

Bachelor's Thesis (UAS)

Degree Programme in Computer Engineering

Information Technology

2012

Abayomi Ayoola, O.

ON-PREMISE CLOUD COMPUTING

DEPLOYING A PRIVATE CLOUD
INFRASTRUCTURE WITH UBUNTU SERVER
10.04 ENTERPRISE CLOUD (UEC)



TURUN AMMATTIKORKEAKOULU
TURKU UNIVERSITY OF APPLIED SCIENCES

BACHELOR'S THESIS | ABSTRACT

TURKU UNIVERSITY OF APPLIED SCIENCES

Computer Engineering | Information Technology

Autumn 2012 | 72

Instructor(s): Junnila Lassi

Abayomi Ayoola, O.

ON-PREMISE CLOUD COMPUTING: DEPLOYING A PRIVATE CLOUD INFRASTRUCTURE WITH UBUNTU SERVER 10.04 ENTERPRISE CLOUD (UEC)

Described in this thesis is the deployment and configuration of a highly customizable on-premise (private) cloud infrastructure, Ubuntu Enterprise Cloud (UEC). UEC is an enterprise private cloud platform built on Eucalyptus, an open source enterprise cloud application. The thesis is in two parts, with the first part discussing the theoretical baseline of Cloud computing such as definitions and platforms while the second part explains in details the required implementation steps.

The research study covers Cloud computing: characteristics, models, architecture, risks and benefits and platforms with focus on Eucalyptus.

The scope of the thesis was the implementation of an on-premise cloud service on Ubuntu server 10.04.

KEYWORDS:

Cloud Computing, Eucalyptus, Virtualization, Private Cloud

FOREWORD

Gratitude to God, the most High for the gift of life all through the course of my studies. To my mother, an angel God has placed in my life and also my father, the very strong bar behind my wall. To all my colleagues at the University who have in one way or the other contributed to success of my programme, I say a very big thank you.

Finally, I would like to say that I am grateful to the following lecturers for their impact on me: Poppy Skarli, you were more than just a teacher, Vesa Slotte, you have transferred a huge amount of knowledge only if you knew, Tero Virtanen, thank you for all those times you shared with me.

December 11th 2012, Helsinki

Abayomi Ayoola, O.

TABLE OF CONTENTS

1 INTRODUCTION	1
1.1 Motivation	1
1.2 Thesis Objective	2
1.3 Organization & Structure	2
2 CLOUD COMPUTING	4
2.1 Definition	4
2.2 Brief History	9
2.3 Architecture	11
2.4 Economic Importance or Impact of Cloud Computing	15
2.5 Risks Attributed to Cloud Computing	18
3 CLOUD COMPUTING SOLUTIONS	23
3.1 Cloud Computing Solutions	23
3.2 Commercial Solutions	23
3.3 Open Source Solutions	29
4 EUCALYPTUS	33
4.1 Variants	33
4.2 Modules	34
4.3 Networking Modes	35
4.4 Management & Administration Tools	37
4.5 Architecture	39
5 SYSTEMS & CONFIGURATION REQUIREMENTS	40
5.1 Systems Design	40
5.2 Systems Requirements	40
5.3 Network Configuration	41
6 INSTALLATION	44
6.1 Building the Cloud Controller	44
6.2 Building the Node Controller	56
6.3 Post-installation Configurations	57
6.4 Installing Images	63
7 INSTANCE MANAGEMENT	67
7.1 Security Group	67
7.2 Keypair	68

7.3 Running Instances	68
8 SUMMARY	70
REFERENCES	71
FIGURES	
Figure 1. Traditional computing environment	8
Figure 2. A Cloud Computing environment	9
Figure 3. Evolution of Cloud Computing	10
Figure 4. Visual Model of NIST Working Definition of Cloud Computing	13
Figure 5. Server stack comparison between on-premise infrastructure, IaaS, and PaaS	14
Figure 6. Levels of abstraction in Cloud Computing	14
Figure 7. Example of Eucalyptus implementation	39
Figure 8. Test network physical topology	42
Figure 9.1 Booting the Ubuntu Server 10.04 from CD	44
Figure 9.2 Choosing a name for Server 1	45
Figure 9.3 No pre-installed cloud controller on the network	45
Figure 9.4 Roles of the front-end server decided	46
Figure 9.5 The entire HDD partitioned	47
Figure 9.6 The only partition available to the installer	47
Figure 9.7 The entire disk space partitioned	48
Figure 9.8 Confirming partition for format	49
Figure 9.10 Installing the base system	49
Figure 9.11 Setting up users full name	50
Figure 9.12 Setting up username	51
Figure 9.13 Setting up users password	51
Figure 9.14 No need for encryption	52
Figure 9.15 No need for proxy information	52
Figure 9.16 Automatic updates are not required	53
Figure 9.17 Cluster name declared	53
Figure 9.18 Input a pool of public IP addresses	54
Figure 9.19 A Fully Qualified Domain Name is required	54
Figure 9.20 Requesting installation of GRUB	55

Figure 9.20 Front-end installation complete	55
Figure 10. Installing the Node controller	56
Figure 10.1 Giving the node controller a hostname	56
Figure 10.2 The role of Server 2	57
Figure 11. CLC web interface	59
Figure 12. Entering the default user credential	60
Figure 13. The new admin credentials are entered after login in the first time	61
Figure 14. Downloading credential certificate	62
Figure 15. Error message displayed on store tab	63
Figure 16. The images finally displays and being installed	65
Figure 17. Registered Images under “Images” tab	66

TABLES

Table 1. Selected definitions of Cloud Computing by industry analyst firms	5
Table 2. Reported outages in Amazon, Google, Microsoft and Salesforce.com public clouds between 2008 and 2009	24 - 25
Table 3. Comparison of Commercial Cloud Solutions	29
Table 4. Comparison of Open Source Cloud Solutions	32
Table 5. Front-end System Requirements	40
Table 6. Node Controller System Requirements	41
Table 7. IP addressing of the cloud infrastructure	43

ACRONYMS, ABBREVIATIONS AND SYMBOLS

UEC	Ubuntu Enterprise Cloud
VMs	Virtual Machines
SaaS	Software as a Service
ACM	Association for Computing Machinery
SLA	Service Level Agreement
PAYG	Pay As You Go
CRM	Customer Relationship Management
PDA	Personal Digital Assisstant
PaaS	Platform as a Service
IaaS	Infrastructure as a Service
EC2	Elastic Cloud Compute
S3	Simple Storage Service
SME	Small and Medium Enterprises
VDI	Virtual Desktop Infrastructure
DRS	Distributed Resource Scheduler
REST	REpresentational State Transfer
AMI	Amazon Machine Image
API	Application Programming Interface
KVM	Kernel-based Virtual Machine
SOAP	Simple Object Access Protocol
DHCP	Dynamic Host Configuration Protocol

1 Introduction

1.1 Motivation

Cloud computing in recent times has generated a lot of buzz and media attention, making analysts and technology enthusiasts positive about the very many opportunities it could and offers businesses and organizations.

IT departments under dynamic business environments today are consistently under pressure to reduce costs and support business IT scalability with a high level of flexibility thus the need for cloud computing. Cloud Computing is not an entirely new technological innovation. It is borne out of virtualization as its bedrock. It allows for increased scalability and a more efficient use of business technology process based on virtualization of IT resources.

This un-rather new area of technology has opened up the possibility of starting and running a new business IT process with almost no initial capital expenditure since it is based on “pay-per-use utility model”, simply pay for what you use. This area of technology has been well received due to the sustainability possibilities achievable as well as cost reduction.

The idea to implement this thesis was derived from the positive attitude and influence Cloud Computing has had on enterprises, how tech start-ups run their business processes and provide services to their consumers by maximizing its use as well as Cloud Camps organized by the author in cities around Africa.

This first part of this thesis will concentrate on the understanding of the fundamentals of Cloud Computing, its architecture and benefits to enterprises, cloud computing solutions and or platforms for on-premise and commercial services, focusing on Eucalyptus as a platform of choice for on-premise Cloud infrastructure. In the second part, Ubuntu Enterprise Cloud is deployed to show the practicality involved in setting up a private cloud environment.

1.2 Thesis Objective

The main aim of this thesis work is to create a private (on-premise) Cloud environment using an open source platform, Ubuntu Enterprise Cloud (UEC), a highly customizable and configurable cloud platform powered by Eucalyptus.

1.3 Organization & Structure

This thesis work is divided into two parts: the theoretical baseline, which deals with the theoretical aspects of cloud computing such as definition, history and architecture while the second part, the practical implementation deals with the deployment of Eucalyptus in a lab environment. Each chapter describes or profer answers to specific questions. The following sections are a brief overview of each of the chapter:

1.3.1 Chapter 1 Introduction

This is the introductory chapter of the thesis which explains the motivation towards the thesis topic with a brief description of the thesis aim and structure or overview of the thesis work.

1.3.2 Chapter 2 Cloud Computing

This chapter elaborates the various definitions of Cloud Computing, brief explanation for proper understanding by the average reader and discussing its history and architecture as well as its economic importance. This is so that the reader could have a fundamental understanding of the subject topic.

1.3.3 Chapter 3 Cloud Computing Solutions

There are various solutions for Cloud Computing infrastructure varying from commercial to open source. These and specific products are what this chapter discusses, citing examples with pros and cons.

1.3.4 Chapter 4 Eucalyptus

This chapter addresses Eucalyptus elaborately, its features, variants, modules networking modes and management & administration tools.

1.3.5 Chapter 5 System & Configuration Requirements

The system requirements and network configuration for deployment are discussed here in detail.

1.3.6 Chapter 6 Installation

The practical implementation is necessary to validate the claims as discussed in previous chapters. It is expected that the audience will gain practical knowledge in deploying Eucalyptus. The aim of this chapter is to achieve that purpose.

1.3.7 Chapter 7 Instance Management

This is an after installation step required for instantiating VMs. Different ways of instantiating and running instances, security groups and key pairs are discussed in this chapter.

1.3.8 Summary & Conclusion

This aim to provide the reader with the summary of the work as presented in the thesis.

2 Cloud Computing

2.1 Definition

2.1.1 Theoretical Definition

With the entire buzz generated in the media concerning the subject topic, there is still no generally acceptable definition for Cloud Computing. In an attempt to give a meaningful insight into the definition of Cloud Computing, analysts, academics, practitioners and IT companies have deduced a number of definitions. Table 1.1 shows a number of definitions by selected industry analyst firms.

Table 1. Selected definitions of Cloud Computing by industry analyst firms.

Source	Definition
Gartner	“a style of computing in which massively scalable IT-related capabilities are provided “as a service” using Internet technologies to multiple external customers” (Gartner 2008b)
IDC	“an emerging IT development, deployment and delivery model, enabling real-time delivery of products, services and solutions over the internet (i.e. enabling cloud services)” (Gens 2008)
The 451 Group	“a service model that combines a general organizing principle for IT delivery, infrastructure components, an architectural approach and economic model –basically, a confluence of grid computing, virtualization, utility computing, hosting and software as a service (SaaS)” (Fellows 2008)
Merill Lynch	“the idea of delivering personal (e. email, word processing, presentations) and business productivity applications (e.g sales force automation, customer service, accounting) from centralized servers” (Merill Lynch 2008)

A conventional characteristic with these definitions is that they all diverge from a singular perspective: the users' perspective focusing on their experience. Thus, the core feature of Cloud Computing is the provision of IT infrastructure and software as a service in a scalable way.

Likewise, academics have also attempted to define and describe the core features of Cloud Computing, one of such definitions include that of the Berkeley RAD Lab.

“Cloud computing refers to both the applications delivered as a service over the Internet and the hardware and systems software in the datacenters that provide those services. The services themselves have long been referred to as Software as a Service (SaaS). The datacenter hardware and software is what we will call a Cloud. When a cloud is made available in a pay-as-you-go manner to the general public, we call it a Public Cloud; the service being sold is Utility Computing. We use the term Private Cloud to refer to internal datacenters of a business or other organization not made available to the general public. Thus, Cloud Computing is the sum of SaaS and Utility Computing, but does not include Private Clouds. People can be users or providers of SaaS or users or providers of Utility Computing” (Armbrust et al. 2009)

In addition to this, there were twenty distinct definitions in a white paper, “*A Break in the Clouds: Towards a Cloud Definition*” published for the Association for Computing Machinery (ACM) Communications Reviews but the authors summarized some of the common features within those into the following definition.

“Clouds are a large pool of easily usable and accessible virtualized resources (such as hardware, development platforms and/or services). These resources can be dynamically reconfigured to adjust to a variable load (scale), allowing also for optimum resource utilization. This pool of resources is typically exploited by a pay-per-use model in which guarantees are offered by the Infrastructure Provider by means of customized SLA's.” (Vaquero, Rodero-Merino, Caceres & Lindner, 2008)

Stanoevska-Slabeva et al. (2010) argued further that, “Summarized *scalability, pay-per-use utility model* and *virtualization*, as the feature set that would most closely resemble a minimum definition of Clouds. However, while the definition of Vaquero et al. (2008) summarizes other definitions with respect to the physical layer very well, it does not stress the integration of hardware with Software-as-a-Service in sufficient manner. (Vaquero et al. 2008)”

An industry analyst thinks otherwise about this definition, in his analysis of the white paper, John Rhoton stated that, “This is a comprehensive definition but it hides the primary discovery of the paper that no single theme seemed to permeate all of the definitions.” In other words, not one single definition was good enough to encapsulate all of the definitions as defined in the research paper.

Another industry analyst, Reese (2009), described that the Cloud can be both application and infrastructure when he noted that, “The [Cloud] service is accessible via a web browser (non-proprietary) or web services API; Zero capital expenditure is necessary to get started; you pay only what you use as you use it.” John Rhoton’s reactions to this definition is, “This definition describes reacts to Cloud Computing from the perspective of a provider, it is deduced from here that the major Cloud component is the datacenter as it contains the hardware resources for computing and storage including software offering using the pay-as-you-go (PAYG) model.”

Finally, the National Institute of Standards and Technology, USA, defined Cloud Computing as, “*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.*” This definition seems to be standard in that it is accepted by a number of academic, practitioners and analyst and appears to be making round than any other definition of Cloud Computing.

Stanoevska-Slabeva et al. (2010) concluded that, “All definitions illustrate that Cloud Computing is a phenomenon that comprises a number of aspects and is related to a new paradigm of IT (hardware and applications) delivery and deployment. Generally, Cloud Computing concerns the delivery of IT capabilities to external customers, or, from the perspective of a user, obtaining IT capabilities from an external provider, as a service in a pay-per-use manner and over the Internet. Further, scalability and virtualization are very often seen as key characteristics of Cloud Computing (e.g., Foster et al. 2008, Sun 2009a, Vaquero et al. 2009). Scalability refers to a dynamic adjustment of provisioned IT resources to variable load, e.g., increasing or decreasing number of users, required storage capacity or processing power. Virtualization, which is also regarded as the cornerstone technology for all Cloud architectures (e.g., Sun 2009), is mainly used for abstraction and encapsulation (Foster et al. 2008). Abstraction allows unifying raw compute, storage, and network resources as a pool of resources and building resource overlays such as data storage services on top of them

(Foster et al. 2008). Encapsulation of applications ultimately improves security, manageability, and isolation (Foster et al. 2008). Another important feature of Clouds is the integration of hardware and system software with applications. Both the hardware and systems software, or infrastructure, and the applications are offered as a service in an integrated manner.”

2.1.2 Cloud Computing Explained

Compared to the traditional computing system where servers are housed locally within the Enterprise data center or facility, cloud computing differs in that computing resources or services such as storage and application hosting are provisioned by third party service companies, e.g., AWS and are accessed remotely via a web browser leaving the management issues to the service provider. It helps businesses become more lean and capable of reacting faster to business challenges and opportunities. Figure 1 and 2 visually describe the contrast between a traditional computing environment where servers, applications, storage etc. are hosted within the company’s IT infrastructure and a cloud-based environment where computing resources are accessed remotely via the Internet as provisioned by a third party provider.

A good illustration in explaining Cloud computing further is that of Mr. A who just moved into a city and is looking for a place to live. There is one of two things on his mind

1. To build his own house or
2. Rent an apartment

If he decides to build a house, there are number of important decisions he would have to make

- The size of the house
- Future plans e.g raising a family and remodeling
- Where and what should be built into it e.g pool, garage etc

These and other factors that would help in determining the expanse of land that needs to be bought as well as customizing e.g landscaping, lighting and floor. After the house has been completed, he has to do maintenance, e.g., pay water bills, pay electricity bills and pay property tax.

On the other hand, he could choose to rent an apartment from a builder, Mr. B, who built a number of apartment units which can be expanded in the future. In this case, he

can make fewer or less important decisions either to expand or downsize later, but does not have many options in customizing, but the infrastructure for expanding is of high quality and provided by Mr. B. There would be no need to consider maintenance when making plans, that is, paying property tax and other forms of maintenance cost but rent and utility cost.

Translating the above scenario to the subject topic, Mr. A would be an already running organization or a startup whose IT infrastructure is either to be built or rented (cloud).

The first scenario is the organization making a decision on factors required to build its IT infrastructure and what it would require in terms of resources thereby paying for power, in-charge of maintenance and recruitment of technical resources. While the second scenario sees the organization renting the IT infrastructure from a third party company, Mr. B, known as the service provider, in which case the organization pays for what it uses and does not have to bother about how the infrastructure is run as long as it enjoys the service it pays for.

