

PROFESSIONAL TRAINING REPORT

Sathyabama Institute of Science and Technology (Deemed to be University)

Submitted in partial fulfilment of the requirements for the award Of Bachelor of
Technology Degree in Information Technology

By

PRAVEEN.N (40120068)



DEPARTMENT OF INFORMATION TECHNOLOGY

SCHOOL OF COMPUTING

**SATHYABAMA INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)**

**Accredited with Grade "A" by NAAC | 12 B Status
by UGC | Approved by AICTE
JEPPIAR NAGAR, RAJIV GANDHISALAI,
CHENNAI – 600119**

APRIL 2023



SATHYABAMA

INSTITUTE OF SCIENCE AND TECHNOLOGY

(DEEMED TO BE UNIVERSITY)


Accredited "A" Grade by NAAC | 12B Status by UGC | Approved by AICTE

www.sathyabama.ac.in


DEPARTMENT OF INFORMATION TECHNOLOGY

BONAFIDE CERTIFICATE

This is to certify that this Project Report is the bonafide work of N. PRAVEEN (40120068) who carried out the project "INSTALLATION OF OPENSTACK IN UBUNTU" under my supervision from January 2023 to April 2023.


Internal Guide


Dr. P. Ajitha


Head of the Department

Dr. R. Subhashini, M.E., Ph.D.



Internal Examiner




External Examiner

DECLARATION

I, Praveen.N (40120068) hereby declare that the Project Report entitled "INSTALLATION OF OPENSTACK IN UBUNTU" is done by me under the guidance of Dr.P.Ajitha and is submitted in partial fulfilment of the requirements for the award of Bachelor of Technology degree in Information Technology.

DATE: 28/4/23

PLACE: Chennai



SIGNATURE OF THE CANDIDATE

ACKNOWLEDGEMENT

I am pleased to acknowledge my sincere thanks to the **Board of Management of SATHYABAMA** for their kind encouragement in doing this project and for completing it successfully. I am grateful to them.

I convey my thanks to **Dr. T Sasikala, M.E., Ph.D.**, Dean, School of Computing
Dr. R Subhashini, M.E., Ph.D., Head of the Department of Information Technology for providing me with necessary support and details at the right time during the progressive reviews.

I would like to express my sincere and deep sense of gratitude to my Project Guide **Dr.P. Ajitha** for her valuable guidance, suggestions and constant encouragement that paved way for the successful completion of my project work.

I wish to express my thanks to all Teaching and Non-teaching staff members of the **Department of Information Technology** who were helpful in many ways for the completion of the project.

CERTIFICATE OF COMPLETION



इलेक्ट्रॉनिक्स एवं
सूचना प्रौद्योगिकी मंत्रालय
MINISTRY OF
ELECTRONICS AND
INFORMATION TECHNOLOGY



Ref.No.: 58 /FSPCC/BC001/04/2023

COURSE COMPLETION LETTER

It is certified that

Praveen N

Sathyabama Institute of Science and Technology

has successfully undergone the Bridge Course Training Program on

Cloud Computing Technology

Under "FutureSkills PRIME" Programme

commenced from 11 Mar, 2023

Shri. V. Sridharan
Chief Investigator
Cloud Computing Technology
Lead Resource Centre

L R PRAKASH

Digitally signed by L R PRAKASH
Date: 2023.04.26 02:18:34
+05'30'

Cdr.(Retd.) L.R. Prakash
Senior Director
Centre Head
C-DAC, Chennai

ABSTRACT

This project report provides a comprehensive guide on installing and configuring Openstack cloud environment that runs on top of Ubuntu on a virtual machine.

The project aims to demonstrate the process of deploying OpenStack which is an Infrastructure as a Service platform that allows users to build and manage private, public, and hybrid cloud environments.

The project's methodology included installing and configuring Ubuntu, on Virtualbox. After successful installing Ubuntu, we will be updating the packages and repositories, followed by finally implementing our methodology to install Openstack with creation of Instance. The results showed that the virtual machine was successfully created, and it enabled the successful installation of Openstack. The project highlights the power and flexibility of Openstack in automating the deployment of cloud resources, which can be beneficial for various industries and organizations.

