets and stores them implicitly for a period of two weeks. Cloud Watch is actually a generic mechanism for measurement, aggregation and querying of historic data. All the measurements are aggregated over a period of one minute.

CloudWatch service is charged separately with a single price per hour, regardless of the resource that is being monitored. Recently, Amazon changed the basic monitoring plan to be free of charge. This includes collection of values every five minutes and storage of these values for a period of two weeks. A detailed monitoring plan is also available that offers value collection at a rate of once per minute and is charged per hour of instance whose resource values are collected.

Microsoft Azure

Monitoring applications deployed on Microsoft Windows Azure, the application developer is given a software library that facilitates application diagnostics and monitoring for Azure applications. This library is integrated into the Azure SDK. It features performance counters, logging, and log monitoring. Performance counters are user-defined and can be any value related to the Cloud application. The logging facilities of the library allow tapping into: Application logs dumped by the application, Diagnostics and running logs, Windows event logs, IIS logs and failed request traces and Application crash dumps. There is no automatic monitoring mechanism for web roles and worker roles running on Microsoft Azure. .

There is no direct fee associated to using them, but there is a fee for storing information in a non-volatile persistence storage service and also in querying that storage service.

GoGrid

GoGrid runs under the name of GoGrid Exchange and presents third-party services that can prove useful to Cloud clients. These services include third-party packages that provides monitoring ranging from platform security monitoring to resource usage monitoring and database monitoring. These services also include the possibility of configurable alerts based on the values of the monitored measures.

RackSpace

The Rackspace offers monitoring capabilities at the whole application level for fixed parameters that include used compute cycle count, used bandwidth and storage. Recently, RackSpace has acquired CloudKick, a multi-cloud virtual platform management tool. CloudKick has a broad range of monitoring features for virtual machines. These include different monitoring metrics from low-level metrics like CPU / RAM / disk utilization to high-level metrics like database statistics, HTTP / HTTPS and others. The monitoring metrics can be extended by custom plug-ins that are able to monitor anything that the user defines. Measured data can be presented in raw form or aggregated by user-defined means. For data visualization, a real-time performance visualization tool is also provided. CloudKick also has alerts feature that have a configurable trigger and repeat interval. The alert prompt can be sent by SMS, email or HTTP.

IV. MONITORING IN OPEN SOURCE CLOUD

Due to the growth of cloud computing, there are several solutions in this area. Some of them are as follows.

Nimbus

Nimbus features a system of Nagios plugins that can give information on the status and availability of the Nimbus head node and worker nodes, including changes of the virtual machines running on the worker node.

Published By:
Blue Eyes Intelligence Engineering
& Sciences Publication

Also, there is active work being done around building a higher level tool that is able to monitor deployed virtual machines and compensate for stress points by using monitor information from sensors and configurable policies.

Eucalyptus

Eucalyptus has introduced monitoring capabilities in version 2.0 for the running components, instantiated virtual machines and storage service. This is done by integrating Eucalyptus monitoring into an existing and running monitoring service. Currently, monitoring has been integrated with Ganglia and Nagios. In Eucalyptus this is done by means of scripts that update the configuration of the running monitoring service to also monitor Eucalyptus components and virtual machines.

As alternative solutions to achieving monitoring at a hardware level, one can employ one of the monitoring systems that have been designed and used in grid environments. We can also use monitoring solution offered by third party providers. Given that Eucalyptus implements the same management interface as Amazon EC2 does, it is relatively easy to find such commercial services. Among the commercial solutions available, we can enumerate the following as compatible with Eucalyptus.

RightScale offers a Cloud management environment for Cloud clients. This management environment also includes a real-time platform monitoring sub-service.

enSTRATUS provides a set of tools for the provisioning, managing and monitoring of Cloud infrastructures with applicability to private and public Cloud platforms.

Makara offers a PaaS (Platform as a service) service on top of an IaaS(Infrastructure as a service) Cloud. They focus on the deployment, management, scaling and monitoring of PHP and Java applications deployed on Clouds. With respect to monitoring, Makara supports real-time and historical performance monitoring.

OpenNebula

The built-in monitoring capabilities of OpenNebula mainly focus on the Cloud provider's interest in the physical resources. This functionality is found in the OpenNebula module called the Information Manager. The Information Manager works by using probes to retrieve information from the cluster's nodes. The probes are actually custom scripts that are executed on the physical nodes and output pairs of Attribute Value on their standard output. Currently, the probes are focused on retrieving only information that underlines the state of the physical nodes and not its running virtual machines (CPU load, memory usage, host name, hypervisor information, etc. The OpenNebula community recommends using a service manager tool that is a separate entity from OpenNebula.

Alternatively, we can also turn to cluster monitoring solutions that come from the open-source world, some of which are the result of long research endeavors. While still under development, the next version of the Information Manager is based on the Ganglia multi-cluster monitoring tool.

ACKNOWLEDGMENT

Hemanshu A Patel wishes to acknowledge Prof.Arvind D. Meniya, and all the staff of Department of Information Technolgy,Shantilal Shah Engineering College, for their kind support throughout the work. Prof. Arvind D.Meniya wishes to acknowledge his family, and all the staff at L. D. College of Engineering.

REFERENCES

1. A Survey on Open-source Cloud Computing Solutions Patrícia Takako Endo , Glauco Estácio Gonçalves, Judith Kelner, Djamel Sadok.Universidade Federal de Pernambuco – UFPE
2. OpenNebula Tutorial, Constantino Vázquez Blanco Borja Sotomayor , DSA-Research.org Distributed Systems Architecture Research Group ,Universidad Complutense de Madrid
3. Cloud Computing with ***Nimbus***. XtreemOS Summer School 2009. Oxford, September 2009. Kate Keahey keahey@mcs.anl.gov,www.nimbusproject.org/files/nimbusxtreemOS_Sept2009.pdf
4. The Eucalyptus Open-source Cloud-computing System Daniel Nurmi, Rich Wolski, Chris Grzegorczyk Graziano Obertelli, Sunil Soman, Lamia Youseff, Dmitrii Zagorodnov Computer Science Department University of California, Santa Barbara Santa Barbara, California93106 open.**eucalyptus**.com/documents/ccgrid2009.**pdf**
5. A Performance Guarantee Approach for Cloud Applications Based on Monitoring Jin Shao, Qianxiang Wang School of Electronics Engineering and Computer Science Peking University Beijing, China.
6. Amazon Elastic Compute Cloud: User Guide Amazon Web Services Copyright © 2012 Amazon Web Services, Inc. and/or its affiliates
7. Windows Azure™ Security Overview By Charlie Kaufman and Ramanathan Venkatapathy
8. Cloud Infrastructure Service Management – A Review Department of Computer Science and Engineering, SCT Institute of Technology, Visvesvaraya Technological University Bangalore, Karnataka, India
9. RACKSPACE LAUNCHES CLOUD MONITORING TO HELP COMPANIES PROACTIVELY TRACK THE HEALTH OF THEIR CLOUD AND WEB INFRASTRUCTURE HOSTED POSTED AUGUST 22ND, 2012 BY RACKSPACE.HTTP://WWW.PRESSRELEASEPOINT.COM/RACKSPACE-LAUNCHES-CLOUD-MONITORING-HELP-COMPANIES-PROACTIVELY-TRACK-HEALTH-THEIR-CLOUD-AND-WEB-INFR