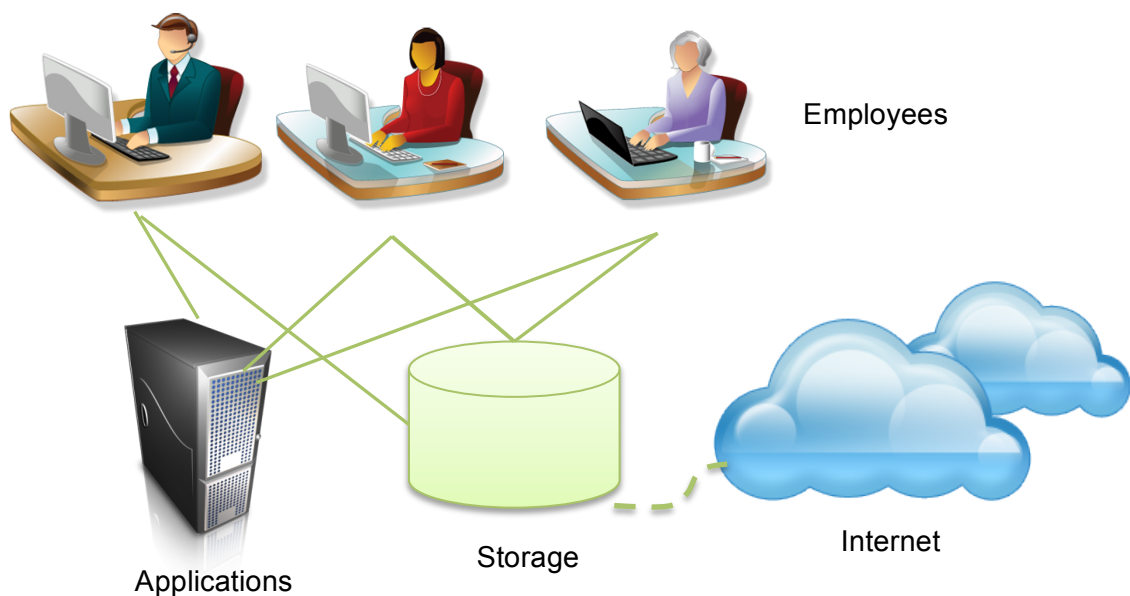


Figure 1. Traditional computing environment.




Figure 2. A Cloud Computing environment.

2.2 Brief History

Cloud Computing is a relatively new term but not an entirely new IT development, rather an extension of the concept of Grid Computing, a concept of applying a large number of computers over a distributed system to a particular, mostly scientific problem.

Having mentioned this, there is the need to distinguish between both Grid Computing and Cloud Computing. Grid Computing leverages the power of several computer resources in parallel to solve a specific scientific problem breaking a huge task into smaller tasks while Cloud Computing leverages on computing resources to deliver services to the end user.

In the nineties, the fundamental thought of virtualization was augmented past virtual servers to higher stage of abstraction—first the virtual platform, embracing storage and web supplies, and subsequently the virtual request for paid job, which has no precise implicit infrastructure. Utility Computing proposed clusters as virtual platforms for computing with a metered enterprise model. More presently, Software as a Service (SaaS) has lifted the stage of virtualization to the application, with an enterprise type of invoicing not by the resources used but by the worth of the application to subscribers.

Precisely in 1999, Salesforce started rendering application services (SaaS) to users over the Internet via their website. Despite its success then, more time was required for it to become widespread.

In 2002, Amazon began Amazon Web Services, providing Cloud services such as object storage and computing. Their major commercial service, Elastic Compute Cloud (EC2) did not start until 2006. The history of Cloud Computing probably would not be complete without mentioning the story of Amazon's CEO, Jeff Bezos, who at the time decided that they started renting out idle compute resources in their datacenters.

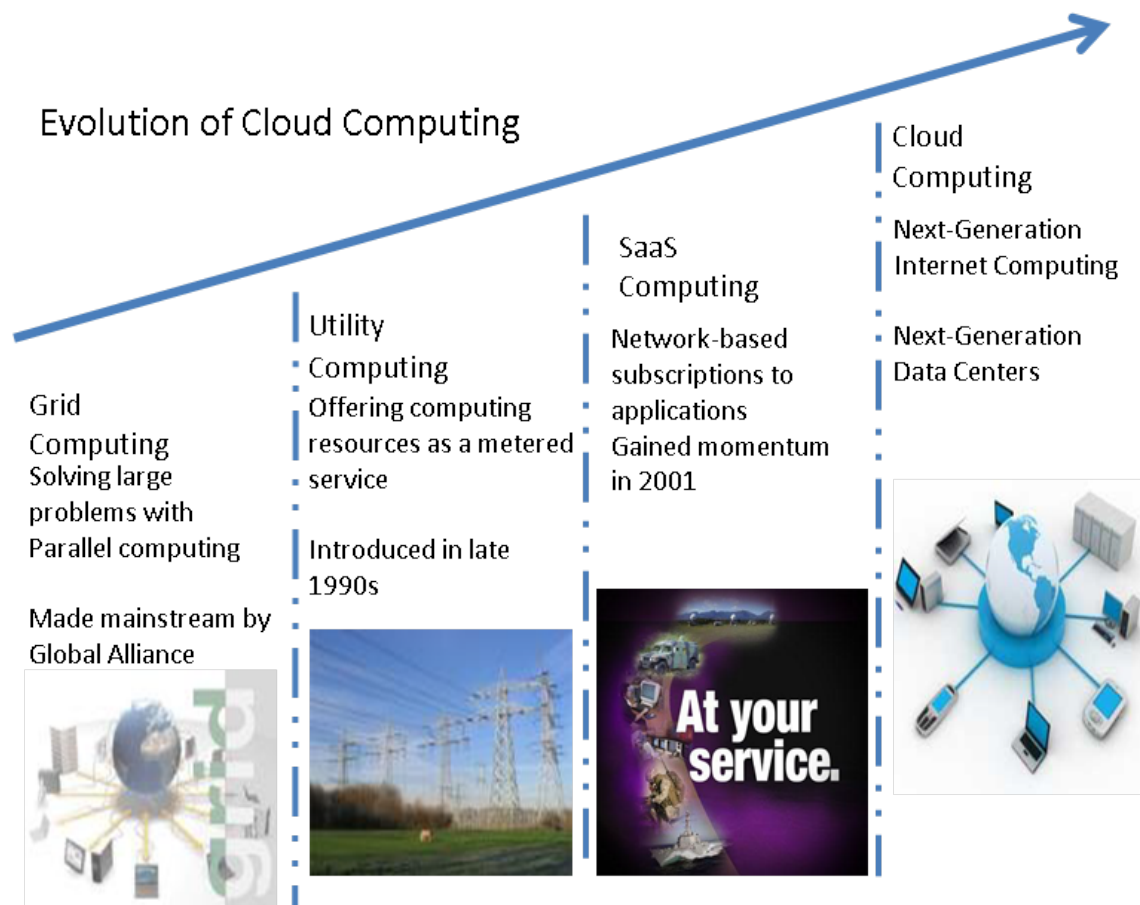


Figure 3. Evolution of Cloud Computing (William 2010)

The year 2009 marked a major turnaround in the evolution of Cloud Computing with Google launching its browser based business and office application known as Google Apps. Big players such as Microsoft, IBM, HP and Oracle have also joined in provisioning of Cloud services such as Windows Azure, Office365, Windows Intune, Microsoft Dynamics CRM, HP Cloud and Oracle Cloud.

2.3 Architecture

Cloud Computing is composed of five essential features, three service models and four deployment models as mentioned by the National Institute of Standards and Technology white paper.

2.3.1 Essential Features

Generally acceptable essential characteristics of Cloud Computing are as follows:

On-demand self-service. Service consumer can procure or enable computing resources such as storage as needed without requiring personnel contact with the service provider, which simply translates to acquiring the usage of computing capabilities without the need of contacting the provider of such service.

Broad network access. These computing resources are available over the network or rather, over the Internet and accessed through various form of standard devices, popularly known as clients just as in a traditional computer networks, such devices are mobile phones, laptops, desktops and PDAs.

Resource pooling. Computing resources such as memory, processing, bandwidth, storage are pooled together by service providers to serve multiple consumers under the multi-tenant model, assigning and re-assigning such resources according to demand. Knowledge of physical location of resources is independent of consumer but isable to specify location at a higher level of abstraction, e.g., country, state.

Rapid elasticity. Resources and capability can be provisioned elastically and rapidly, and in most cases automatically so that manual configurations are not required; this is so that services can scale rapidly outward and rapidly released to scale inward on demand. With this, resources appear to be unlimited and up for purchase in any quantity and at any time.

Measured service. Cloud systems leverage a metering capability to automatically manage and optimize resource usage at level of abstraction suitable to that level of service. Monitoring, control and report is done both ways by providers and consumers thereby providing transparency.

Most enterprises already running varied applications in their datacenters could use cloud computing as a mean of extending their existing infrastructure and as such could be dedicated to certain tasks or processes either already running or new ones.

2.3.2 Service Models

Software as a Service (SaaS). Service provisioning under this model is for users or consumers to use a provider's application running on its infrastructure at a remote location on demand, in that, a single instance of the application is available to multiple end-users. In other cases, different instances of the same application are deployed for different consumers. These applications are accessible via various client devices as stated below:

- Zero Clients. These are the latest computers that include only display devices with every other service located on the server.
- Mobile Clients. These include Mobile phones, Tablets and PDAs.
- Thin Clients are computers without hard drives but rather displays information via an interface as output by the server or IT infrastructure.
- Thick Clients are the traditional computer hardware that connects to the cloud via web browsers such as Firefox, Chrome and Internet Explorer.

The end-users do not manage the infrastructure and resources (operating system, network and servers) used in provisioning this service except for a limited number of user-specific configurations. Examples of such applications in this category include Google Apps, Salesforce CRM, and Microsoft Office 360.

Platform as a Service (PaaS). The service provision on this model is the deployment of applications acquired or developed by end-user onto the infrastructure of service provider using tools and programming languages supported by such. Consumers do not manage or control underlying infrastructures behind the applications but have total control over applications deployed and its hosting environment configurations. Examples of such services include Google App Engine, Force.com, Heroku and Github.

Infrastructure as a Service (IaaS). The provisioning of computing resources such as operating system, storage, servers, processing, bandwidth, datacenter space and sometimes applications are pooled together and made available by service providers to handle tasks and processes. Consumers manage and have a total control over

operating systems, storage and applications deployed. Examples include AWS EC2 and S3, GoGrid and 3Tera.

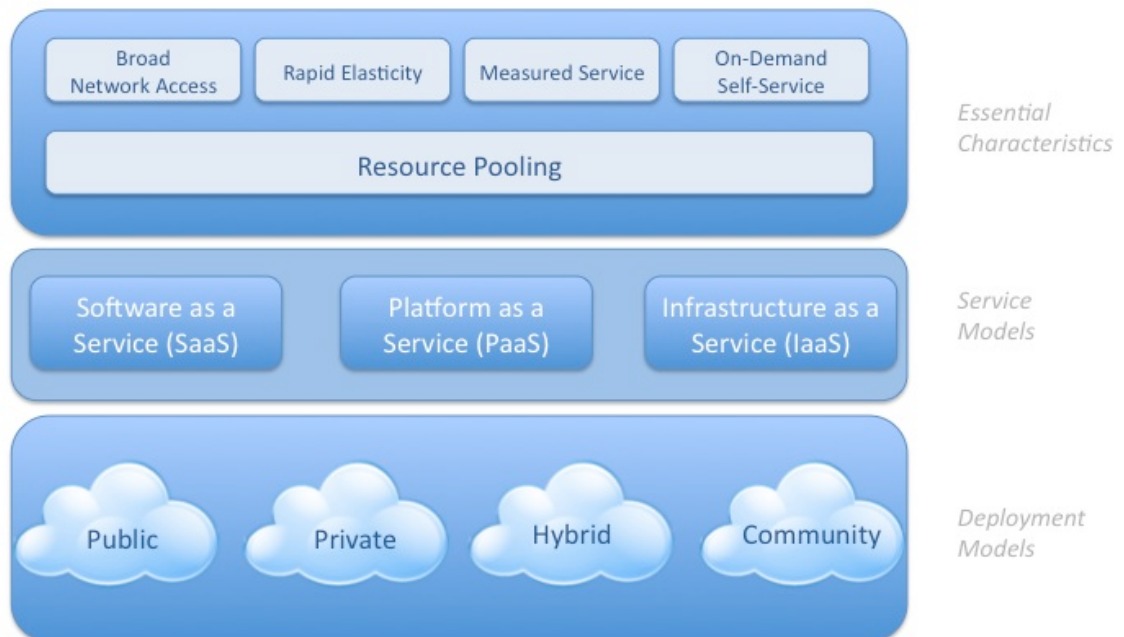


Figure 4. Visual Model of NIST Working Definition of Cloud Computing (Cloud Security Alliance Security Guidance for Critical Areas of Focus in Cloud Computing Version 2.1(2009))

2.3.3 Deployment Models

Private Cloud. A cloud infrastructure built entirely for a single organization and its business units and managed either by its own human resource or third party organization.

Community Cloud. A cloud infrastructure set up for several organizations with common interest and concern. Like private cloud, it can be managed either by member of its staff or third party organization.

Public Cloud. The provisioning of cloud services is rendered by an organization to the public.

Hybrid Cloud. The provisioning of two or more cloud deployment models networked by standards or proprietary protocols that allow portability of traffic such as cloud bursting activity between clouds.

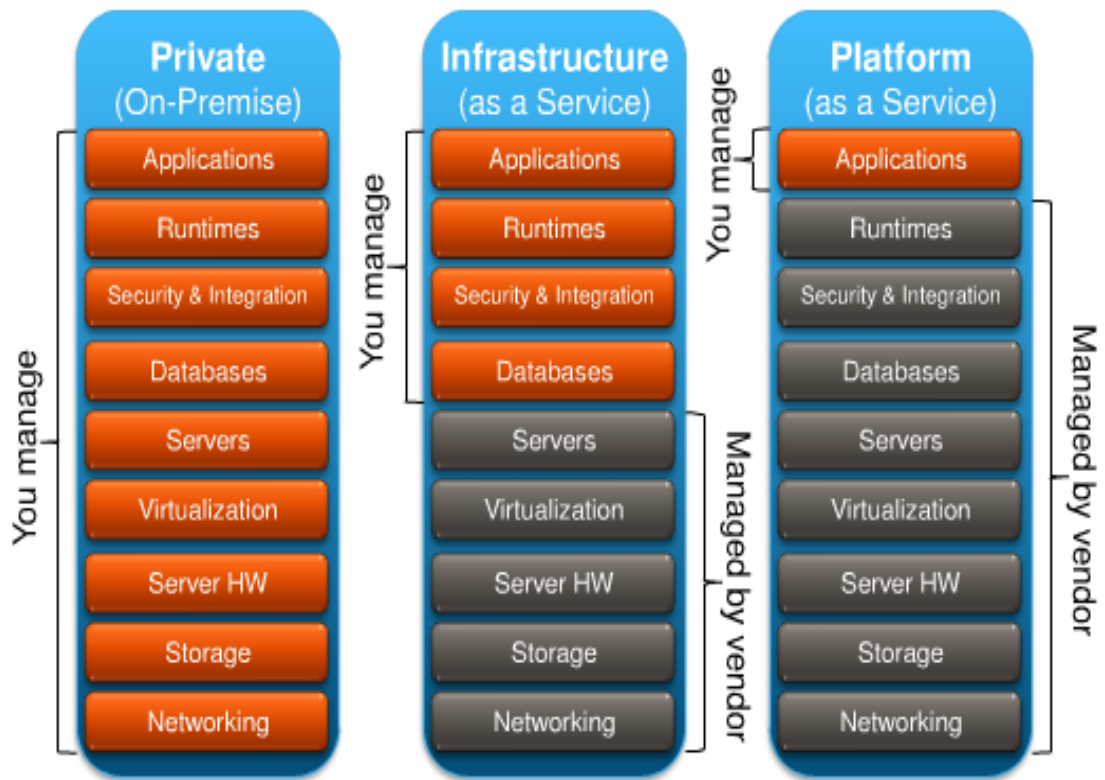


Figure 5. Server stack comparison between on-premise infrastructure, IaaS, and PaaS. (Exploring the Limits of Cloud Computing, 2010)




Figure 6. Levels of abstraction in Cloud Computing (Strategi Consulting, 2009)

2.4 Economic Importance or Impact of Cloud Computing

Financial

The major aim of any business is to reduce capital expenditure, create value and maximize profit. Virtually every process within today's organization requires IT to function efficiently thereby creating the need for more expenditure and a concern for CEO's and Financial Directors not only in terms of financial cost but also human resources for administrative management.

Cloud Computing has contributed to reduction of cost of IT in varied ways, one of which permits infrastructure management and expenditure to become operational rather than capital thereby conserving capital for other purposes. The financial benefit of Cloud Computing is more public cloud-focused than its other models, which is as a result of the ease of use attributed to the pay-per-use utility model.

Pay-per-use model

Cloud computing allows for the unpredictable scalability and elasticity that business IT pass through every now and then. Consumers of public cloud services split the cost of multi-tenant IT infrastructure thereby making affordable the consumption and subscription cost of the services they enjoy. The principle of billing is virtually the same with all service providers but cost models varies. As an example, the cost of using a SaaS depends on the number of users from within the organization; the size and usage of the application you develop determines the cost or cost increases as PaaS usage increases and IaaS costs is for the usage of storage and servers. Small and medium scale enterprises are more likely to benefit from Cloud computing on the long run compared to large enterprises.

Capital vs. Operational Expenditure

The purchase of computing hardware is part of capital expenditure which increases tax burden while payment for Cloud computing services are considered as operational expenditure since the services are rented and does not accumulate assets giving room for cost to be deducted directly from profits, a practice considered important by most businesses. A business looking to come up with its own infrastructure would require a significant amount in cost to acquire hardware and software compared to a business

considering cloud computing which allows for growth and shrinking in IT assets in proportion to the business.

IT Administration and Management

Management and administration of IT infrastructure from within the business generates overhead in terms of cost, which is a significant factor in the total cost of ownership. Computing infrastructure requires technical human resources employed to manage, support, and administrate thereby generating cost overhead. All of these are reduced by cloud computing since these overheads are part of the service provided by the service providers. Compared to SMEs, organizations with large IT business units or departments have fewer issues in terms of overhead cost, with fewer resources and can benefit largely from cloud computing.

Technological

On Demand Scalability

Part of the essential features of cloud computing are on-demand self-service and rapid elasticity. These accelerate the provisioning of IT services.

Accessibility

Cloud services are web-based and accessible via the Internet and as such all that is required is a web browser using variety of devices as described earlier. This allows for it to be accessible anywhere and anytime provided there is an Internet connection.

Future Enabled

This is so in that application updates are automatic costing consumers nothing and making service providers more competitive by running up-to date software.

Operational

Operational benefits of Cloud computing are as follows:

IT Service Management

This makes IT a service providers problem while businesses can focus on their core offerings. The whole essence of using third parties services is to have non-core business offerings handled by professionals in the field especially IT services, which is

often a distraction to businesses with reference to small businesses in particular where resources are limited.

Remote Accessibility

With Cloud computing, business data, applications and storage can be accessed from a remote location thereby giving employees the independence of working remotely regardless of the type of device in use. This helps businesses deal with the issue of premise availability before anything can be done.