TABLES OF CONTENTS

NO	TITLE	PAGE NO
	ABSTRACT	vi
	LIST OF FIGURES	x
	LIST OF ABBREVIATION	x
1	INTRODUCTION	1
	1.Introduction	1
2	CLOUD COMPUTING	2
	2.1 Multicloud & Hybrid Multicloud	2
	2.2 Cloud Computing Architecture	3
	2.3 Types of Cloud	3
	2.4 SPI Model	4
	2.5 Virtualisation	7
	2.6 Types of Virtualisation	8
	2.6.1 Operating System Virtualisation	8
	2.6.2 Server & Hardware Virtualisation	9
	2.6.3 Storage Virtualisation	10
	2.7 Advantages of Virtualisation	11
	2.8 Disadvantages of Virtualisation	12
3	VM VIRTUALBOX & UBUNTU	13
	3.1 Pros & Cons of VM Virtualbox	14
	3.2 Ubuntu OS	16
	3.2.1 Pros & Cons of Ubuntu	18
4	DEVSTACK & MICROSTACK	20
	4.1 System Requirements	21
	4.1.1 Windows	21

	4.1.2 Ubuntu	21
	4.1.3 Openstack	21
5	Openstack	22
	5.1 Openstack	22
	5.2 Work for Openstack	24
6	PROJECT IMPLEMENTATION	26
	6.1 Openstack Installation	26
	6.2 Initialisation	27
	6.3 Verification	27
	6.4 Create an Instance	28
	6.5 Connecting to an Instance	29
	6.6 Access the Cloud Dashboard	30
7	CONCLUSION & FUTURESCOPE	31
	7.1 Conclusion	31
	7.2 Future Scope	31
	REFERENCES	33

LIST OF FIGURES

FIGURE NO.	FIGURE NAME	PAGE NO
2.1	Cloud Computing	2
2.4.1	Infrastructure as a Service	4
2.4.2	Platform as a Service	5
2.6.1	Types of Virtualization	8
3.1	VM Virtualbox	13
3.2	Oracle VM Virtualbox Manager	14
3.2.1	Ubuntu Home Screen	17
3.2.2	Opening Terminal in Ubuntu	17
5.1.1	Openstack	22
5.1.2	Overview	23
5.2.1	Login Page of Openstack	25
5.2.2	Openstack Dashboard	25
6.1.1	Microstack Installation	26
6.1.2	Information of Installed Snap	27
6.2	Initialisation of Openstack Services	27
6.3.1	Checking for Image List	28
6.3.2	Viewing Flavor list	28
6.4	Creation of the Instance	29
6.5	Accessing the Instance	30
6.6	Openstack Login Dashboard	30

CHAPTER 1

INTRODUCTION

Cloud computing has revolutionized the way businesses and individuals store, access, and manage their data. The increasing demand for cloud-based solutions has led to the development of several cloud platforms, such as Openstack. Openstack is a free and open-source cloud computing platform that provides a scalable and flexible infrastructure for building private, public, and hybrid cloud environments. It is widely used by organizations of all sizes to create and manage cloud-based solutions. One of the key components of Openstack is the Heat Orchestration Module.

Heat provides a framework for defining and deploying complex cloud environments using templates. Heat templates are YAML files that define the resources needed to create a complete cloud environment, such as virtual machines, volumes, networks, and security groups. Heat templates allow users to automate the process of deploying resources in Openstack, which saves time and reduces errors.

The Heat Orchestration Module is a powerful tool that enables users to automate the deployment of complex cloud environments in Openstack. It is designed to simplify the process of creating and managing cloud-based solutions by providing a simple and flexible framework for defining and deploying resources. Heat is widely used by organizations of all sizes to deploy cloud-based solutions, including web applications, databases, and storage systems.

In this project, we will provide a comprehensive guide on installing and configuring the Heat Orchestration Module in an Openstack cloud environment running on top of Ubuntu. We will also demonstrate the process of deploying a virtual machine with a Cinder volume using a NovaInstance. This project aims to showcase the power and flexibility of Openstack and the Heat Orchestration Module in automating the deployment of cloud resources, which can be beneficial for various industries and organizations.

CHAPTER - 2

CLOUD COMPUTING

In the context of technology, "cloud" refers to a network of remote servers that are hosted on the internet and used to store, manage, and process data. These servers are typically operated by third-party service providers, who offer their resources to businesses and individuals who need them.

Cloud computing allows users to access and use computing resources on-demand, without needing to own or maintain their own hardware or software. This means that businesses and individuals can scale their computing resources up or down as needed, without having to invest in expensive infrastructure.




Fig 2.1 Cloud Computing

2.1 MULTICLOUD AND HYBRID MULTICLOUD

Multicloud is the use of two or more clouds from two or more different cloud providers. Having a multicloud environment can be as simple using email SaaS from one vendor and image editing SaaS from another. But when enterprises talk about multicloud, they are typically talking about using multiple cloud services including SaaS, PaaS, and IaaS services from two or more of the leading public cloud providers.

Hybrid multicloud is the use of two or more public clouds together with a privatecloud environment.