Environment

Although cloud computing has often been referred to as “green computing” thereby generating a lot of arguments for and against its environmental benefits, it is argued that the enhanced technology available in the cloud requires more computing resources than the combined resources used apart.

Resource Sharing

Cloud computing has contributed immensely to “Green IT” thereby creating the right atmosphere for energy efficiency.

- With the multi-tenancy model of application sharing by consumers, cloud computing shares a pool of resources.
- Service providers are coming up with better and energy efficient data centers, which reduce pollutions and every form of greenhouse effects.
- Increased server utilization due to server virtualization technology. Fewer items of physical hardware are needed to meet the same larger computing transactions.

Beyond these, increase in internet traffic, data replication in the cloud and demand created by new services have been argued to be non-green besides, “the electricity consumed by cloud computing globally will increase from 623 billion kilowatt hours in 2007 to 1,964 billion kWh by 2020” (Greenpeace, 2010).

Commuting

Nowadays, employees can work from home without traveling to the office on a daily basis thereby reducing the environmental impact of business activities.

2.5 Risks Attributed to Cloud Computing

There are various risks of immense magnitude to be considered before moving to the cloud, otherwise major catastrophes are bound to happen. Such risks include data protection risks; cloud security risks, system outages; data loss and more.

Internal Security Breach

There are bound to be internal security breaches with running business applications on the cloud. According to the 2009 Identity Theft Resource Centre (ITRC) report, there were about sixteen percent of “insider thefts” of reported cases of data breaches in the United States in 2008. This is possible through the use of sniffing tools to steal employee access credentials (usernames and passwords) and rogue employees who continue to have access to a company’s data despite having left the body of staff due to the fact that their access information was not deleted from the system. In order that this is kept in check, governing policies should be put in place for this reason. Three of such policies are

In-house IT process

Most businesses have a governing IT policy in place to check these security breaches. Having a strong password is one of such, which requires that employees use very strong passwords for authentication as well as have an internal process such as a checklist to follow up on activated and deactivated users account within the company’s network.

Two-way authentication

A very good solution to also check these breaches is a two-way authentication, which requires another identifying factor immune to identity theft. Examples include one-time passwords generated by a token; one-time password sent to a user’s mobile phone after entering the usual credentials; biometric authentication such as retina scan etc.

Single Sign-on

Maintaining various passwords by a single user on different systems within the same network can be cumbersome and prone to internal security breaches and so this poses password management problems. There is a number of commercial and open source

federated identity applications to deal with this. SSO allows users to use a set of credentials on multiple cloud systems through a single secured application.

External Security Breach

Businesses should take their data seriously otherwise if it gets into the wrong hands, it could cause more harm than just security problems and as well as legal issues which could go a long way in destroying the company reputation as well as those of its clients. One of the major security issues that can be experienced in the cloud is the compromise of important business data in the event that there is failure in service providers security technology. In this case, service providers should take security more seriously as well as have businesses subscribed to them. The following key security questions are answered.

Security Failures

Beyond the bug found in Google Docs (SaaS) in March of 2009 which led to small part of documents to be seen by unauthorized users (Mazzon, 2009), no significant example of such has been discovered at least as at the time of writing, says much about security concerns in the cloud and how many resources service providers make available for such.

Operational Failures

As discussed earlier, strong passwords are recommended, as weak passwords usually the easiest point of hacking a user's account. A significant example was when a hacker gained access to President Barack Obama's Twitter account in January 2009 (Twitter, 2009). There has been a number of such incidents, which is due to the fact that users use very mild password to protect their accounts.

Data Protection

In order for cloud service consumers to ensure that their data is handled in the most legal way possible, there is the need to be able to check the data handling procedure of the service providers. This is so because data protection breach poses the following risk to businesses

- Loss of data by service providers
- Unauthorized access to data which may lead to litigation

- Hacking or security activities targeted at service providers. e.g., compromising data due to poor internal IT security

As at the time of writing, it is still difficult for consumers to gain access to detailed system access logs and have to depend on wording as specified in contracts or SLA's for compliance on regulations related to data protection. As a precaution, businesses with data stored or moved around the cloud stretching beyond international borders should ensure that the country hosting the data adheres to data protection acts relating to that of the consumer.

Lock-in

Vendor lock-in has long existed before cloud computing. Analysts have argued that it is unavoidable for consumers, be it any IT solution, cloud computing "on-premise" or "as a service." Little or much less have been achieved for the interoperability between clouds as consumers are still been locked in to particular provider thereby leaving them dependent on the particular service provider for provisioning of services. As a check to this, a number of standard organizations are being founded to work out modalities for the interoperability between clouds. Such industry consortia with different area of proposal include:

- The Cloud Computing Interoperability Forum (<http://www.cloudforum.org/>)
- The Distributed Management Task Force's Open Cloud Standards Incubator (<http://www.dmtf.org/about/cloud-incubator>)
- The Open Cloud Manifesto (<http://www.opencloudmanifesto.org/>)
- The Open Grid Forum's Open Cloud Computing Interface Working Group (<http://www.occi-wg.org/>)

Service providers are continuously under pressure to provide an interoperability tool for consumers who may wish to move to other service provider or back "on-premise." Consumers should always have an exit strategy from a provider in consideration for moving between clouds and back "on-premise."

Service Outage

There has been a number of outage cases in the cloud in recent past due to events usually beyond providers control. The four biggest names in cloud services have been hit with such events and are listed below:

Table 2. Reported outages in Amazon, Google, Microsoft and Salesforce.com public clouds between 2008 and 2009 (Williams, 2010)

Service outage	Date	Duration
Amazon S3	15 Feb 2008	4 hours
Amazon EC2	Apr 2008	1 hour
Amazon S3	20 Jul 2008	8.5 hours
Amazon EC2	11 Jun 2009	7 hours
Amazon EC2 7	9 Dec 2009	1 to 5 hours
Google App Engine	17 Jun 2008	7 hours
Google Gmail	16 Jul 2008	1.5 hours
Google Apps & Gmail	6 Aug 2008	about 15 hours
Google Gmail	11 Aug 2008	1.5 hours
Google Gmail	15 Aug 2008	more than 24 hours
Google Gmail	16 Oct 2008	30 hours
Google Gmail	24 Feb 2009	2.5 hours
Google Gmail	9 Mar 2009	up to 22 hours
Google network	14 May 2009	2 hours
Google App Engine	2 Jul 2009	6 hours
Google Gmail	1 Sep 2009	2 hours
Google Gmail	24 Sep 2009	2.5 hours
Microsoft Windows Live	26 Feb 2008	About 6 hours
Microsoft Hotmail	12 Mar 2009	5 hours
Microsoft Azure	13 Mar 2009	22 hours

Microsoft Sidekick	4 Oct 2009	6 days + total loss of contact data
Salesforce.com	11 Feb 2008	6 hours
Salesforce.com	6 Jan 2009	1 hour

3 Cloud Computing Solutions

3.1 Cloud Computing Solutions

Major IT vendors such as VMware, Microsoft and Citrix have cashed in on the market demands of Cloud computing products and services to develop solutions for it, with VMware churning up the largest market share since 2006. Likewise, the open source community is not left behind in developing free solutions as alternative to the commercial vendors solutions. This section aims to discuss the flexibility; strength and functionality of already existing commercial as well as open source Cloud computing solutions. Discussed here are the three of the major commercial Cloud solutions and three major Open source solutions.

3.2 Commercial Solutions

Big players in the industry have varied Cloud computing solutions inclusive of emerging architectural models ranging from VDI solutions, server virtualization to management applications. The following are some of the popular and newer ones:

VMware vSphere 5 (ESXi)

The VMware vSphere 5, formally known as VMware Infrastructure 4, is an enterprise level, bare-metal hypervisor virtualization product by VMware Inc., built utilizing the VMware ESXi, which replaced the ESX. The VMware vSphere 5, the latest edition of vSphere was released on the 12th of July, 2011. The vSphere is available in three enterprise editions, Standard, Enterprise and Enterprise+, and non-enterprise editions, vSphere Kits (Essential and Acceleration Kits), vSphere Desktop, vSphere Kits for Retail and Branch office and vSphere for SMB's. A free edition is the vSphere Hypervisor.

Features

The VMware vSphere comprises of features that turns a traditional hardware into a shared virtual environment with built-in services for applications. These features are based on such services, Infrastructure Services (Compute, Storage and Network) and Application Services (Availability, Security and Scalability).

Infrastructure Services

Compute resources are virtualized and aggregated across clusters such that various clients can share them. Storage is utilized more efficiently and network is simplified for efficient management across physical and virtual environments.

Compute

- Auto Deploy – This mechanism enables a new host to be deployed more quickly.
- ESXi Convergence – vSphere has been updated from ESX to ESXi, a more advanced and secured hypervisor architecture.
- Support for Apple Products – Supports running OSX 10.6 server on Xserve server as a guest OS.
- New VM format (version 8) – This allows support for:

3D graphics for Windows Aero

USB 3.0

32 vCPUs per VM

1TB of RAM per VM

Storage

- Storage DRS – Initially available for CPU and memory only but presently available for storage, DRS is a vSphere mechanism that continuously monitors resource pool utilization across VM's and allocates unused resources to them according to demand, thereby,

Decreasing downtime and improving reliability.

Increasing energy efficiency by booting down unused servers.

And enabling

Load balancing. Migrating VM's in the event that storage cluster becomes full.

Initial Placement. Placing new VM's on storage with the largest space

Aggregation of physical server resources. Managing CPU's, memory and storage as a shared pool of resources.

- Profile-Driven Storage – Ensuring the right storage for a specific VM is the essence of this feature.
- Storage I/O Control – Helps to improve administration of storage by extending limits in Network File System (NFS).
- Storage API Program – Allows for storage to report status and capability.

Network

- Network I/O Control – Enables the creation of network resource pool, QoS Tagging and, Shares and Limits to traffic types.
- Distributed Switch – This improves the visibility of VM's and enhances monitoring and troubleshooting.

Application Services

Application services allow for easier administration of applications running within VM's thereby giving System Administrators control over application service levels using in-built tools.

Availability

- High Availability - Provides for:
 - Within minutes reboot in the event of hardware or OS failure
 - Fault Tolerance in the event of an unplanned systems failure and
 - Data Recovery without the loss of connectivity to users.
- vMotion – Eliminates scheduling of application downtime thereby allowing live migration of VM's and disk across servers without disruption to users.

Security

- ESXi Firewall – Application level security and network segmentation of users and sensitive business data is possible with vShield Zones and VMsafe.

Scalability

- Larger VM's – Enables

Hot add of CPUs and memory

Hot plug virtual storage and network devices and

Hot extend of virtual disks without disruption or downtime.

Microsoft Hyper-V Server 2008 R2

Hyper-V, codenamed Viridian is Microsofts standalone server and cloud platform based on the hypervisor (Virtual Machine Manager). The Hyper-V was initially released in an installable role in Windows Server 2008 and Windows Server 2008 R2 and later released as a standalone version as Hyper-V Server 2008. Its present release status is release 2 (R2) and runs only on 64 bits CPU architecture.

Features

Flexibility

- Live Migration – Running virtual machines can be moved around Hyper-V physical hosts without disruption to service. This helps increase performance and consolidation without affecting users; thereby reducing costs and increasing scalability.
- Cluster Share Volumes – This feature allows for multiple Servers to access SAN storage using a single consistent namespace for all volumes on all hosts thereby simplifying and enhancing the shared storage usage. This is useful in a cluster configuration scenario as it enables faster live migration and easier storage management.
- Hot Swappable of Storage – This makes it possible to add or remove VHD files and pass-through disks while a virtual machine is still running.
- Dynamic Memory – This help administrators better manage memory resources in that it dynamically allocates memory resources to virtual machines based on demand from each VM.
- Multi-site Cluster for fault tolerance and disaster recovery

Performance

- Performance & Power Consumption – Hyper-V comes with enhancement that reduces VM's power consumption and Second Level Address Translation

(SLAT), which uses features on newer CPU's to improve performance in VM environment.

- USB Support
- Networking Support – The Virtual Machine Queue (VMQ) feature in Hyper-V Server 2008 establishes a dedicated queue on the physical network interface cards (NIC) to use direct memory access (DMA) to place packet's contents directly into VM's memory thereby increasing input/output performance. Another feature is the Jumbo Frame support, which increases throughput and maximum utilization of CPU resources while large files are transferred. The Jumbo Frame supports up to 9014 bytes if physical network allows it.

Scalability & Extensibility

- Hardware Support for Hyper-V VM – Hyper-V support up to sixty-four logical processors.
- Extensibility – Hyper-V API allows ISVs and developers to build and customize tools for the virtualization environment.

Limitation

A major limitation of Hyper-V Server 2008 R2 is Application Failover; this is available in VMs to prevent single-point of failure, a situation where different instances of the same application are running in virtual environments on different hardware in which one is acting as a standby for the other. Another limitation is the support for Virtual Desktop Infrastructure (VDI), an architectural model where a client's operating system runs in server-based virtual machines (VMs) in the data center and interacts with the user's client device such as a PC or a thin client, which is available with the Hyper-V Server 2008 R2 SP1.

Citrix XenServer 6

The Citrix XenServer is a virtualization product by Citrix built on Xen, a research project by Ian Pratt out of University of Cambridge and commercialized by XenSource Inc. until it was bought over by Citrix in October of 2007. Just like the VMware ESXi, the Citrix XenServer is a bare-metal hypervisor virtualization application. The XenServer is available in four different editions, XenServer (Free Edition), XenServer Advanced Edition, XenServer Enterprise Edition and XenServer Platinum Edition. The recent release is the XenServer 6.

Features

Architectural – XenServer 6 is based on Xen 4.1, a more improved Xen hypervisor.

Virtual Appliance Support (vApp) and V2V – This allows for the creation of multi-VM within XenCenter and boot sequenced virtual appliances (vApps) integration between VM's during Site Recovery. Importation and exportation of vApps are made easier with the aid of Open Virtualization Format (OVF).

Microsoft System Center Integration – Beginning with the 6.0 release, it now has the option of managing XenServer hosts and VM's using Microsoft System Center Virtual Machine Manager (SCVMM) 2012 and System Center Operations Manager (SCOM) 2012. This is made possible with a special supplemental pack from Citrix.

XenDesktop Enhancement – XenServer 6.0 includes with it High Definition User Experience (HDX) enhancement for optimized user experience with virtual desktops, GPU Pass-Through and the delivery of other graphical applications via virtual desktops.

Enhanced Guest OS Support – Formal support for Ubuntu 10.04 (32/64) bit. Updated support for Debian Squeeze 6.0 64-bit, Oracle Enterprise Linux 6.0 (32/64) bit and SLES 10 SP4 (32/64) bit and experimental VM templates for CentOS 6.0 (32/64) bit, Ubuntu 10.10 (32/64) bit and Solaris 10.

Improved Performance and Scale – Supported limits for XenServer hosts have been increased to 1 TB, up to 16 virtual processors and 128 GB of virtual memory for VMs.

Disaster Recovery – XenServer 6.0 has capabilities for remote data replication between arrays initially using the StorageLink Gateway Site Recovery feature for this purpose but has replaced with the Integrated Site Recovery.

Comparison

The three major cloud computing platforms mentioned above are compared on a number of essential capabilities such as scalability, performance, high availability and business continuity in Table 3 below.

Table 3. Comparison of Commercial Cloud Solutions				
		Microsoft Hyper-V Server 2008 R2	VMware vSphere 5	Citrix XenServer 6
CAPABILITY				
Host	Logical Processor	64	160	160
	Memory	1TB	32GB	1TB
	vCPU	512	2 048	Undocumented
	Memory Overcommitment	Yes	Yes	Yes
Networking	VLAN Support	Yes	Yes	Yes
	Network I/O Control	No	Yes	Yes
	DHCP Snooping Protection	No	No	No
	Port Monitoring	No	Yes	Per Port Group
Cluster	Nodes	16	NA	16
	Maximum # of VMs	1000	NA	800-960
VM	vCPU/VM	4	8	16
	Memory/VM	64GB	32GB	128GB
	VM/Host	384	512	50-130
Storage	Local Storage	Yes	Yes	Yes
	Storage Encryption	No	No	No
	Storage Replication	No	Yes	No
	Live Storage Migration	No	Yes	No
Business Continuity	Fault Tolerance	No	Yes	Yes
	Disaster Recovery	Yes	Yes	Yes
	Live Migration	Yes	Yes	Yes

Newer versions of these applications have since been released as at the time of writing and features have been improved to meet certain criteria as need arises. Most importantly, usability is mostly determined by business need.

3.3 Open Source Solutions

OpenStack

OpenStack is an open source cloud computing platform pioneered and funded by Rackspace Cloud, a web/cloud hosting company, and National Aeronautics and Space Administration (NASA) in July of 2010 by integrating code from NASA's Nebula platform with Rackspace Cloud files. At the time of writing, there are 3029 individuals ranging from researchers, cloud computing experts, developers and 173 companies that are part of the OpenStack project. Its first official release, code-named Austin was made available on the 21st of October 2010. Its present community release, Essex, which is the fifth release, was made available on the 4th of April 2012.

Features (Essex Release)

The OpenStack project runs on a modular architecture and initially has three components as listed below:

OpenStack Compute (Nova) – This is an application (Cloud computing fabric i.e. the main part of an IAAS system) designed to manage large networks of VMs, helping to create a redundant and scalable cloud computing platform. It provisions the software, control panels, and APIs required in managing a cloud which include running instances, managing networks, and controlling users and project access. OpenStack Compute currently supports a variety of standard hardware configurations and seven major hypervisors.

- Role-based Access Control – This ensures security based on user access, role and security group.
- VM Image Management – This feature supports enabling new standardized VM, import and query of images.
- Management of Virtual Resources – It enables management and utilization of virtual server resources such CPU, memory and storage devices
- vLAN Management – It dynamically allocates IP addresses for provisioning of network capabilities.

OpenStack Object Storage (Swift) – This is an OpenStack project component designed for the sole purpose of creating redundant and scalable data storage using clusters of standardized servers to store huge accessible data. It is a long term storage system for huge amounts of static data that can be retrieved, leveraged updated.

Object Storage ensures data replication and distribution across variety of storage devices.

OpenStack Image Service (Glance) – provision discovery, registration, and delivery services for virtual disk images. The Image Service API server provides a standard REST interface for querying information about virtual disk images stored in a variety of back-end storage. Clients can register new virtual disk images with the Image Service, query for information on publicly available disk images, and use the Image Service's client library for streaming virtual disk images.

Glance supports a multi-format image registry, i.e., it allows uploads of images in a variety of formats, such format include

- Raw image format
- Machine (kernel/ramdisk outside of image, a.k.a. AMI)

- VHD (Hyper-V)
- VDI (VirtualBox)
- qcow2 (Qemu/KVM)
- VMDK (VMWare, others)

Two more components have since been included in the Essex release as listed below:

OpenStack Dashboard (Horizon) - provides administrators and users baseline interface for managing OpenStack services.

OpenStack Identity (Keystone) - provides authentication and implements Identity API for use across OpenStack projects integrating with already existing identity management solution.

CloudStack

CloudStack, just like OpenStack, is an open source cloud computing application originally developed by cloud.com, a company formally known as VMOps and recently acquired by Citrix Systems, precisely on the 12th of July, 2011. Most of CloudStack was released in May 2010 under the GNU General Public License version 3 (GPLv3) until it was recently moved to the Apache License 2.0, and the rest was not released until August of 2011. The current version, CloudStack 3.0 was released in February 2012. Citrix Systems has since discontinued its development as it was donated to the Apache Software Foundation in April 2012.

CloudStack supports various hypervisors such as vSphere, KVM, Oracle VM and Xen server to virtualize OS environments.

Features

Multiple Hypervisor Support – CloudStack cloud deployment supports multiple hypervisor implementations of various types such as Citrix XenServer, vSphere and KVM, allowing users the possibility of choosing the most convenient for them.

Support for massive infrastructure management – This feature allows the management of thousands of servers across geographically located datacenters through a central management server.

Simple and easy-to-use web interface – This enables an easier and simple to use cloud infrastructure via a feature-rich interface deployed on the CloudStack API. It

allows the real-time view of the aggregated computing resources such as CPU and memory thereby giving better control over the implemented cloud.

RESTful API – The CloudStack API with its unique and innovative features implements industry standard APIs on top of it.

Comparison

Compared to the commercial versions of the cloud computing solutions, the open source alternatives are mostly modular in that they were built in modules on top of one another with each carrying out a specific function, yet they possess peculiar features. Table 4 below illustrates this. The features are distinguished based on security, high availability, administration and architecture.

		CloudStack 3.0	Eucalyptus 1.6	OpenStack (Essex)
CAPABILITY				
Architecture	Build	Monolithic	Modular, 5 Modules	Modular, 6 Modules
	AWS API Compability	Yes	Yes	Yes
	Memory Overcommitment	Yes	Yes	Yes
Administration	Configuration	Web Interface, Command Line Interface	Command Line Interface	Multiple Command Line Interface
	Management Console	Yes	No	Yes
	Active Directory Integration	No	No	No
Security	VLANs	Yes	Yes	Yes
	Technique	Baseline	Security Group	Keystone
High Availability	High Availability Component	Yes	Yes, 2x Failover	No, Swift only
Implementation	Installation	Medium	Medium	Difficult
	Availability	Free	Free/Commercial	Free

4 Eucalyptus

Eucalyptus is an acronym for “Elastic Utility Computing Architecture for Linking Your Programs To Useful Systems” and like other leading Cloud platforms, it is an open source application built from the ground up for the implementation of on-premise or private cloud computing infrastructure thereby allowing the enterprise full control of its own cloud. It is entirely developed and sponsored by Eucalyptus Systems, Inc.

Eucalyptus runs entirely on Linux Operating System but supports VMs for Microsoft Windows images as well as a number of OS. One of its main features is the capability for hybrid cloud since it is compatible with Amazon EC2, S3 and IAM API services.

Features

As at Eucalyptus version 3, it has the following features out of the box:

Resource Reporting and Accounting

LDAP and Active Directory Integration

Elastic Resource Orchestration and Management (Placement, Addition and Removal)

Flexible Network Management, Security Groups and Traffic Isolation

Resource Policies and Quotas

4.1 Variants

Eucalyptus version 1.0, which was the initial release, was released as a research project at the University of California, Santa Barbara while the first open source version was released by Eucalyptus Systems, Inc. as version 1.5.2. It has since been scheduled to be released biannually in March and August.

Eucalyptus Community Edition (Ubuntu Enterprise Cloud (UEC))

Eucalyptus had been the core element of the Ubuntu Enterprise Cloud since version 1.5 until recently, precisely in May 2011, when Canonical decided to switch to Openstack as the core element of UEC beginning with Ubuntu 11.10. Although it is still possible to install Eucalyptus on UEC, the default and supported platform would be Openstack.

Eucalyptus Enterprise Edition (Commercial)

Eucalyptus Enterprise Edition is the commercial provisioning of Eucalyptus by Eucalyptus Systems, Inc. It supports more features compared to the community version.

4.2 Modules

The Eucalyptus Cloud is a Linux-based, stack of five modules built by Canonical into the Ubuntu Server Edition making up the UEC. Details are as described below:

4.2.1 Cloud Controller (CLC)

The Cloud Controller (CLC) represents the front-end or the entry point to the entire cloud infrastructure for admin and users. All a cloud infrastructure requires is a single CLC to acts as its front-end. It provides an AWS EC2/S3 compliant external web interface for user interaction and also interacts with the rest of the entire Eucalyptus components on the other side.

The CLC monitors running instances, authenticates users, decides which cluster a requested instance will be allocated and monitors the entire availability of resources in the cloud.

4.2.2 Node Controller (NC)

A Eucalyptus node is a VT-x enabled server capable of running a hypervisor, in the case of UEC; KVM is the default. Eucalyptus does support other hypervisors such as Xen. Each VM running on the hypervisor is referred to as an "instance."

The NC controls the life cycle of instances running on the node. The NC interacts in two ways, with the OS and the hypervisor running on the node with the Cluster Controller (CC). It is also responsible for gathering data received from querying the OS and using it to map out the node physical resources and reporting same to the CC.

4.2.3 Cluster Controller (CC)

Physical hosts or nodes containing NC form a cluster. This cluster is managed by a CC. Cluster Controller (CC) acts as an intermediary between the CLC and the NC, communicating with the CLC on one side and the NC on the other. CC is either

installed on a dedicated host accessible to the nodes and the CLC simultaneously or on the node in the case of a single node cluster. Its functions include:

- Deploying instances on one of its related nodes upon request from the CLC.
- Resource arbitration: deciding in which physical node a given instance will be deployed.
- Managing the networking for the running instances and control the virtual network available to the instances.
- Creating a single virtual network or grouping instances on virtual networks depending on the Eucalyptus networking mode established.
- Collection of resource information from NCs and reporting this information to the CLC.

4.2.4 Walrus Storage Controller (WS3)

The Walrus Storage Controller (WS3) provisions an AWS S3 API compatible persistent simple storage service, which is dependent on REST and SOAP API. Persistent means that it is not attached or linked to a particular instance and as such stored contents persists even when instances are terminated.

The WS3 can be referred to as simple file storage system. Due to its stability, it is required to store machine images, templates used to deploy VM instances, snapshots and serving files using the Amazon S3 API.

4.2.5 Storage Controller (SC)

This is similar to the Amazon Elastic Block Storage (EBS) service in that it provisions persistent block storage for use by instances.

The SC is used mainly for the creation of snapshot or state for later access.

4.3 Networking Modes

Eucalyptus runs in four different networking modes, which have distinguished levels of control and management features and are as discussed in increasing order of configuration features below.

Worthy of note here is that these levels of networking modes are particular to Eucalyptus while other cloud platforms offer similar control and management techniques.

4.3.1 System

Eucalyptus provisions the least networking feature in this mode, assigning random MAC addresses to VM instances and attaching its virtual Ethernet interface to the physical Ethernet interface via the Node's local bridge. VM instances are assigned IP addresses over DHCP just in the same way physical hosts are dynamically assigned addresses. Apart from being the least in features, it is also the easiest to configure in that it does not require any pre-requisites to setup for a running DHCP server in the network.

This mode has similarity to the "Bridged Networking" mode, hypervisors such as VirtualBox offers.

4.3.2 Static

In this mode, administrators are offered more control over VM IP address assignment compared to the System mode. Administrators are required to configure a map of MAC addresses to IP addresses; this enables Eucalyptus to setup a static entry in its controlled DHCP server on the CLC, assign the next available MAC/IP address to an instantiated VM and attach its virtual Ethernet interface to the physical Ethernet interface via the nodes bridge as is done with System mode.

Administrators with access to a pool of MAC/IP pairs are more likely to use this mode of networking, as they do not have to rely on a DHCP server for addressing.

It is important to state here that running Eucalyptus in either System or Static mode disables major functionalities that would allow for a more manageable implementation. Such functionalities include

- VM instances ingress traffic filtering (Security Groups)
- Elastic IPs, i.e., user controlled dynamic IP assignment to instances.
- Network traffic isolation to and from VM instances

- Meta-data service availability i.e. using the URL <http://169.254.169.254/> to obtain instance specific information.

4.3.3 Managed No-VLAN

Managed No-VLAN mode is the default Eucalyptus networking mode and is identical in features to the Managed networking mode in that it supports security groups, which is applicable to inbound traffic (all outbound traffic are allowed) and dynamically assigned IP addressing but it does not provision instance isolation.

4.3.4 Managed

Compared to other modes of networking in Eucalyptus, Managed mode offers the most features as well as coming with the most potential constraints on setup. The main difference between this mode and Managed-NoVLAN mode is that it is able to create multiple VLANs by tagging of packets. The Managed mode allows the administrator to define a large private and unroutable network from which IP addresses are assigned to VM instances created and maintained by a DHCP server running on the CLC with static mappings for each VM.

The concept of security group is a very powerful security tool especially in this mode due to its multiple VLAN capability, which allows for ingress traffic filtering and instance isolation. This is made possible by the CLC dividing the entire block of IP addresses defined and breaks them into segments that are allocated to security groups created by users. This in turn makes it possible for instances in the same security group to use IP addresses from the same segment thereby giving room for applying different set of rules to each VLAN created.

4.4 Management & Administration Tools

4.4.1 Euca2ools

Euca2ools is an open source version of the AWS EC2 API and AMI Tools, a set of management utilities and command line tools used for interacting with web services.

Features

- Query of availability zones (i.e. called “clusters” in Eucalyptus)
- SSH key management (add, list, delete)

- VM management (start, list, stop, reboot, get console output)
- Security group management
- Volume and snapshot management (attach, list, detach, create, bundle, delete)
- Image Management (bundle, upload, register, list, deregister)
- IP address management (allocate, associate, list, release)

4.4.2 Hybridfox

Hybridfox is an open source GUI and Mozilla Firefox browser extension that has the capabilities of allowing

- Launching and management of instances
- Management of Security Groups
- Management of Elastic IPs etc.

Apart from Eucalyptus, it supports AWS, OpenStack, OpenNebula, HP Cloud etc.

4.5 Architecture

The following is a diagram exemplifying the components of Eucalyptus.




Figure 7. Example of Eucalyptus implementation
(<http://yoyoclouds.files.wordpress.com/2011/12/eucalyptus-cloud-setup.png>)

5 Systems & Configuration Requirements

5.1 Systems Design

Compared to the older version of Eucalyptus, the latest version could run entirely on limited computing resources, i.e., a single physical host. Although not a system requirement, Eucalyptus places a limit on setting up all components on a single host such that most networking features are not functional. This mode is referred to as SYSTEM networking mode. Implementing this environment, two physical hosts of 64bits computing architecture were used.

- Front-end Components
Cloud Controller, Cluster Controller, Walrus Storage Controller and Storage Controller
- Node Controller

A multiple host setup is required to fully test Eucalyptus.

5.2 Systems Requirements

For a simple and non-HA testing environment such as the project in discussion, the following is the recommended requirements compared to the available resources.

Table 5. Front-end System Requirements.

	Minimum Requirement	Suggested	Available Resource
CPU	1GHz	2 * 2GHz	2GHz, Duo
Memory	1GB	4GB	2GB
HDD	40GB	200GB	150GB
Networking	100Mbps	1000Mbps	1000Mbps

Table 6. Node Controller System Requirements.

	Minimum Requirement	Suggested	Available Resource
CPU	1GHz, VT-x enabled	VT-x, 64-bits, Multicore	2GHz, Duo
Memory	1GB	4GB	2GB
HDD	40GB	100GB	150GB
Networking	100Mbps	1000Mbps	1000Mbps

Eucalyptus is space sensitive and as such a large disk space is highly recommended. The NC requires a 64-bits VT-x enabled CPU architecture since it would be running the VM, otherwise the hypervisor is unable to run thereby rendering the hardware useless. 64-bits CPU architecture would run 32-bits images as well. A recommended 100Mbps network is required for Eucalyptus to function efficiently as images would be copied over the network.

5.3 Network Configuration

The following section describes the network as implemented for the thesis while the physical topology implemented is as shown in Figure 6.




Figure 8. Test network physical topology.

Two servers and a client were interconnected over a private network (LAN) setup entirely for the purpose of this thesis using a Cisco switch as an access switch and a Cisco router connected to the Cisco Networking Academy's lab LAN (internet) via the Fast Ethernet interface 0 and to the Eucalyptus private network via the Fast Ethernet interface 1, which also doubles as the private LAN gateway. The two servers are marked Server 1, prepared to contain the CLC, CC, SC and WS3, while Server 2 was prepared for the NC. The client was to be used to test access from within the network. Further, clients connected to the Academy's Network should be able to access the cloud, launch VMs and communicate with them should a security group be created for that purpose. The test uses the default networking configuration mode, Managed No-VLAN mode, which support the dynamic IP assignments to VMs and allows control of ingress traffic by the creation of profiles known as security groups. With this, it is possible to control what group of users has access to the cloud thereby making it one of the major security features of Eucalyptus.

For best practices, traffic to and from the NC should be through the CLC which means that the NCs gateway address should be that of the CLC such that the CLC has two

Ethernet interfaces, one for the inside or private network and the other connected to the outside or public network, but due to limited hardware resources, hosts with multiple Ethernet interfaces could not be accrued for the purpose of this thesis.

IP Addressing

The entire cloud infrastructure is setup to be built on the class 'C' private network block of address, 192.168.1.0/24. In Table 1.0 below is the addressing for the cloud infrastructure.

Table 7. IP addressing of the cloud infrastructure

	Interface	IP Address	Gateway	Name Server
Server 1	Eth0	192.168.1.6	192.168.1.5	193.166.135.5
Server 2	Eth0	192.168.1.8	192.168.1.5	193.166.135.5
Client	Eth0	192.168.1.9	192.168.1.5	193.166.135.5
Router	Fa0/0	DHCP	-	-
Router	Fa0/0	192.168.1.5	-	-

6 Installation

Setting up the cloud environment required the steps as described below.

6.1 Building the Cloud Controller

There are two major ways of installing the UEC environments:

- Installing from the CD/DVD or
- Installing package by package

Another way would be to install from a server prepared for that purpose or cloning. Installing from the CD/DVD was the preferred method as it is much easier this way.

The CLC was built by first booting the system from the Ubuntu 10.04 CD and choosing the “Install Ubuntu Enterprise Cloud” from the list of options.

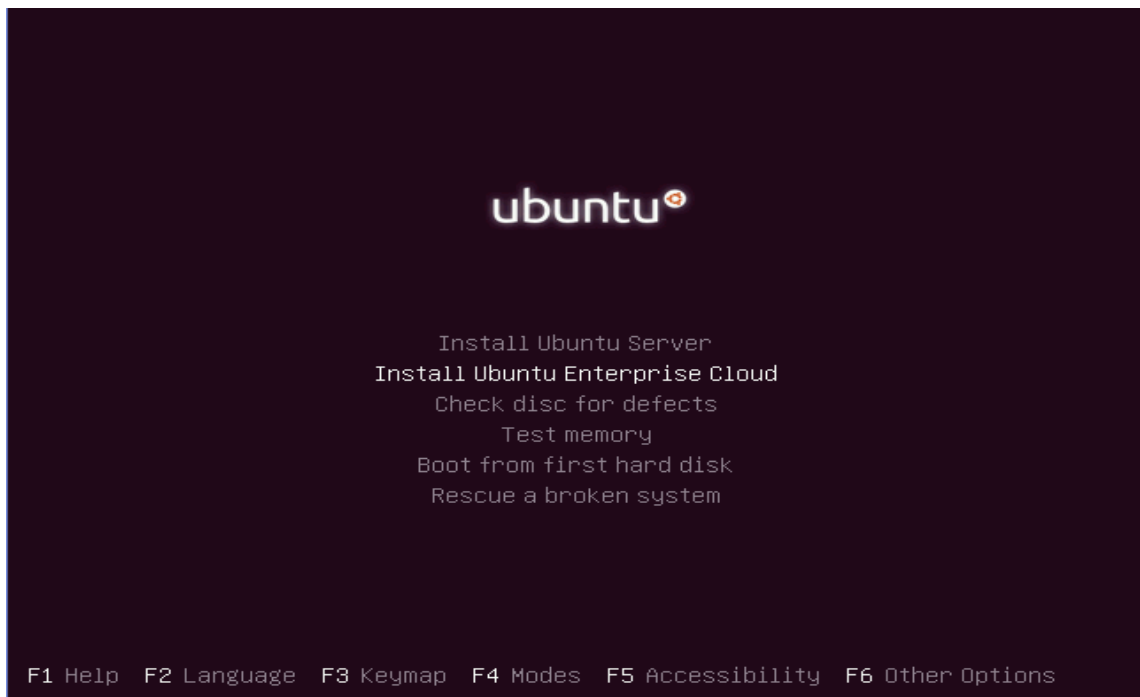


Figure 9.1 Booting the Ubuntu Server 10.04 from CD.

After choosing the appropriate region, language and keyboard layout for which screenshots are not provided, the system prompts for the configuration of the network interface after which it requests for a hostname. The name “cc” was chosen for the CLC.