Organizations choose multicloud to avoid vendor lock-in, to have more services to choose from, and to access to more innovation. But the more clouds you use each with its own set of management tools, data transmission rates, and security protocols the more difficult it can be to manage your environment. Multicloud management platforms provide visibility across multiple provider clouds through a central dashboard, where development teams can see their projects and deployments, operations teams can keep an eye on clusters and nodes, and the cyber security staff can monitor for threats.

2.2 CLOUD COMPUTING ARCHITECTURE

Cloud computing architecture refers to the components and sub-components required for cloud computing.

These components typically refer to:

1. Front end (fat client, thin client)
2. Back-end platforms (servers, storage)
3. Cloud-based delivery and a network (Internet, Intranet, Intercloud).

Hosting a cloud: There are three layers in cloud computing. Companies use these layers based on the service they provide.

1. Infrastructure
2. Platform
3. Application

2.3 TYPES OF CLOUD

There are several different types of cloud services, including:

Infrastructure as a Service (IaaS) - which provides access to computing resources such as virtual machines, storage, and networking.

Platform as a Service (PaaS) - which provides a platform for developers to build and deploy their applications, without needing to manage the underlying infrastructure.

Software as a Service (SaaS) - which provides access to software applications over the internet, without needing to install or maintain the software locally.

1. INFRASTRUCTURE AS A SERVICE (IAAS) :

IaaS stands for Infrastructure as a Service, which is a type of cloud computing service that provides users with access to virtualized computing resources over the internet.

With IaaS, users can access and manage virtualized computing resources such as servers, storage, and networking components, without needing to own or maintain their own physical hardware. Instead, the IaaS provider manages the underlying infrastructure, while users have control over the operating systems, applications, and storage they deploy on the virtual machines.

IaaS allows users to quickly and easily scale up or down their computing resources as needed, without having to invest in expensive hardware. This makes it a flexible and cost-effective option for businesses and individuals who need access to computing resources for a variety of purposes, such as hosting websites or running complex applications.

Some popular examples of IaaS providers include Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform.




Fig 2.4.1 Infrastructure As a Service

2. PLATFORM AS A SERVICE (PAAS) :

PaaS stands for Platform as a Service, which is a type of cloud computing service that provides users with a platform for developing, deploying, and managing their applications over the internet.

With PaaS, users can access a pre-configured platform that includes all the necessary infrastructure, middleware, and tools needed to develop and deploy applications, without needing to manage the underlying hardware or software.

PaaS providers typically offer a range of services, such as application hosting, development tools, database management, and integration with other cloud services. Users can access these services through a web-based interface, API, or command line interface, depending on the provider.

PaaS allows users to focus on developing and deploying their applications, without having to worry about the underlying infrastructure. This makes it a popular choice for developers who want to quickly and easily create and deploy applications, without needing to invest in expensive hardware or software.



Fig 2.4.2 Platform As a Service

3. SOFTWARE AS A SERVICE (SAAS) :

SaaS stands for Software as a Service, which is a type of cloud computing service that provides users with access to software applications over the internet.

With SaaS, users can access and use software applications hosted by a third-party provider, without needing to install or maintain the software locally. Users typically access the software through a web browser or mobile app, and pay a subscription fee to the provider for access.

SaaS providers manage the underlying infrastructure, including servers, storage, and networking, as well as software updates and maintenance. This means that users can focus on using the software to accomplish their goals, without having to worry about the underlying technology.

SaaS is a popular choice for businesses and individuals who need access to software applications for a variety of purposes, such as productivity tools, customer relationship management (CRM) software, and accounting software. It allows users to quickly and easily access the software they need, without having to invest in expensive hardware license.



2.5 VIRTUALIZATION

Virtualization refers to the creation of a virtual version of a computing resource, such as a server, operating system, storage device, or network.

Virtualization allows multiple virtual resources to be created and used on a single physical computing resource, such as a server, which can improve resource utilization and flexibility. For example, a single server can be virtualized into multiple virtual servers, each of which can run different operating systems and applications, allowing multiple users or applications to share the physical computing resources.



Fig 2.5 Virtualization

2.6 TYPES OF VIRTUALIZATION

1. Operation system virtualization
2. Server virtualization
3. Hardware virtualization
4. Storage virtualization




Fig 2.6.1 Types Of Virtualization

2.6.1 Operating system virtualization:

Operating system virtualization, also known as containerization, is a type of virtualization that allows multiple isolated user-space instances, or containers, to run on a single operating system kernel.

In this type of virtualization, the host operating system runs on the physical hardware, and multiple containers are created on top of the host operating system. Each container is isolated from the other containers and from the host operating system, and appears to have its own isolated environment with its own resources, including its own file system, network interface, and user accounts.