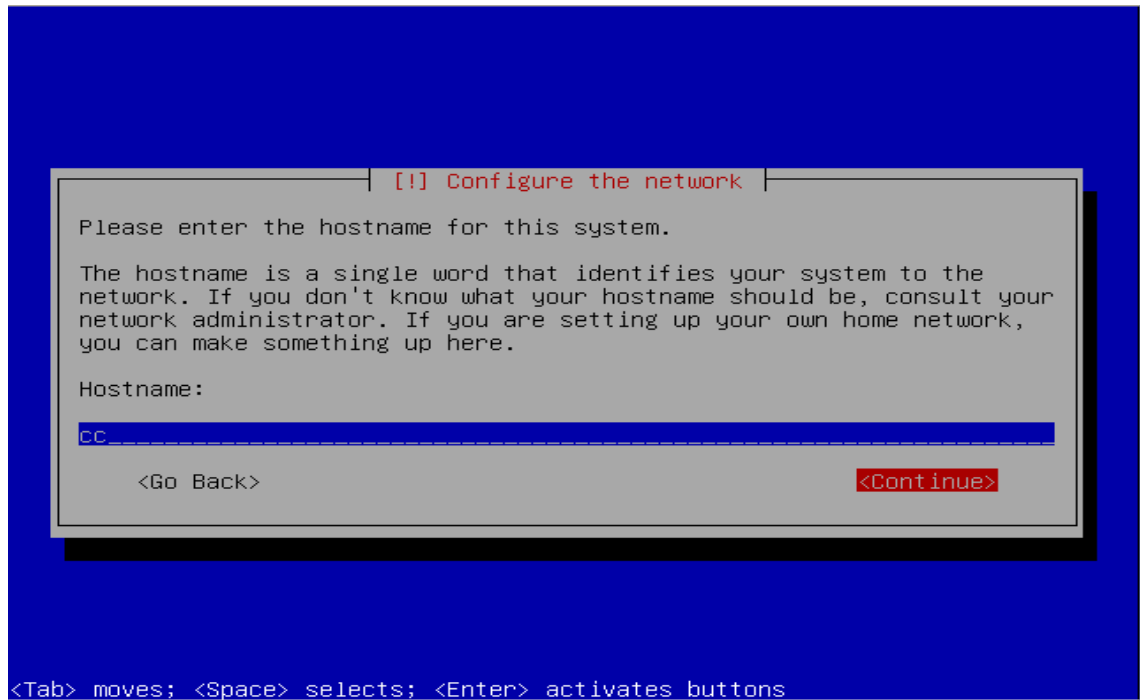


Figure 9.2 Choosing a name for Server 1

The following screenshot is so because there is no pre-installed CLC within the network so the installations continued by selecting continue.




Figure 9.3 No pre-installed cloud controller on the network

The following screenshots is where the role of Server 1 in the cloud was decided. For larger and or complex installations, each role shown in the list of options should be

installed on separate server but for this installation, all other roles except the NC was installed on Server 1.

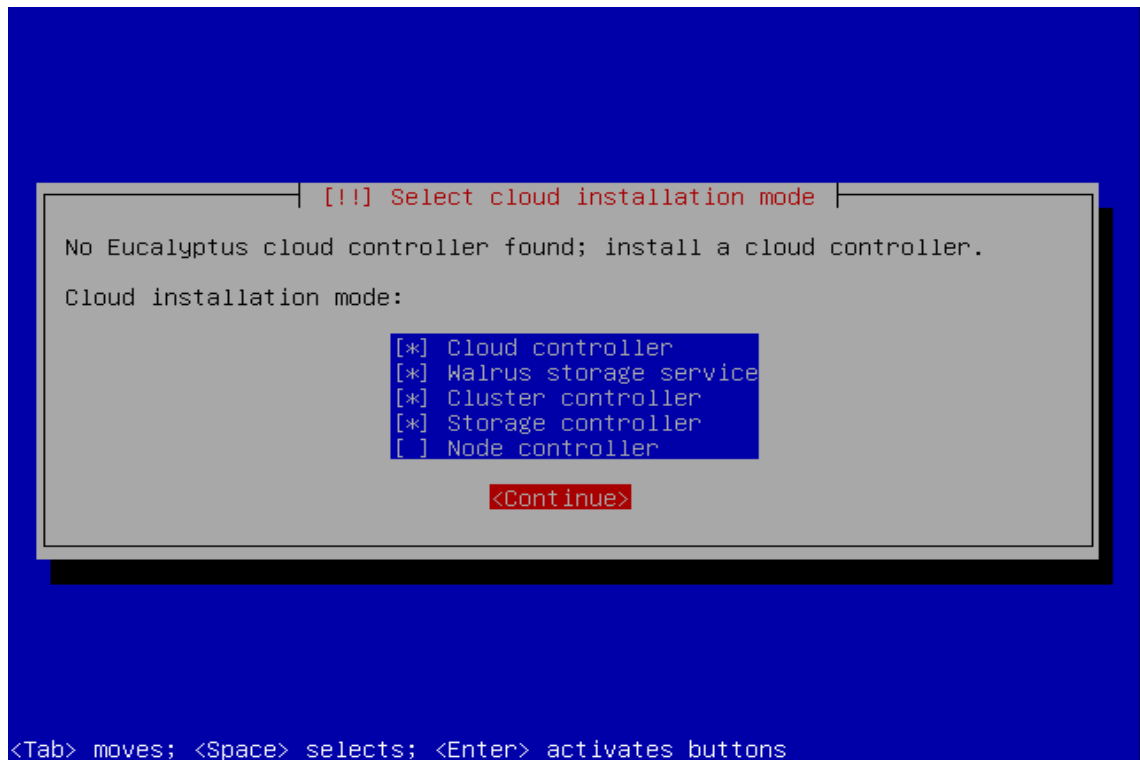


Figure 9.4 Roles of the front-end server decided

The next few shots show accepting the default HDD partition options.

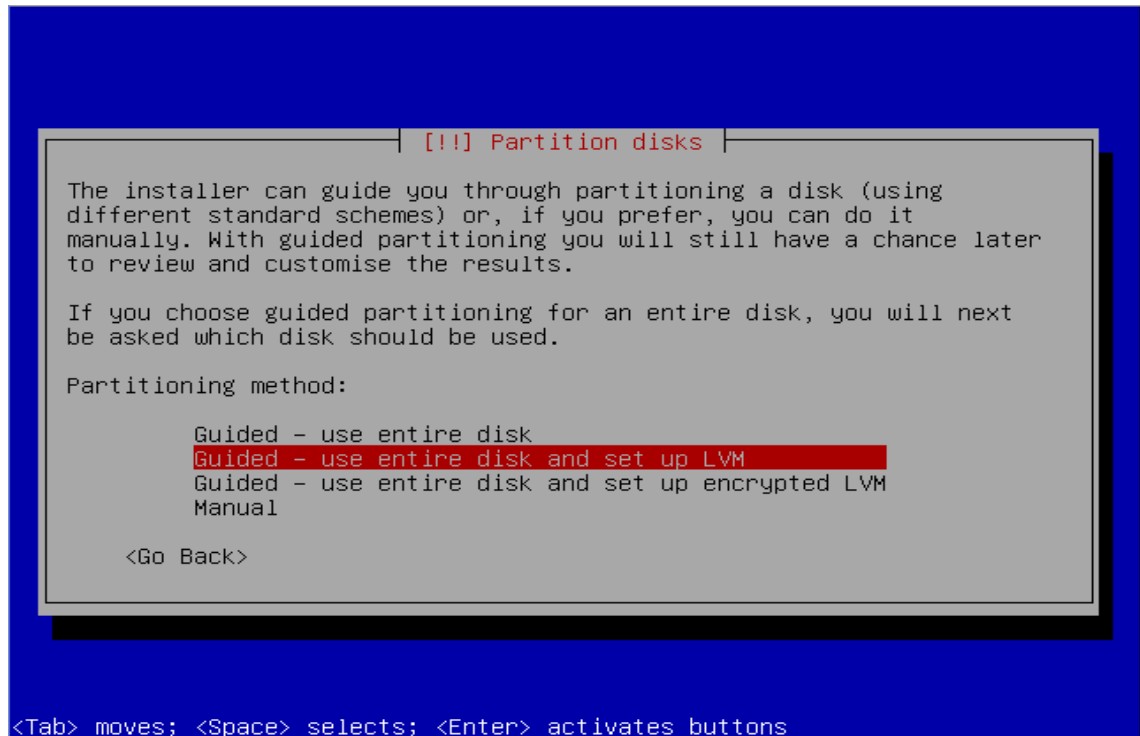


Figure 9.5 The entire HDD partitioned




Figure 9.6 The only partition available to the installer

In order to keep this simple, the entire volume group of the HDD was used for guided partitioning.

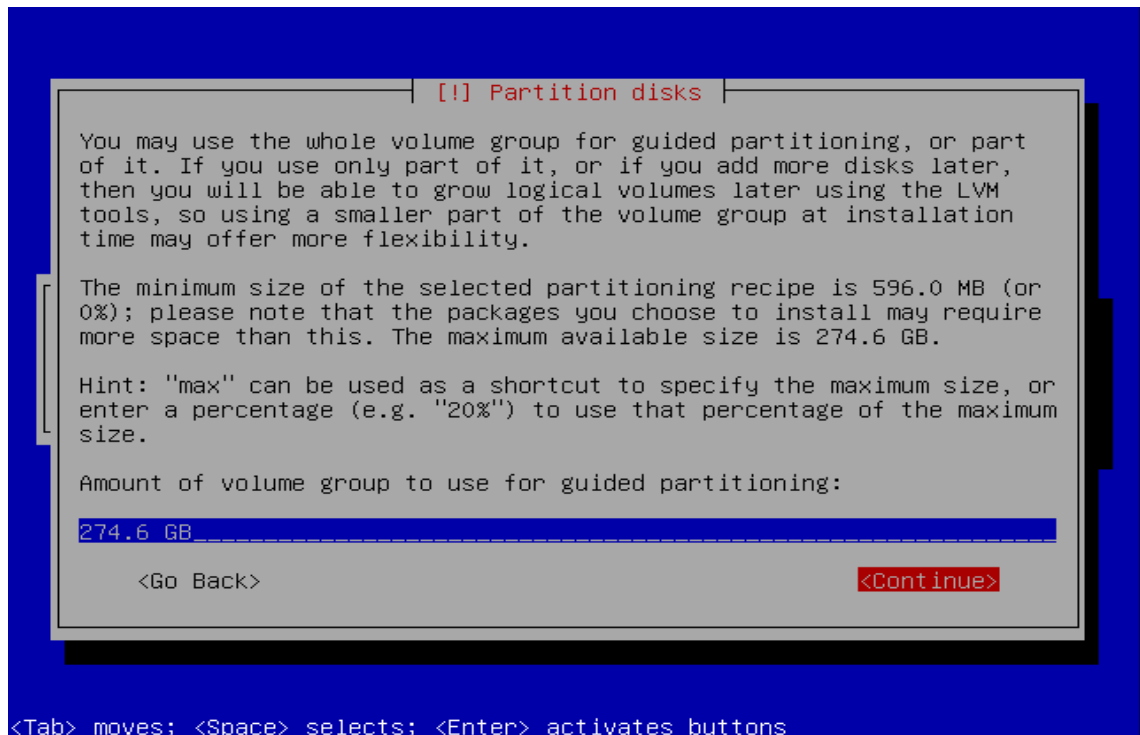


Figure 9.7 The entire disk space partitioned

Since there are no data currently available on the HDD, the default partition and configured Linux Volume Manager was accepted and was written to disk.

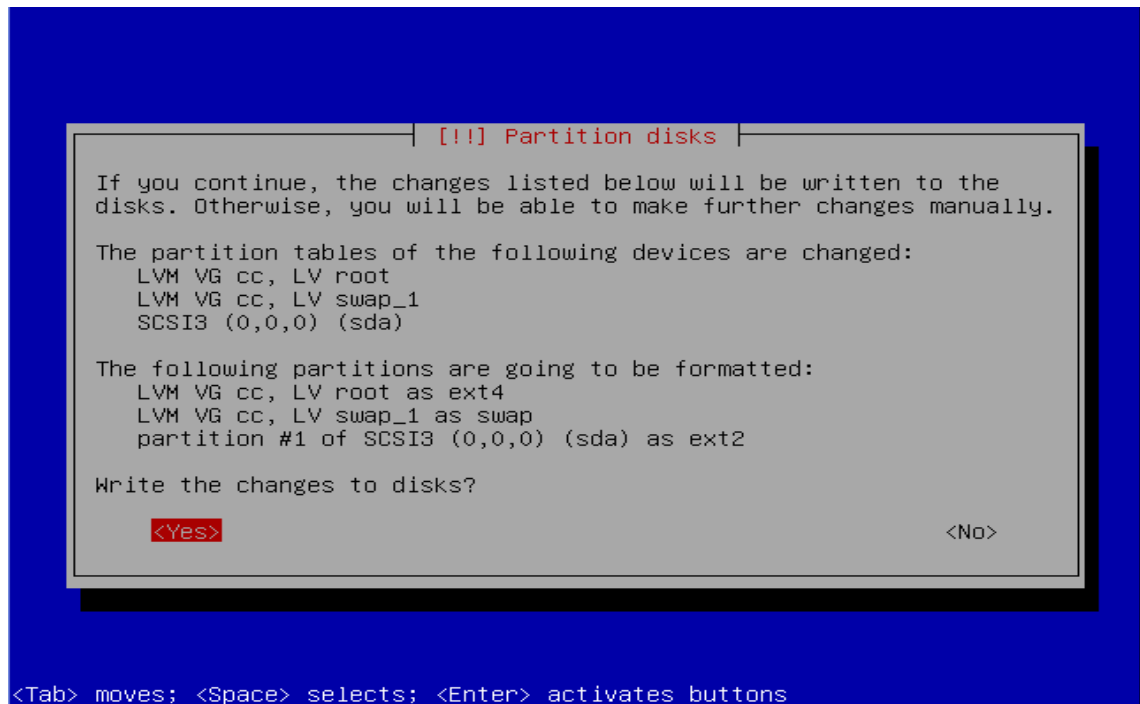



Figure 9.8 Confirming partition for format

At this point, the partitions were formatted while the base system was installed afterwards.



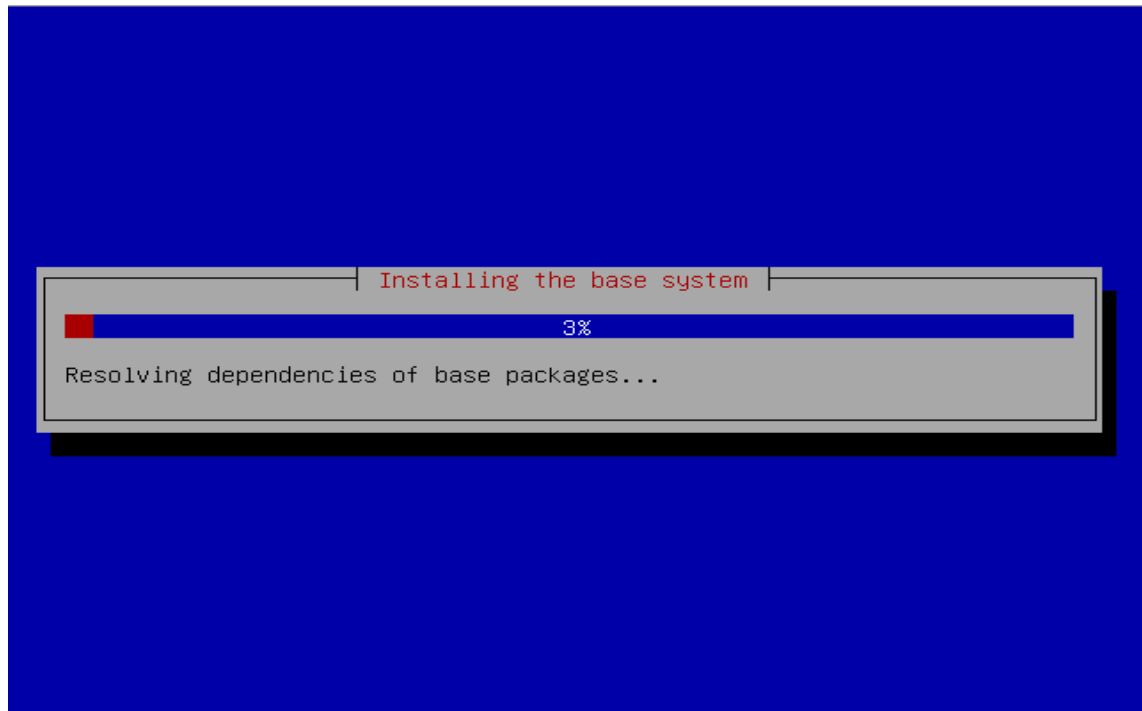


Figure 9.10 Installing the base system

Here the user credentials are prompted for creation. Noteworthy is that the account created here would have administrative or sudo (Linux) privileges. The user's real name is as follows.

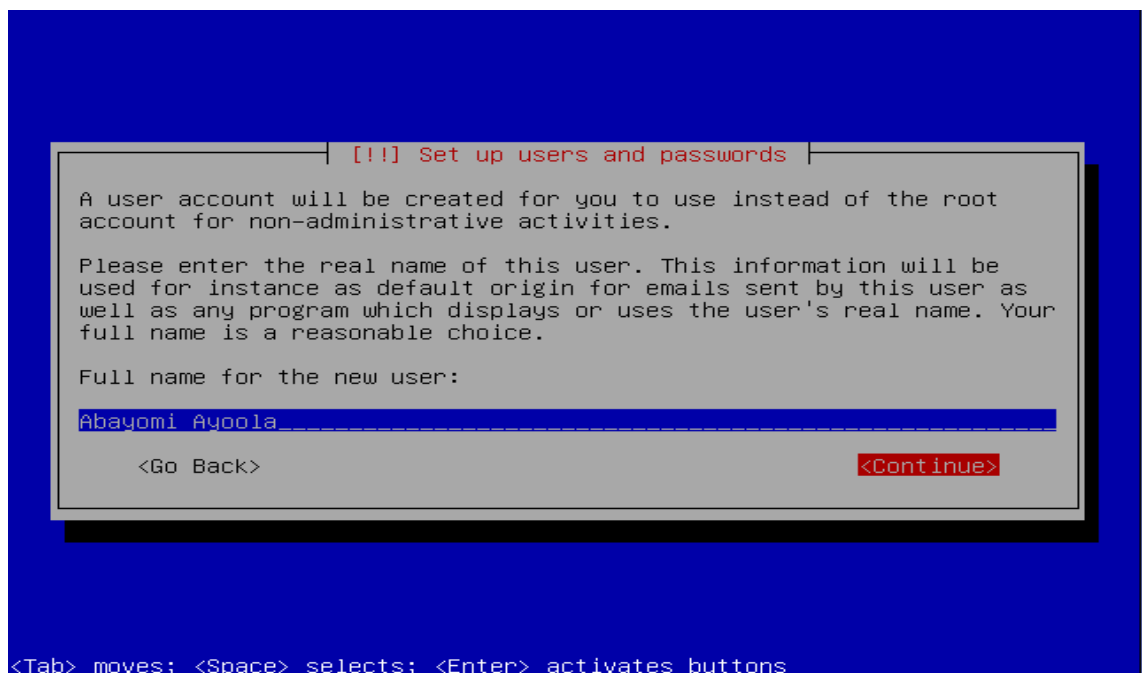


Figure 9.11 Setting up users full name

And the username as “ccadmin”

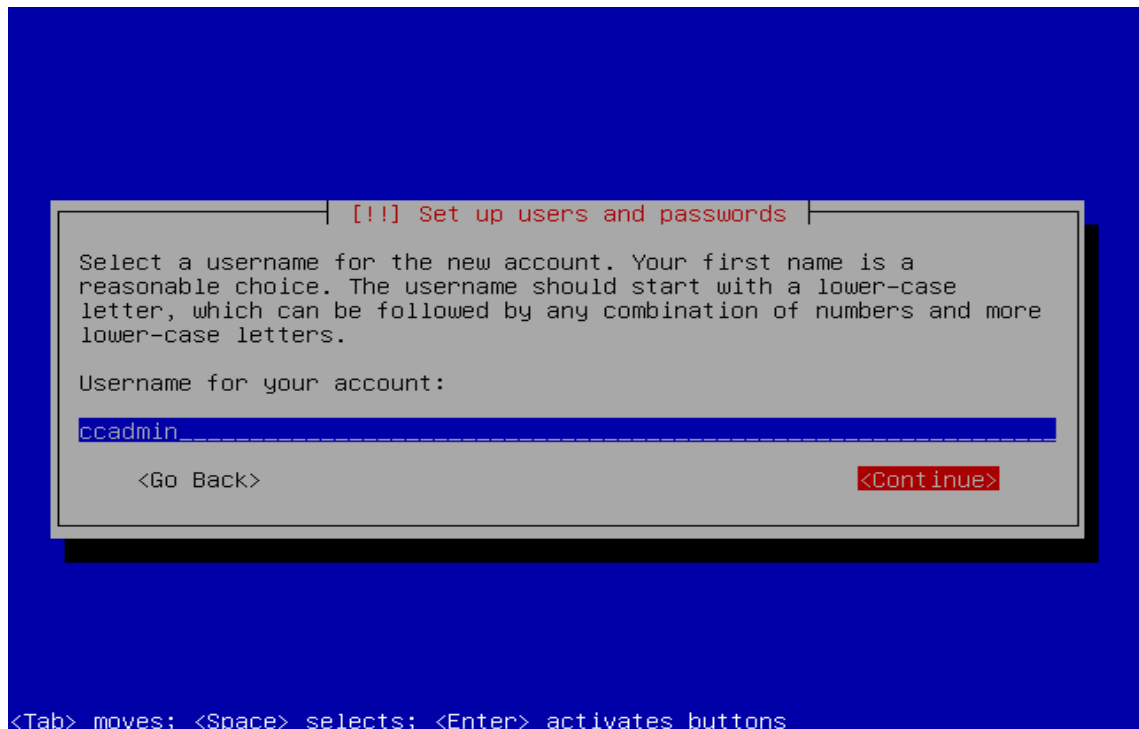


Figure 9.12 Setting up username

And password “MYworld247,”




Figure 9.13 Setting up users password

There would not be any need to encrypt the home directory and so the default option was chosen and also for the immediate following shot.

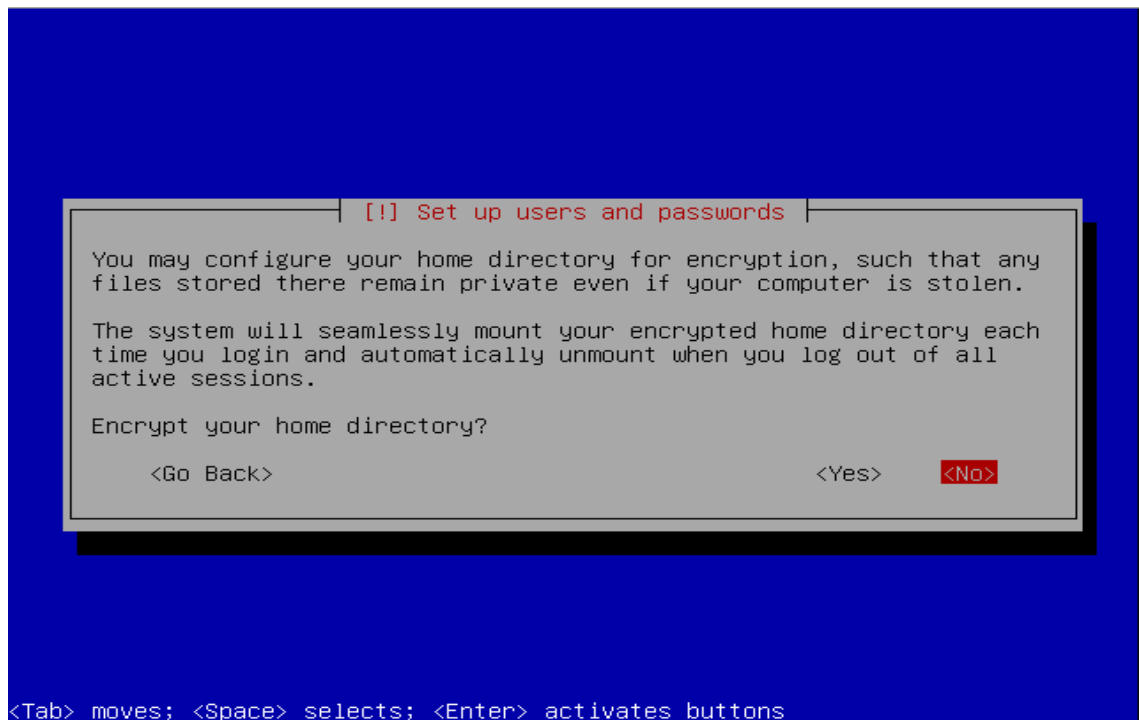


Figure 9.14 No need for encryption

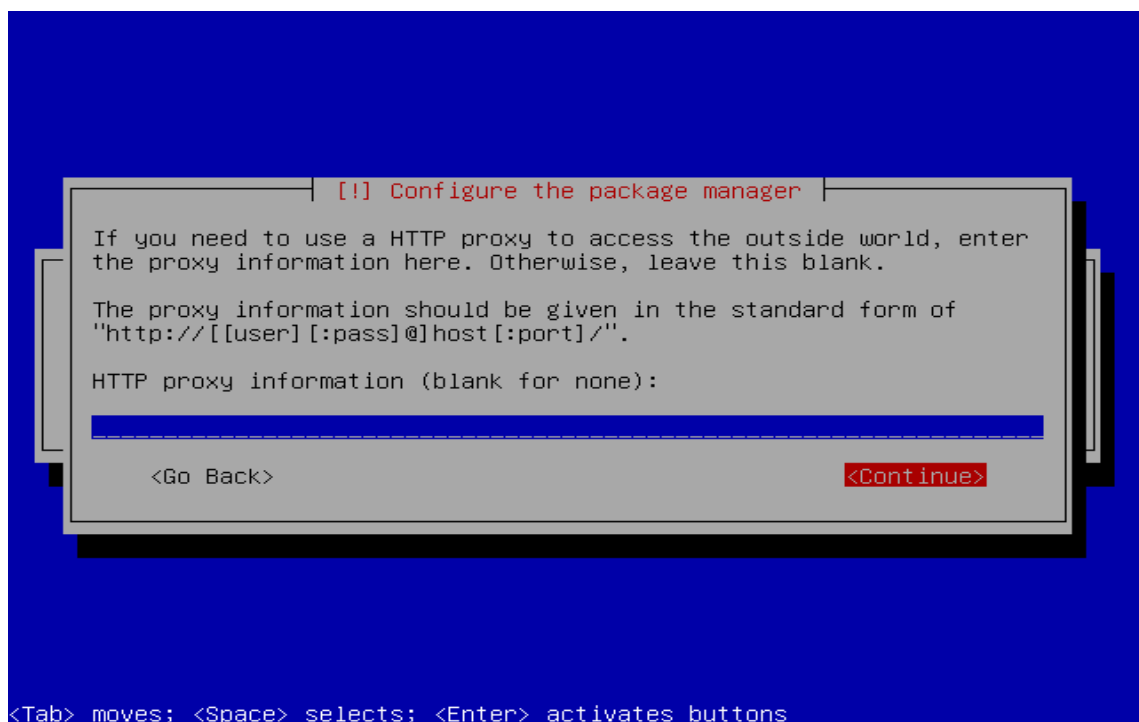


Figure 9.15 No need for proxy information

Although recommended, the “No automatic updates” option was chosen.

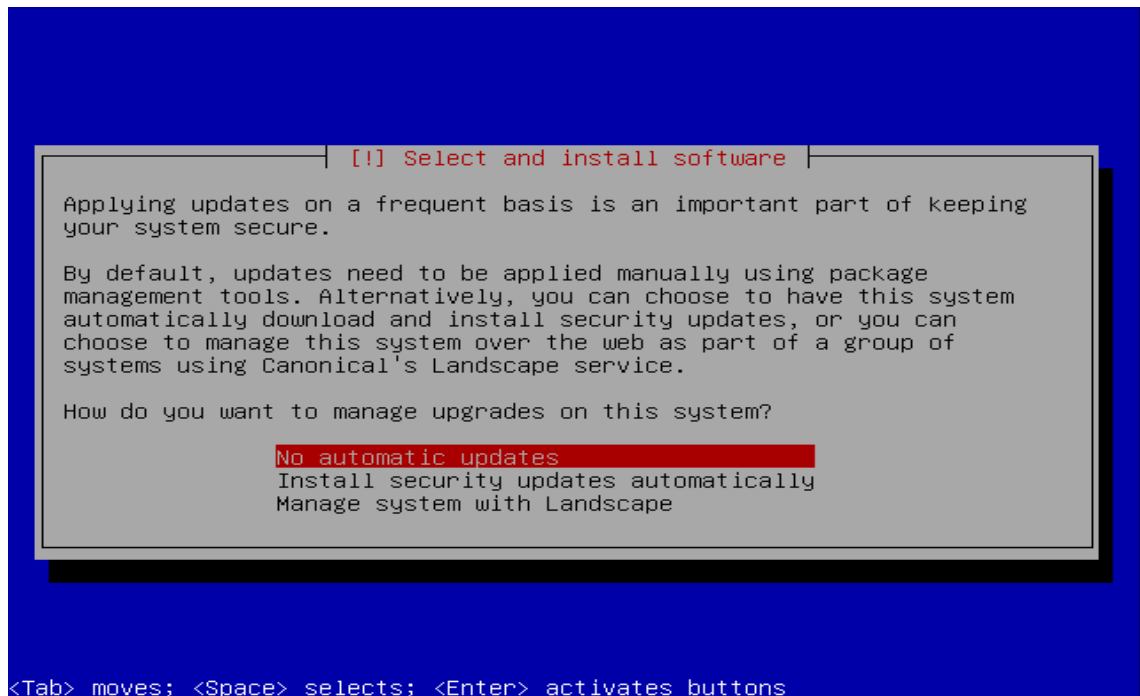


Figure 9.16 Automatic updates are not required

From this point on, the Eucalyptus Cluster Controller would be configured. The initial configuration is the cluster name, “cluster1” and can also be named anything.



Figure 9.17 Cluster name declared

A pool of IP addresses that would be dynamically assigned (Public) to VMs as they are created are provided at this point giving accessibility to cloud instances from outside the cloud network. In this case, 192.168.1.101 – 192.168.1.199

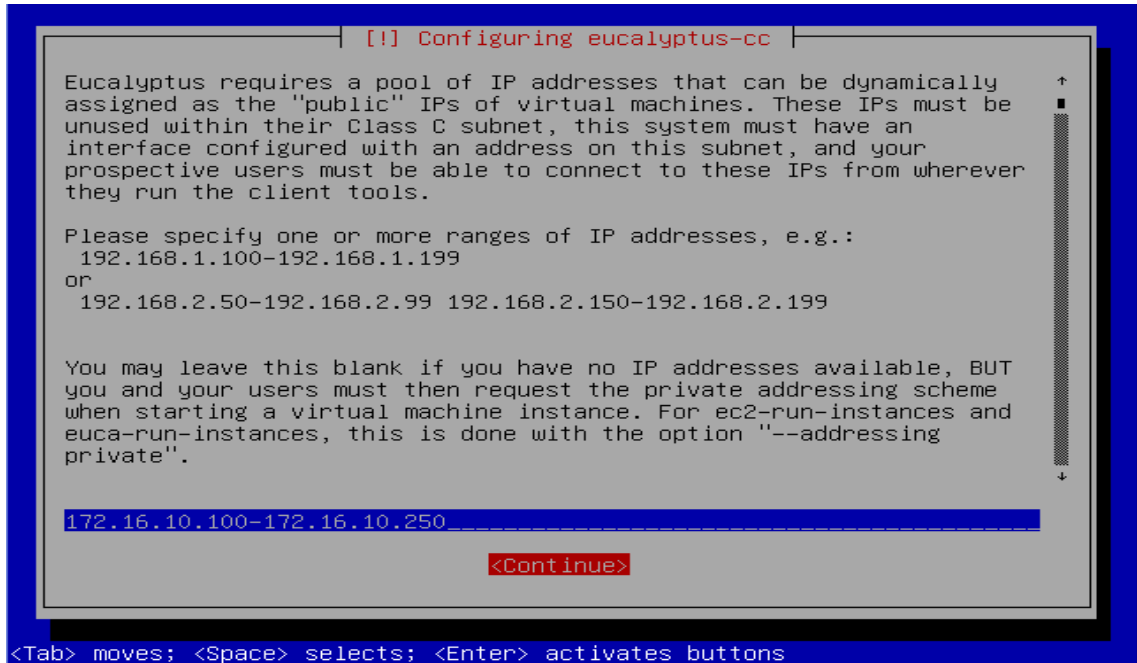


Figure 9.18 Input a pool of public IP addresses

A FQDN is required at this stage, in this case "cc.lns-servers.com"

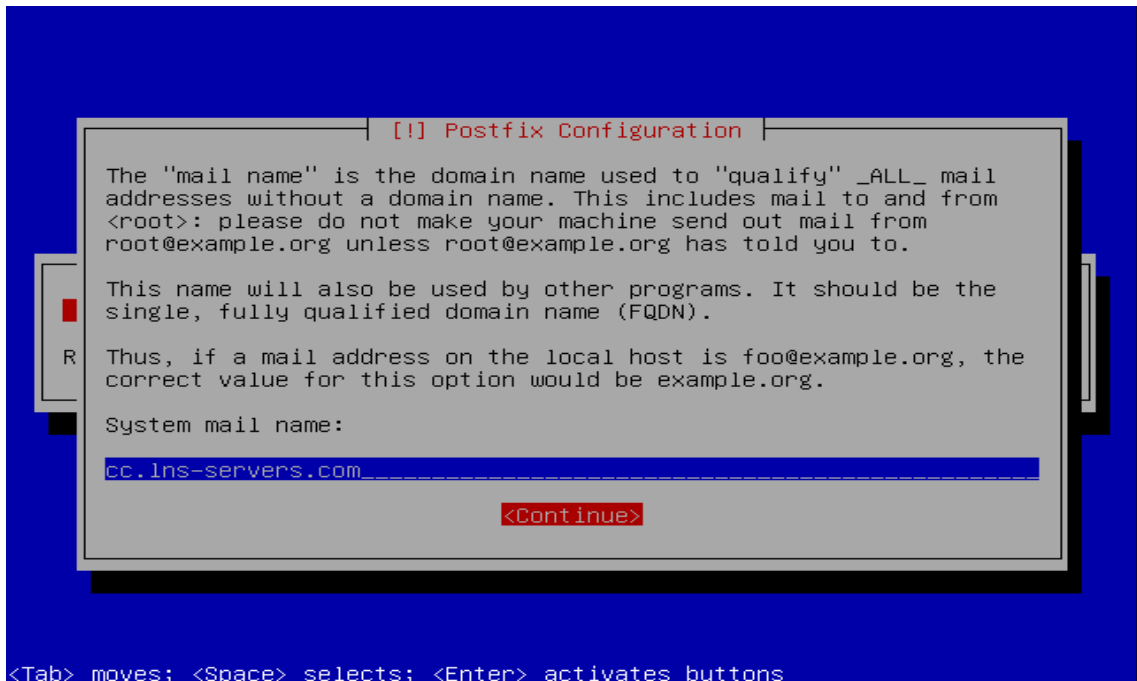


Figure 9.19 – A Fully Qualified Domain Name is required

And finally, the boot loader is installed

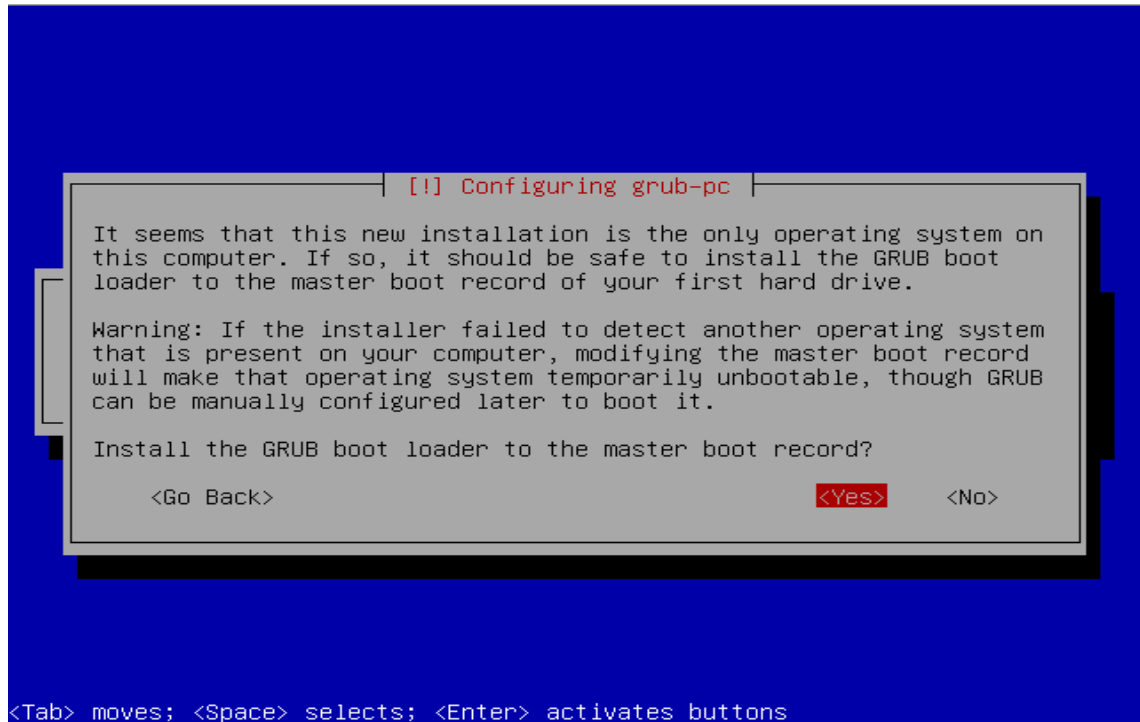


Figure 9.20 – Requesting installation of GRUB

This concludes the installation of the CLC.