Operating system virtualization is commonly used in cloud computing and in containerization platforms such as Docker, Kubernetes, and LXC. It provides many benefits, including improved resource utilization, faster deployment and scaling of applications, and better isolation and security of application environments.

Unlike traditional virtualization, where each virtual machine requires its own operating system and resources, operating system virtualization allows multiple applications to share the same operating system kernel, which reduces overhead and improves performance. Additionally, since each container is isolated from the others, application dependencies can be easily managed and conflicts can be avoided, making it easier to manage and deploy applications at scale.

2.6.1 Server virtualization:

Server virtualization is a type of virtualization that allows multiple virtual servers to run on a single physical server, by creating multiple isolated virtual environments, each of which appears to be a dedicated physical server.

In server virtualization, a hypervisor is installed on the physical server, which enables the creation of multiple virtual machines, each with its own operating system, applications, and resources. The hypervisor abstracts the underlying hardware and presents it to each virtual machine as if it were dedicated hardware, allowing each virtual machine to operate.

Server virtualization provides many benefits, including improved resource utilization, flexibility, and scalability. By consolidating multiple physical servers into a single virtual server, organizations can reduce hardware costs, power consumption, and data center footprint. Additionally, virtual machines can be easily migrated between physical servers, which allows for better load balancing and disaster recovery.

Server virtualization is commonly used in cloud computing, data centres, and virtual desktop infrastructure (VDI). Some popular server virtualization platforms include VMware, Hyper-V

2.6.2 Hardware virtualization:

Hardware virtualization is a type of virtualization that enables the creation of multiple virtual machines (VMs) on a single physical machine, by abstracting and partitioning the underlying hardware resources.

Hardware virtualization requires the use of a hypervisor, which is a layer of software that creates and manages the virtual environment. The hypervisor sits between the physical hardware and the virtual machines, and enables each virtual machine to access a portion of the underlying hardware resources, such as CPU, memory, storage, and network connectivity.

Hardware virtualization provides many benefits, including improved resource utilization, better isolation and security, and increased flexibility and scalability. By creating multiple virtual machines on a single physical machine, organizations can reduce hardware costs, energy consumption, and data centre footprint.

Hardware virtualization is commonly used in data centres, cloud computing, and virtual desktop infrastructure (VDI), among other applications. Some popular hardware virtualization platforms include VMware, Hyper-V, and KVM.

2.6.3 Storage virtualization:

Storage virtualization is a type of virtualization that abstracts physical storage resources, such as disks, tapes, and flash drives, and creates a logical view of those resources that can be accessed by users or applications.

Storage virtualization allows for the creation of a pool of storage resources that can be shared by multiple users or applications, regardless of the physical location of the storage devices. This improves storage utilization and allows for more efficient use of storage resources.

Storage virtualization can be implemented using hardware or software. Hardware-based storage virtualization uses specialized hardware devices, such as storage arrays, to virtualize storage resources. Software-based storage virtualization uses software applications, called virtual storage appliances, to provide virtualized storage resources.

Storage virtualization provides many benefits, including improved storage utilization, easier management of storage resources, and increased flexibility and scalability. It can be used in a variety of applications, including cloud computing, data centers, and virtual desktop infrastructure (VDI).

Some popular storage virtualization platforms include VMware Virtual SAN, Microsoft Storage Spaces, and IBM Spectrum Virtualize.

2.7 ADVANTAGES OF VIRTUALIZATION

Virtualization and cloud computing are often used together to deliver scalable and efficient IT services. The advantages of virtualization and cloud computing combined include:

Scalability: Virtualization allows for the efficient use of resources, and cloud computing provides scalable computing resources. Together, they enable the quick and easy provisioning of resources to meet changing business needs.

Cost Savings: By consolidating multiple virtual machines on a single physical server, virtualization reduces the number of physical servers needed. This leads to lower hardware costs, and when combined with cloud computing, organizations can benefit from a pay-as-you-go model, paying only for the resources they need.

Flexibility: Cloud computing and virtualization provide flexibility by enabling resources to be dynamically allocated and adjusted as needed. This allows businesses to quickly respond to changing business needs and demands.

Increased Security: Virtualization provides a secure, isolated environment for applications and data, and cloud computing provides robust security measures to protect against cyber threats.

Disaster Recovery: Virtualization simplifies disaster recovery by allowing virtual machines to be easily backed up, replicated, and restored. Combined with cloud computing, it provides a robust disaster recovery solution, ensuring business continuity in the event of a disaster or hardware failure.

Improved Resource Utilization: Virtualization and cloud computing together provide better resource utilization by enabling multiple virtual machines to share the same physical resources. This leads to improved performance and reduced energy consumption.