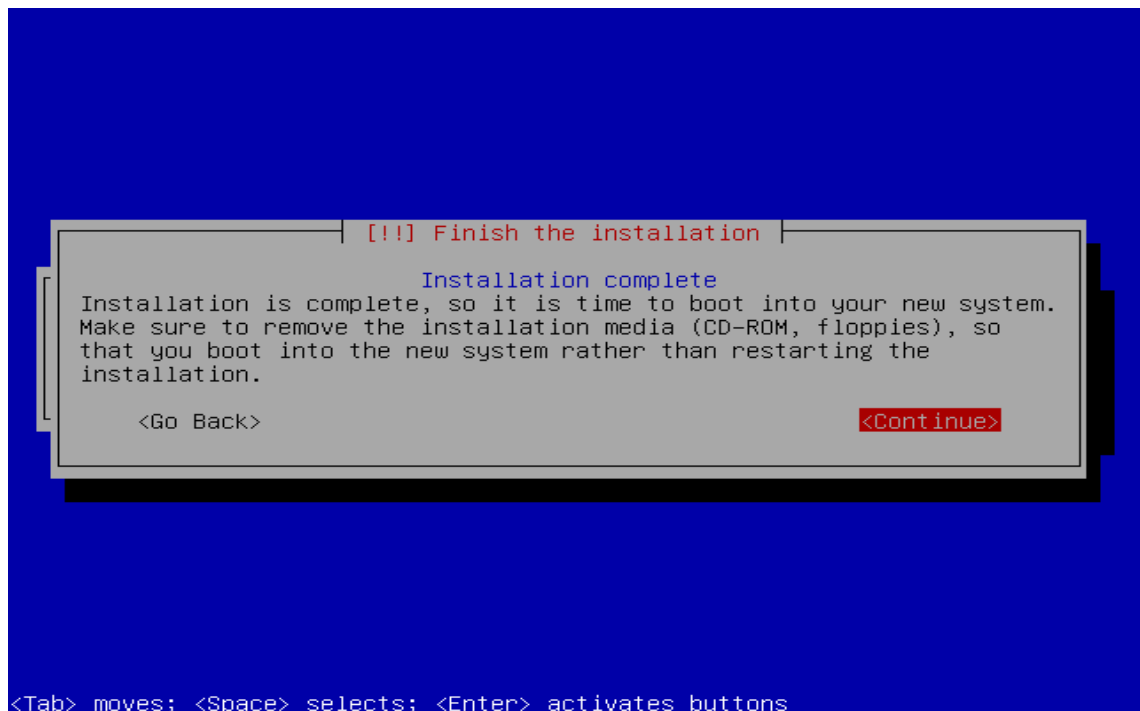


Figure 9.21 – Front-end installation complete

6.2 Building the Node Controller

The process of building the Node Controller is as the same as the process required in building the Front-end. The Node Controller was built on Server 2. Booting the Ubuntu Server 10.04 disk, the “Install Ubuntu Enterprise Cloud” option is chosen.



Figure 10. Installing the node controller

After making all necessary selections and options just as in the case of the Front-end, a hostname is required for Server 2.

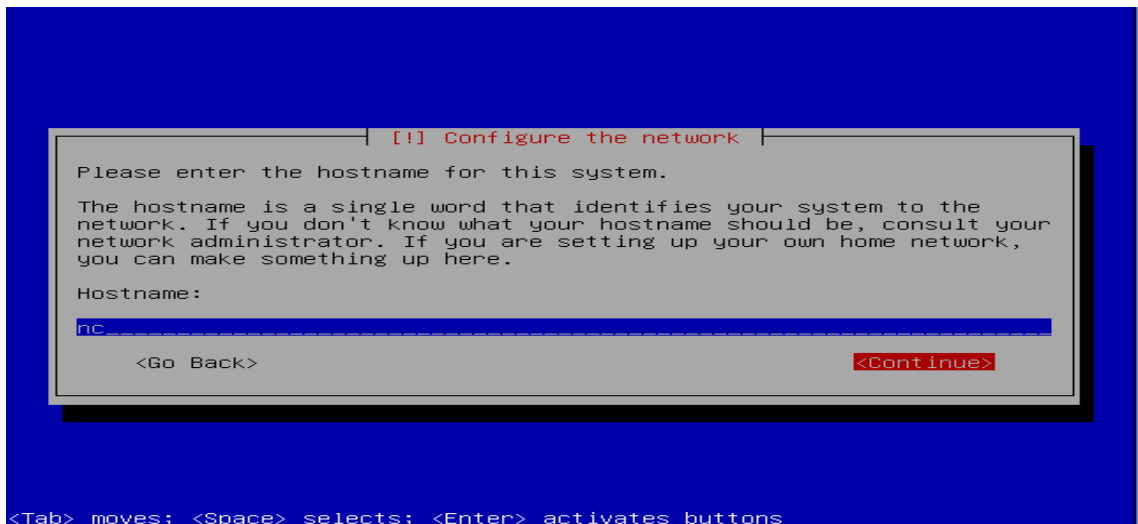


Figure 10.1 – Giving the node controller a hostname

The Node Controller should automatically detect the Front-end already running within the same network and as such would be the default option.

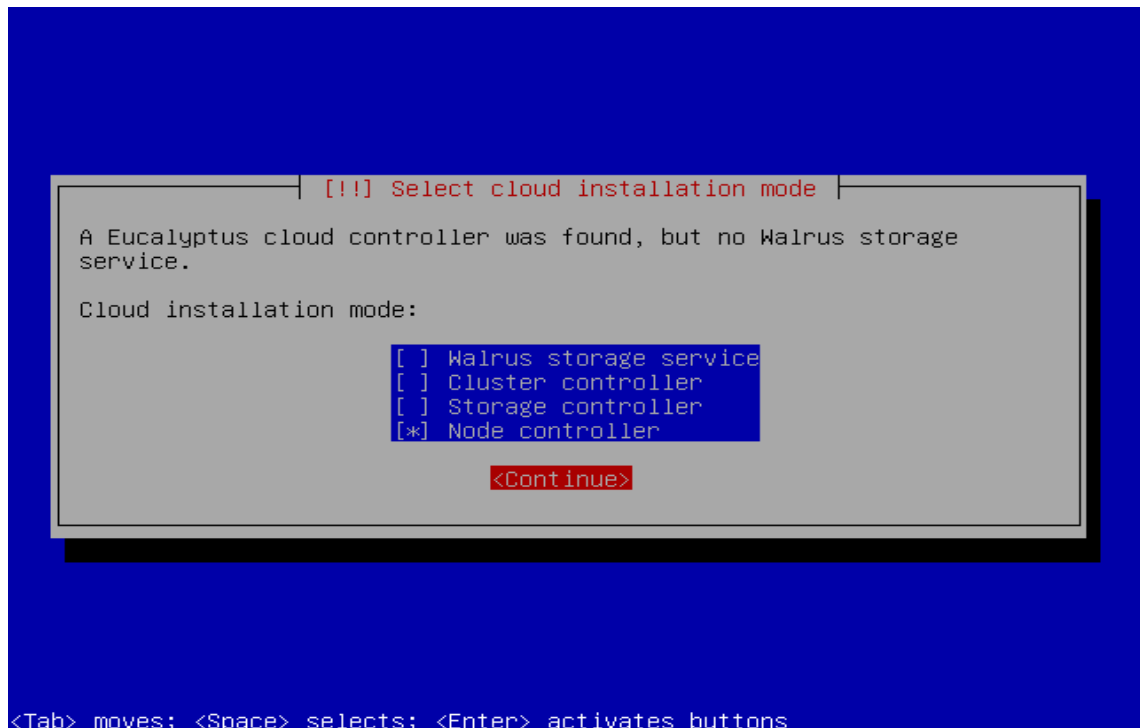


Figure 10.2 The role of Server 2

The installation screens after this will present the same partitioning options as when installing the cloud controller and the same settings were used. The installation is ended and the node controller is rebooted.

For a successful functioning of the environment, a number of post installation configurations were carried out in the following texts.

6.3 Post-installation Configurations

Should all service be properly configured as in the installation steps above, all components registration should be automatic. In order to verify registration, the following command should be entered on the CLC command interface.

```
$ cat /var/log/eucalyptus/registration.log
```

The output should look like the following:


```
$ 2012-04-08 10:46:36-05:00 | 24243 -> Calling node cluster1 node 10.1.1.75
$ 2012-04-08 10:46:36-05:00 | 24243 -> euca_conf --register-nodes returned 0
$ 2012-04-08 10:48:47-05:00 | 25858 -> Calling walrus Walrus 10.1.1.71
$ 2012-04-08 10:48:51-05:00 | 25858 -> euca_conf --register-walrus returned 0
$ 2012-04-08 10:49:04-05:00 | 26237 -> Calling cluster cluster1 10.1.1.71
$ 2012-04-08 10:49:08-05:00 | 26237 -> euca_conf --register-cluster returned 0
$ 2012-04-08 10:49:17-05:00 | 26644 -> Calling storage cluster1 storage 10.1.1.71
$ 2012-04-08 10:49:18-05:00 | 26644 -> euca_conf --register-sc returned 0
```

IP addressing on the physical hardware Front-end and the NC were configured as static. The following IP configuration were added to /etc/network/interfaces and the DNS server details to /etc/resolv.conf.

Front-end:

```
auto eth0
iface eth0 inet static
address 192.168.1.6
netmask 255.255.255.0
network 192.168.1.0
broadcast 192.168.1.255
gateway 192.168.1.5
```

DNS Server:

```
nameserver 193.166.135.5
```

NC:

```
auto eth0
iface eth0 inet static
address 192.168.1.8
netmask 255.255.255.0
network 192.168.1.0
broadcast 192.168.1.255
gateway 192.168.1.5
```

DNS Server:

```
nameserver 193.166.135.5
```

For the changes to take effect, the networking was restarted by running the command

```
ccadmin@clc:~$ sudo /etc/init.d/networking restart
```

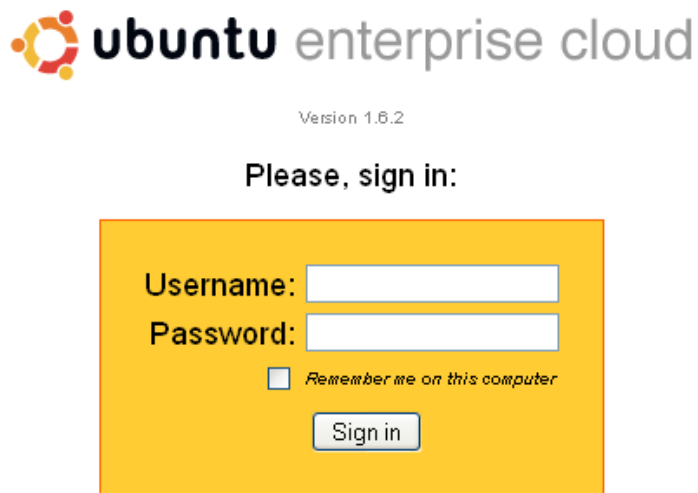
6.3.1 Installing Administrative Credentials

In order to be able to use the command-line tools to interact with the cloud, credentials which include a x.509 certificate and environment variable is to be installed. This can be done either remotely or locally on the CLC. It was done remotely on a client within the same network as the CLC in this case and the following steps were arrived at.

From a browser, the address of the CLC was entered as follows

`https://IP-address-of-CLC:8443`. The IP address of the CLC in question is 192.16.1.6

It should bring up the login page as seen below.



ubuntu enterprise cloud

Version 1.6.2

Please, sign in:

Username:

Password:

Remember me on this computer

Sign in

Figure 11. CLC web interface

The default user credentials are the following: username = admin and password = admin.



Version 1.6.2

Please, sign in:

A login form with a yellow background. It contains the following elements:

- A label "Username:" followed by a text input field containing the text "admin".
- A label "Password:" followed by a password input field with six black dots.
- A checkbox with the label "Remember me on this computer".
- A "Sign in" button.

Figure 12. Entering the default user credential

After login in the first time, new admin credentials are required.



First-time Configuration

Please, supply the following parameters to activate your Eucalyptus installation:

Administrator's new password:

The password, again:

Administrator's email address:

User signup requests will be sent to this address and messages from Eucalyptus will have this address in the *From:* field. (If you want to change this behavior, edit the appropriate values in the *eucalyptus-web-properties* file of your Eucalyptus installation. We still need an email address now, though.)

Figure 13. The new admin credentials are entered after login in the first time.

The new admin credentials password was "MYworld247". Most of the administration would be done from the command line using the euca2ools which has been installed by default on the CLC and as such the x.509 certificate has to be installed. The certificate credentials are accessed by clicking on the credential tab on the web interface and clicking "Download Credentials" as in the figure below



The screenshot shows the Ubuntu Enterprise Cloud management interface. At the top, there is a navigation bar with the Ubuntu logo and the text "ubuntu enterprise cloud". On the right side of the navigation bar, it says "Logged in as admin | Logout". Below the navigation bar, there is a secondary navigation bar with tabs for "Credentials", "Images", "Store", "Users", "Configuration", "Services", and "Extras". The "Credentials" tab is currently selected. In the main content area, there is a section titled "User account Information" with the following details: Login: admin, Name: (blank), and Email: a.ayoola@gmail.com. Below this information, there is a paragraph of text: "Feel free to change the account information (except the login) and the password whenever you want. The cryptographic credentials for the Web services associated with this account, shown below, will not be affected by these changes." Underneath the text are two buttons: "Edit Account Information" and "Change Password". Below these buttons is another section titled "Credentials ZIP-file" with a paragraph of text: "Click the button to download a ZIP file with your Eucalyptus credentials. Use the public/private key pair included therein with tools that require X.509 certificates, such as Amazon's EC2 command-line tools." At the bottom of this section is a button labeled "Download Credentials".

Figure 14. Downloading credential certificate

After the download was completed on the client, the downloaded file was copied via ftp to the home directory, /home/ccadmin, of the "ccadmin" on the CLC. At this point, the downloaded file, euca2-admin-x509.zip, needed to be extracted and placed in the proper folder to provision its functions. From this point, the administration is moved to the CLC where a hidden folder was created in which the unextracted certificate credential is to be moved to. This is achieved with the following commands:

```
ccadmin@clc:~$ mkdir ~/.euca          create a hidden directory named ".euca"
ccadmin@clc:~$ mv euca2-admin-x509.zip ~/.euca  move certificate into ".euca"
ccadmin@clc:~/.euca$ unzip /euca2-admin-x509.zip  uncompress certificate file with
unzip utility tool
```

In the next step, the compressed file is removed and permissions is applied to the ".euca" directory and its contents as follows:

```
ccadmin@clc:~/.euca$ rm euca2-admin-x509.zip  remove compressed file
ccadmin@clc:~/.euca$ chmod 0700 ~/.euca      apply permission to the .euca folder
ccadmin@clc:~/.euca$ chmod 0600 ~/.euca/*  apply permission to the .euca contents
```

For the proper initialization of essential environment when login in, the following line, “. ~/.euca/eucarc” was added to the ~/.bashrc file on the CLC.

```
ccadmin@clc:~/euca$ echo “. ~/.euca/eucarc” >> ~/.bashrc
```

and in order to ensure the proper functioning of the configurations, the .bashrc file was sourced with the command below and then the user “ccadmin” was logged out and logged back in.

```
ccadmin@clc:~/euca$ source ~/.bashrc
```

6.4 Installing Images

By default, the image tab has no listed image as none come installed but images can either be bundled from the Canonical cloud image store; uploaded and registered with the cloud or installed directly via store tab from the same image store over the internet. The latter method was chosen because of its simplicity. A major issue encountered while trying to do this, was that the remote server hosting the images failed to verify the CLC thereby displaying the error message in Figure 15. Researching deeply into this reveals that the certificate issued by Canonical and signed by GoDaddy is not trusted by the CLC. Every trusted CA is listed in /etc/ssl/certs and because the CLC could not find the certificate in the list, it is unable to verify it. A way around this was to include the same certificate in the trusted list. To do this, the certificate was downloaded from <https://imagestore.canonical.com> and added to the list.