Simplified Management: Virtualization and cloud computing provide a centralized management interface for all virtual machines and resources. This simplifies management, making it easier to deploy, manage, and monitor multiple virtual machines and cloud resources from a single console.

2.8 DISADVANTAGES OF VIRTUALIZATION

While virtualization and cloud computing offer many advantages, there are also some disadvantages to consider. Some of the disadvantages of virtualization and cloud computing include:

Complexity: The deployment of virtualization and cloud computing can be complex, requiring specialized skills and expertise. The implementation process can also be time-consuming, leading to potential delays and disruptions.

Cost: While virtualization and cloud computing can offer cost savings, there may be additional costs associated with hardware, software, and maintenance. Organizations must carefully evaluate the costs and benefits of virtualization and cloud computing to determine if they are cost-effective.

Security Risks: Virtualization and cloud computing create new security risks that must be carefully managed. Virtual machines can be more vulnerable to cyber-attacks, and cloud providers can be targeted by hackers, leading to data breaches and other security incidents.

Performance Issues: Virtualization and cloud computing can lead to performance issues, especially if resources are not allocated properly. This can result in slower application performance and decreased productivity.

Dependency on Internet Connectivity: Cloud computing requires a stable internet connection to access services and data stored in the cloud. In the absence of a stable internet connection, the user may experience downtime or loss of access to their data.

Vendor Lock-In: Moving data and applications from one cloud provider to another can be challenging and costly. This can create a dependency on the cloud provider and limit the user's flexibility in terms of choosing other providers.

CHAPTER 3

VM VIRTUALBOX & UBUNTU

VirtualBox is a free and open-source virtualization software that enables users to run multiple operating systems on a single physical computer. It was developed by Oracle and is available for Windows, Mac OS X, Linux, and Solaris.

VirtualBox works by creating a virtual machine (VM) on the host computer, which acts as a separate computer system with its own hardware resources, including CPU, RAM, storage, and network interfaces. Users can install and run any operating system, such as Windows, Linux, or macOS, on the virtual machine, just as they would on a physical computer.

VirtualBox provides a range of features, including support for 3D graphics, USB device integration, and snapshots, which allow users to save the state of the virtual machine and restore it later. It also supports different networking modes, allowing users to connect virtual machines to the internet, local networks, or isolated networks.

VirtualBox is a popular tool for developers, system administrators, and IT professionals, as it enables them to test and run multiple operating systems and applications on a single physical computer. It is also used by individuals who want to run legacy applications or experiment with different operating systems without affecting their primary system.



Fig:3.1 VM virtual box



Fig:3.2 oracle VM Virtual Box Manager

3.1 PROS AND CONS of VM virtual box

Pros:

- **Cost-effective:** VirtualBox is free and open source, which makes it an attractive option for individuals and small businesses that want to use virtualization technology without incurring additional expenses.
- **Easy to use:** VirtualBox has a user-friendly interface and is relatively easy

- **Resource-efficient:** VirtualBox uses minimal system resources, which means it can run on almost any hardware configuration. This makes it a good option for users who want to run virtual machines on older or less powerful hardware.
- **Flexible:** VirtualBox allows users to create and configure virtual machines with a high degree of flexibility, including customizing hardware settings and network configurations.
- **Snapshot and cloning:** VirtualBox allows users to take snapshots of their virtual machines, which can be useful for testing and troubleshooting purposes. Users can also clone virtual machines to create multiple copies with the same configuration.

Cons:

- **Performance:** VirtualBox may not perform as well as other virtualization solutions, such as VMware or Hyper-V, especially when running resource-intensive applications.
- **Limited features:** VirtualBox has fewer advanced features compared to other virtualization solutions, such as support for nested virtualization, which can limit its functionality for some users.
- **Security:** VirtualBox may pose security risks if not properly configured. For example, if the virtual machines are not isolated from the host system, malicious code or viruses may be able to spread between the two systems.
- **Support:** While VirtualBox has an active user community, official support from Oracle is limited, which may make it difficult to get help with more complex issues.

3.2 Ubuntu

Ubuntu is a free and open-source Linux-based operating system that is widely used for desktop, server, and cloud computing applications. It is developed by Canonical and is based on the Debian architecture. Ubuntu is known for its user-friendly interface and ease of use, making it popular among users who are new to Linux or open-source software. It is also highly customizable and supports a range of software and development tools, including Python, Java, and Ruby.