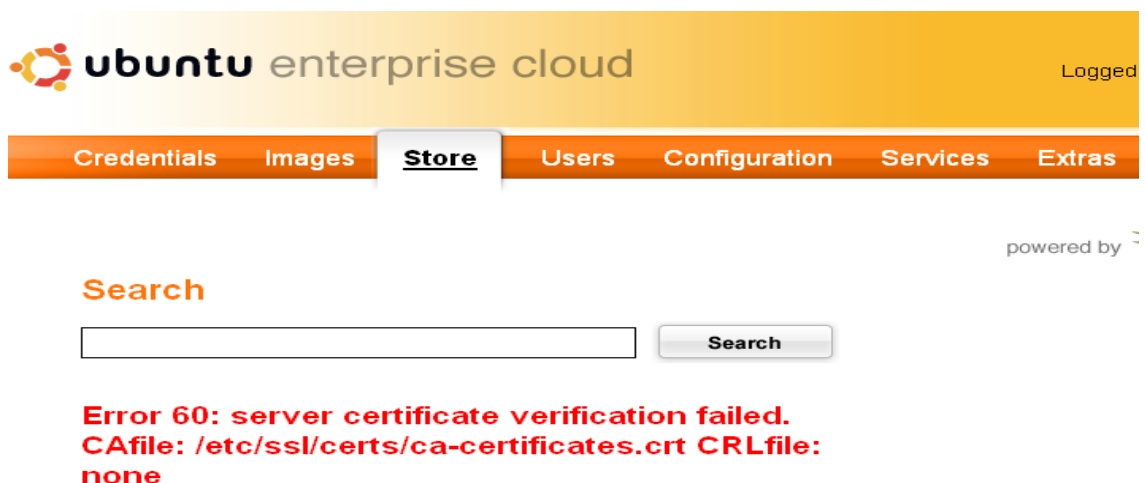


Figure 15. Error message displayed on store tab

From a Firefox browser, it was navigated to <https://imagestore.canonical.com> and the exception was added as requested. Although the page displays a “Cannot be found”, the icon before web address was clicked, this displays the certificate file which was saved and exported to the CLC precisely to the following folder, `/usr/share/ca-certificates/godaddycertificate.crt` and following command was run to update the trusted list.

```
ccadmin@clc:~$ sudo dpkg-reconfigure ca-certificates
```

The new certificate was chosen at this point and the “enter” key was pressed for it to be installed. The issue persisted after doing this. A solution that worked is not recommended to be tried in a production environment as it would open the cloud infrastructure to security attack. The option was to connect to the image store insecurely that is to disable certificate verification as there is a problem. This is done by editing the following file,

`/usr/lib/python2.6/dist-packages/imagestore/lib/fetch.py` and including the lines listed:

```
curl.setopt(pycurl.SSL_VERIFYPEER, 0)
```

```
curl.setopt(pycurl.SSL_VERIFYHOST, 0)
```

The image store proxy was restarted after this by running the following command:

```
ccadmin@clc:~$ sudo /etc/init.d/image-store-proxy restart
```

Afterwards, the GoDaddy certificates was updated issuing the following set of commands:

```
ccadmin@clc:~$ sudo wget -P /usr/local/share/ca-certificates/ --no-check-certificate
```

```
https://certs.godaddy.com/repository/gd-class2-root.crt
```

```
https://certs.godaddy.com/repository/gd_intermediate.crt
```

```
https://certs.godaddy.com/repository/gd_cross_intermediate.crt
```

```
ccadmin@clc:~$ sudo update-ca-certificates
```

This appears to work as images were immediately displayed under the “Store” tab as displayed in Figure 16 and images were confirmed to have been registered with the cloud under the “Images” tab as displayed in Figure 17.

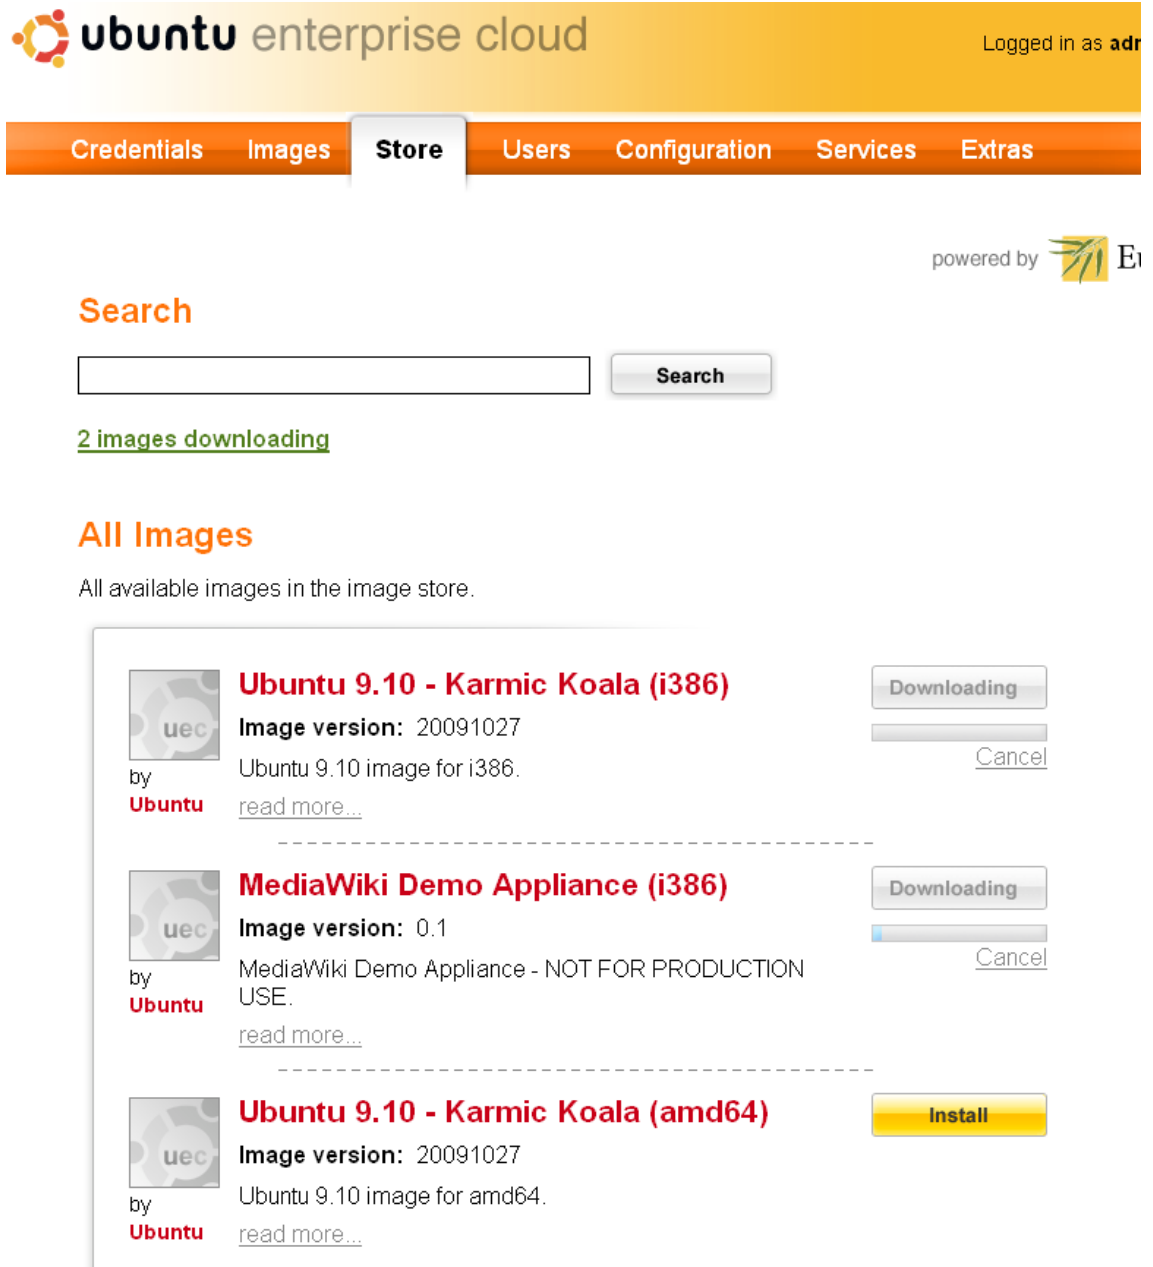


Figure 16. The images finally displays and being installed

The cloud infrastructure is ready at this point and VMs can be instantiated.

ubuntu enterprise cloud Logged in as admin | Lo

Credentials **Images** Store Users Configuration Services Extras

Id	Name	Kernel	Ramdisk	State	Actions
eki-F3C110E1	image-store-1340710590/kernel.manifest.xml			available	Disable
emi-DD521051	image-store-1340710590/image.manifest.xml	eki-F3C110E1	eri-07E3113B	available	Disable
eri-07E3113B	image-store-1340710590/ramdisk.manifest.xml			available	Disable
eri-07C61145	image-store-1340710462/ramdisk.manifest.xml			available	Disable
emi-DD4D1061	image-store-1340710462/image.manifest.xml	eki-F35310D4	eri-07C61145	available	Disable
eki-F35310D4	image-store-1340710462/kernel.manifest.xml			available	Disable

powered by Eucaly

Figure 17. Registered Images under “Images” tab

Worthy of note here is the image ID, emi-xxxxx, under the “Id” column which would be used as identifier in running instances in the following chapter. This is the Amazon Machine Image (AMI) equivalent from AWS which has disk image and pointer to a kernel bundled in it.

7 Instance Management

There are three ways of instantiating VMs in Eucalyptus. They are as listed below:

1. Command Line
2. UEC Compatible Management Tool e.g Landscape
3. Firefox Extension i.e Hybridfox

The command line is the preferred method as it gives total control of the server. It is necessary to be aware of sufficiency of resources available in the NC before running any instance, otherwise there might be issues. In order to access the availability of resources, the following command was ran:

```
ccadmin@clc:~$ euca-describe-availability-zones verbose
```

and the output is as follows

```
AVAILABILITYZONE    cluster1    192.168.1.6
AVAILABILITYZONE    |- vm types  free / max  cpu  ram  disk
AVAILABILITYZONE    |- m1.small  0004 / 0004  1   192  2
AVAILABILITYZONE    |- c1.medium 0004 / 0004  1   256  5
AVAILABILITYZONE    |- m1.large  0002 / 0002  2   512  10
AVAILABILITYZONE    |- m1.xlarge 0002 / 0002  2  1024  20
AVAILABILITYZONE    |- c1.xlarge 0001 / 0001  4  2048  20
```

These zones are modifiable under the “Administration” tab on the web interface.

7.1 Security Group

These is a set of packet filtering rules which controls connection requests coming from outside the cloud and targeted towards instances within the cloud. Eucalyptus default security group does not contain any rule and so allows every external connection to instances thereby opening up the cloud to security risks. The default security group can either be modified or a new group can be created but for this research, a new group was created and named, “Linux_Access”

```
ccadmin@clc:~$ euca-add-group Linux_Access -d access to linux instances
```

```
ccadmin@clc:~$ euca-authorize Linux_Access -P tcp -p 22 -s 0.0.0.0/0
```

```
ccadmin@clc:~$ euca-authorize Linux_Access -P tcp -p 80 -s 0.0.0.0/0
```

To confirm that the security group was built,

```
ccadmin@clc:~$ euca-describe-groups
```

This should display the newly built group.

7.2 Keypair

Keypairs are required to be injected into instances so that the instantiated VM is accessible via SSH.

```
ccadmin@clc:~$ euca-add-keypair access_key > ~/.euca/access_key.priv
ccadmin@clc:~$ chmod 0600 ~/.euca/access_key.priv
```

7.3 Running Instances

A single line of command is required to run an instance. At this point, an instance of Ubuntu 9.10 assigned to the Linux_Access security group is to be instantiated so that incoming connections is allowed on ports 22 (ssh) and 80 (http).

Running the command below was supposed to instantiate an instance of Ubuntu 9.10 assigned to the security group, Linux_Access and availability zone, m1.large which by default create a single instance but can be increased by including -n (number) followed by a number in the parameter.

```
$ euca-run-instances -g Linux_Access -k access_key -t emi-DD521051 m1.large
```

The instance would not run as the monitoring command suggested. By the way, the euca-describe-instances command can be run independently but the “watch” utility is used with the -n parameter to specify the number of seconds to refresh the euca-describe-instances command; that way it is easier to follow up with the instance from booting to running

```
cladmin@clc:~$ watch -n1 euca-describe-instances
```

And the output was as follows:

```
RESERVATION r-5CB0475C 965590394582
INSTANCE i-523E3E64 emi-DD521051 192.168.1.120 192.168.1.120 pending
access_key 0 m1.large 2012-08-08T23:15:46Z Lab-01 eki-F3C110E1 eri-07E3113B
```

And the output after about 5 seconds later:

```
RESERVATION r-5CB0475C 965590394582
INSTANCE i-523E3E64 emi-DD521051 192.168.1.120 192.168.1.120 terminated
access_key 0 m1.large 2012-08-08T23:25:50Z Lab-01 eki-F3C110E1 eri-07E3113B
```

The following line was detected in the NC log, `/var/log/eucalyptus/nc.log`, which after plenty of research suggests that the NC architecture does not support Eucalyptus default hypervisor, KVM as the CPU architecture is not VT-x enabled or optimized for virtualization and is unable to run VMs.

```
cladmin@clc:~$ cat /var/log/eucalyptus/nc.log
```

And the output:

```
[Mon Aug 28 11:42:32 2012][003524][EUCAERROR ] libvirt: internal error no supported architecture for os type 'hvm' (code=1)
```

A way to mitigate this is by deploying Eucalyptus using the Xen hypervisor rather than the default KVM but that is beyond the scope of this thesis.

8 Summary

The thesis work was splitted into two major parts with the first part explaining the overall technology of discussion while the second part was focused on the major topic of research, Ubuntu Enterprise Cloud (UEC) 10.04. For the readers who have little or no knowledge of Cloud Computing, Chapter 2 should be a good point of entry to establish the appropriate understanding to read through the thesis.

The overall building of the infrastructure was presented in Chapters 5 and 6. They explains in details, the requirements, configurations and about how Ubuntu Enterprise Cloud is built and configured with the production environment in mind. Chapter 5 discusses Eucalyptus in detail through which the readers can evaluate and probably adapt it to their own use. Chapter 6 provides detailed information on the implementation, installation and post-installation configuration phases. Although a major issue was encountered in this chapter, the resulting solution is, however, not recommended in a production environment as this would completely open the infrastructure to a possible external attack.

Chapter 7 discusses instance management as it relates to Ubuntu Enterprise Cloud. This chapter opens the reader not only to the steps required in instatiating an instance but also to the technical difficulties and challenges experienced while trying to run the cloud. The final result of implementation and configuration phase is clearly discussed.

REFERENCES

Brunette, G. & Mogull, R. (2009). *Security Guidance for Critical Areas of Focus in Cloud Computing Version 2.1*. Cloud Security Alliance Publication

Citrix XenServer Features. [www-document]. Available at: <http://blogs.citrix.com/2011/09/30/xenserver-6-0-is-here/> Accessed on 3rd May 2012

Citrix XenServer Features. [www-document]. Available at: <http://www.archy.net/2011/09/27/citrix-xenserver-6/> Accessed on 3rd May 2012

Jamil, D. & Zaki, H. (2011). *Cloud Computing Security*. International Journal of Engineering Science and Technology.

Johnson, D., Kiran, M., Murthy, R., Suseendran, R.B, & Yogesh, G. (2010). *Eucalyptus Beginner's Guide - UEC Edition (Ubuntu Server 10.04 - Lucid Lynx)*. CSS Corp. Pvt. Ltd.

Make IT Green: Cloud Computing and its Contribution to Climate Change. A Green Peace Report. [www-document]. Available at: <http://www.greenpeace.org/usa/Global/usa/report/2010/3/make-it-green-cloud-computing.pdf>. Accessed on 14th April 2012

Marks, E. A., Lozano, B. (2010). *Cloud Computing*. Hoboken, New Jersey: Wiley

Mazon, J (2009). [www-document]. 'On yesterday's email', Official Google Docs Blog. Available at: <http://googledocs.blogspot.com/2009/03/on-yesterdays-email.html>. Accessed on 8th March 2012.

Mell, P. & Timothy, G. (2011). *The NIST Definition of Cloud Computing: Recommendation of the National Institute of Standards and Technology, 800 – 145*. U.S. Department of Commerce Special Publication.

Microsoft Hyper-V Features. [www-document]. Available at: <http://www.learnmsexchange.com/hot-topics/541-microsoft-hyper-v-server-2008-r2-features>. Accessed on 11th April 2012

Microsoft Hyper-V Features. [www-document]. Available at: <http://www.microsoft.com/en-us/server-cloud/windows-server/hyper-v-features.aspx>. Accessed on the 11th April 2012

OpenStack Features. [www-document]. Available at: <http://www.openstack.org/software/essex/> Accessed on 3rd May 2012

Reported cloud outages for Amazon, Google, Microsoft and Salesforce.com in 2008 and 2009. [www-document]. Available at: <http://blog.muoncloud.com/2010/01/31/reported-cloud-outages->

for-amazon-google-microsoft-and-salesforce-com-in-2008-and-2009/. Accessed on 10th April 2012.

Rhoton, J. (2009). *Cloud Computing Explained*. (pp 48 & 50) United States & United Kingdom: Recursive Press

Stanoevska-Slabeva, K., Wozniak, T. & Ristol, S. (2010). *Grid and Cloud Computing: A Business Perspective on Technology and Applications*. (2nd Edition pp 48 – 49 & 52). London, New York: Springer Heidelberg.

Twitter Official blog. Monday Morning Madness. (2009). Available at: <http://blog.twitter.com/2009/01/monday-morning-madness.html>. Accessed on 20th April, 2012.

Vaquero, M. L. Rodero-Merino, L, Caceres, J & Lindner, M. (2009). *A Break in the Clouds: Towards a Cloud Definition*. Published in ACM SIGCOMM Computer Communication Review Newsletter. New York: Association for Computing Machinery.

Victor, D. *Exploring the limits of cloud computing*. Masters Thesis, Kungliga Tekniska Högskolan (KTH), 2010

VMware vSphere Features. [www-document]. Available at: <http://vmguy.com/wordpress/index.php/archives/1674> Accessed on 14th April 2012

VMware vSphere Features. [www-document]. Available at: <http://www.vmware.com/files/pdf/products/vsphere/vmware-what-is-new-vsphere5.pdf> Accessed on 14th April 2012

VMware vSphere Features. [www-document]. Available at: <http://www.vmware.com/products/drs/overview.html> Accessed on 14th April 2012

VMware vSphere Features. [www-document]. Available at: <http://www.vmware.com/products/vsphere/mid-size-and-enterprise-business/features.html> Accessed 14th April 2012

Wikipedia, Eucalyptus. [www-document]. Available at: [http://en.wikipedia.org/wiki/Eucalyptus_\(computing\)](http://en.wikipedia.org/wiki/Eucalyptus_(computing)). Accessed on 10th April 2012.

Wikipedia, Hyper-V. [www-document]. Available at: <http://en.wikipedia.org/wiki/Hyper-V>. Accessed on 10th April 2012.

Williams, Mark I (2010). *Quick Start Guide to Cloud Computing: Moving Your Business into the Cloud*. London: Kogan Page Ltd.