Ubuntu provides a wide range of applications and tools, including LibreOffice, Firefox, and the GNOME desktop environment. It also supports a wide range of hardware, making it suitable for use on a variety of devices, from desktops and laptops to servers and IoT devices. Ubuntu releases new versions every six months, with Long Term Support (LTS) versions released every two years. LTS versions are supported for five years, while regular versions are supported for nine months. Ubuntu has a strong community of users and developers who contribute to its development and support. It is widely used in enterprise environments, as well as by individuals and organizations who value open-source software and the benefits it provides.

Ubuntu has a software centre that makes it easy to find and install applications. The software centre has a wide range of applications, including popular ones like LibreOffice, Firefox, and VLC media player. Ubuntu has built-in security features, including the AppArmor security framework, which helps protect against malware and other security threats. Ubuntu has a large and active community of users and developers who provide support and contribute to its development. This community provides a wealth of resources, including forums, documentation, and tutorials.

Ubuntu is highly customizable, allowing users to personalize their experience and tailor it to their needs. This includes the ability to customize the desktop environment, themes, and icons.



Fig:3.2.1 ubuntu home screen




Fig:3.2.2 Opening terminal in ubuntu

3.3 Pros and Cons of Ubuntu

Pros:

- Free and open-source: Ubuntu is free to download and use, and its open-source nature means that the code is transparent and can be modified and shared.
- User-friendly: Ubuntu has a user-friendly interface and is relatively easy to set up and use, making it a good option for both new and experienced users.
- Wide range of software: Ubuntu has access to a wide range of open-source software that can be installed with a few clicks. This makes it a versatile option for many users, from web developers to gamers.
- Stable: Ubuntu is known for its stability, with regular updates and security patches that help to keep the system running smoothly.
- Customizable: Ubuntu can be customized to suit individual user preferences, from the desktop environment to the software installed.

Cons:

- Compatibility issues: While Ubuntu has a wide range of software available, not all software is compatible with Linux. This can be a problem for users who rely on specific programs or applications that are only available for Windows or macOS.
- Limited support: While Ubuntu has an active community of users who can offer help and advice, official support from Canonical, the company behind Ubuntu, is limited and may not be sufficient for more complex issues.

- Learning curve: Ubuntu may require a bit of a learning curve for users who are not familiar with Linux or command-line interfaces. This can be a barrier for some users who are used to more user-friendly operating systems.
- Hardware compatibility: Ubuntu may not be compatible with all hardware configurations, especially older or less common hardware. This can make it difficult to use Ubuntu on some systems.
- Updates: While Ubuntu's regular updates are a pro for stability and security, they can also cause compatibility issues with some hardware or software configurations. Users may need to check for updates and troubleshoot any issues that arise after an update.

CHAPTER 4

ABOUT DEVSTACK & MICROSTACK

DevStack is a tool for deploying OpenStack, an open-source cloud computing platform, on a single machine for development and testing purposes. It provides an easy and efficient way to set up a complete OpenStack environment on a local machine or a virtual machine. DevStack is maintained by the OpenStack community and is designed to be used by developers, testers, and users who want to experiment with OpenStack without the need for a full-scale deployment. It can be used to test new features, develop applications that integrate with OpenStack, and learn about the platform.

DevStack installs and configures all the necessary components of OpenStack, including Keystone (identity service), Glance (image service), Nova (compute service), Neutron (networking service), and Horizon (dashboard). It also supports a range of optional components, including Swift (object storage), Cinder (block storage), and Heat (orchestration). DevStack is designed to be easy to install and configure, with a simple script that automates the process. It is also highly customizable, allowing users to configure various aspects of the OpenStack environment, such as the networking topology, the number of compute nodes, and the storage backend.

Overall, DevStack is a powerful tool for deploying and testing OpenStack, providing a simple and efficient way to set up a complete OpenStack environment on a single machine. It is widely used by developers and testers in the OpenStack community and is a valuable tool for anyone interested in learning about or working with OpenStack.

Although we have looked at DevStack tool for deploying and testing OpenStack, for our ease, we will be using MicroStack, which is a pure upstream Openstack distribution designed for small scale and edge deployments, that can be installed and maintained with a minimal effort.

4.1 System Requirements

4.1.1 WINDOWS

- Microsoft® Windows® 10 64-bit
- 4 GB RAM minimum, 8 GB RAM recommended
- GB of available disk space minimum,
- GB Recommended 500 MB for IDE + 1.5 GB for Android SDK and emulator system image
- 1280 x 800 minimum screen resolution

4.1.2 UBUNTU

- 2Ghz dual core processor
- 8Gib RAM (system memory)
- 25 GB of hard drive space (or USB stick, memory card or external drive but see LiveCD for an alternative approach)
- VGA capable of 1024x768 screen resolution
- Either a CD/DVD drive or a USB port for the installer media
- Internet access is helpful

4.1.3 OPENSTACK

- multi-core processor
- at least 8 GiB of memory
- 100 GiB of disk space.
- Ubuntu 18.04 LTS or Ubuntu 20.04 LTS.

CHAPTER 5

5.1 OPENSTACK

OpenStack is an open-source cloud computing platform that allows users to manage and control resources, such as compute, storage, and networking, through a web-based dashboard. It was first launched in 2010 as a joint project between Rackspace Hosting and NASA.

OpenStack is designed to be modular and scalable, and can be used to build public or private clouds. It consists of several core services, including Nova for compute, Cinder for block storage, Neutron for networking, Glance for image management, and Keystone for authentication and authorization.

OpenStack is often used by organizations that require flexibility and control over their cloud infrastructure, as it allows them to customize and manage their resources to meet specific requirements. Additionally, being open-source, it provides a high degree of transparency and encourages collaboration among developers.



Fig:5.1.1 openstack

Private cloud: OpenStack can be used to build a private cloud infrastructure for organizations to manage their own resources, applications and data.

Public cloud: OpenStack can be used by service providers to build and offer public cloud services to customers.

Hybrid cloud: OpenStack can be used to build a hybrid cloud infrastructure that combines private and public cloud resources, allowing organizations to scale up or down as needed.

DevOps: OpenStack can be used to provide developers and IT teams with on-demand access to infrastructure resources, enabling faster application development and testing.

Big Data: OpenStack can be used to provide scalable and flexible infrastructure for big data processing, storage and analysis.

Network function virtualization (NFV): OpenStack can be used to build an NFV infrastructure that allows for virtualization of network functions, such as firewalls and load balancers.

High-performance computing (HPC): OpenStack can be used to provide high-performance computing infrastructure, which can be used for scientific research and simulations.

Overall, OpenStack is a versatile and flexible platform that can be used for a wide range of cloud computing use cases




Fig: 5.1.2 Overview

5.2 WORK FOR OPENSTACK

Here are some of the main components and work of OpenStack:

Compute (Nova): This component provides the virtualization layer and manages the creation, scheduling, and termination of virtual machines (VMs) on the cloud infrastructure.

Storage (Cinder and Swift): These components provide different types of storage services. Cinder provides block storage, which is used to store persistent data for VMs, while Swift provides object storage, which is used for storing large volumes of unstructured data.

Networking (Neutron): This component provides network connectivity for the VMs and enables the creation of virtual networks, routers, and firewalls.

Identity (Keystone): This component provides authentication and authorization services and manages user and project information for the OpenStack environment.

Dashboard (Horizon): This component provides a web-based graphical user interface (GUI) for managing OpenStack resources.

Orchestration (Heat): This component provides a framework for automating the deployment and scaling of OpenStack resources.

Telemetry (Ceilometer): This component provides monitoring and metering services to track resource usage and performance.

OpenStack works by providing a set of APIs that enable users to interact with the various OpenStack components and manage their cloud resources. Users can deploy OpenStack on their own hardware or use a public cloud service that runs on OpenStack. OpenStack is highly flexible and can be customized to meet the specific needs of different organizations, making it a popular choice for building and managing cloud infrastructure.



openstack.


Log in

User Name

Password

Sign In

Fig:5.2.1 Login Page of OpenStack


admin ▾
admin ▾

Project ▾
Compute ▾
Overview
Instances
Volumes
Images
Key Pairs
API Access
Network >
Orchestration >
Admin >
Identity >

Project / Compute / Instances

Instances

Instance ID ▾
Filter
Launch Instance
Delete Instances
More Actions ▾

☐	Instance Name	Image Name	IP Address	Flavor	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
No items to display.											

Fig 5.2.2 OpenStack Dashboard

CHAPTER 6

PROJECT IMPLEMENTATION

6.1 Openstack Installation

- Since Microstack is in BETA Phase, we will be using the beta channel:
 - Sudo snap install microstack --beta




Fig:6.1.1 Microstack installation


- To check the information on the installed snap, it can be viewed with the following command:
 - `snap list microstack`

```
idris@idris:~$ snap list microstack
Name          Version  Rev  Tracking    Publisher  Notes
microstack    ussuri   245  latest/beta canonical** -
```

Fig:6.1.2 Microstack information on installed snap

6.2 Initialisation

- The initialisation step automatically deploys, configures, and starts OpenStack services.
- In particular, it will create the database, networks, an image, several flavors, and ICMP/SSH security groups.
- This can all be done within 10 to 20 minutes depending on your machine:
 - `sudo microstack init --auto --control`



```
idris@idris:~$ sudo snap install microstack --beta
[sudo] password for idris:
Run install hook of "microstack" snap if present
Run install hook of "microstack" snap if present
microstack (beta) ussuri from Canonical** installed
idris@idris:~$ snap list microstack
Name          Version  Rev  Tracking    Publisher  Notes
microstack    ussuri   245  latest/beta canonical** -
idris@idris:~$ sudo microstack init --auto --control
2023-04-26 21:03:51,408 - microstack_init - INFO - Configuring clu
stering ...
2023-04-26 21:03:51,530 - microstack_init - INFO - Setting up as a
control node.
2023-04-26 21:03:53,382 - microstack_init - INFO - Generating TLS
Certificate and Key
2023-04-26 21:03:53,992 - microstack_init - INFO - Configuring net
working ...
2023-04-26 21:03:58,094 - microstack_init - INFO - Opening horizon
dashboard up to *
2023-04-26 21:03:58,731 - microstack_init - INFO - Waiting for Rab
bitMQ to start ...
Waiting for 10.0.2.15:5672
2023-04-26 21:04:04,040 - microstack_init - INFO - RabbitMQ starte
```

Fig:6.2 Initialisation of Openstack Services

6.3 Verification

- The purpose of the verification step is to confirm that the cloud is in working order and to discover some of the defaults used by MicroStack. Verification will consist of the following actions:
 - perform various OpenStack queries
 - create an instance
 - connect to the instance over SSH
 - access the cloud dashboard
- Firstly, we can check for the image list by using the following command:
 - `microstack.openstack image list`

```
idris@idris:~$ microstack.openstack image list
+-----+-----+-----+
| ID | Name | Status |
+-----+-----+-----+
| 74aecb01-9ad3-48cd-8d14-42dc823c1af3 | cirros | active |
+-----+-----+-----+
```

Fig: 6.3.1 Checking for image list

- Last but not the least, we can check for the flavor list by using the following command:
 - `microstack.openstack flavor list`

```
idris@idris:~$ microstack.openstack flavor list
+-----+-----+-----+-----+-----+-----+-----+
| ID | Name | RAM | Disk | Ephemeral | VCPUs | Is Public |
+-----+-----+-----+-----+-----+-----+-----+
| 1 | m1.tiny | 512 | 1 | 0 | 1 | True |
| 2 | m1.small | 2048 | 20 | 0 | 1 | True |
| 3 | m1.medium | 4096 | 20 | 0 | 2 | True |
| 4 | m1.large | 8192 | 20 | 0 | 4 | True |
| 5 | m1.xlarge | 16384 | 20 | 0 | 8 | True |
+-----+-----+-----+-----+-----+-----+-----+
```

Fig: 6.3.2 Viewing the flavor list

6.4 Create an Instance

- MicroStack comes with a convenient instance creation command called `microstack launch`. It uses the following defaults for its instances:
 - keypair 'microstack'
 - flavor 'm1.tiny'
 - floating IP address on subnet '10.20.20.0/24'
- To create an instance named 'test' based on the 'cirros' image:
 - `microstack launch cirros -n test`

```
idris@idris:~$ microstack launch cirros -n test
Creating local "microstack" ssh key at /home/idris/snap/microstack/common/.ssh/id_microstack
Launching server ...
Allocating floating ip ...
Server test launched! (status is BUILD)

Access it with `ssh -i /home/idris/snap/microstack/common/.ssh/id_microstack cirros@10.20.20.83`
You can also visit the OpenStack dashboard at https://10.20.20.1:443
```

Fig:6.4 Creation of the Instance

- The `microstack launch` command also supports arguments `--key`, `--flavor`, `--image`, and `--net-id`, in which case you will need to create objects using the standard client if non-default values are desired.

6.5 Connecting to the Instance

- Output from the `microstack launch` command includes all the information needed to connect to the instance over SSH:
 - We can use the command `ssh -i /home/idris/snap/microstack/common/.ssh/id_microstack cirros@10.20.20.83`

```
Access it with `ssh -i /home/idris/snap/microstack/common/.ssh/id_microstack cirros@10.20.20.83`  
You can also visit the OpenStack dashboard at https://10.20.20.1:443
```

Fig:6.5 Accessing the Instance

6.6 Access the Cloud Dashboard

- You can log in to the web UI by pointing your browser to the following URL:
 - <https://10.20.20.1>




Fig: 6.6 Openstack Login Dashboard

And then we will get the login page of openstack

CHAPTER 7

CONCLUSION & FUTURE SCOPE

7.1 CONCLUSION

OpenStack is a popular open-source software platform used to build and manage private and public cloud infrastructures. The project was launched in 2010 and has grown to become one of the most widely used cloud computing platforms in the world. In conclusion, OpenStack offers many benefits to organizations looking to build their own private or public cloud infrastructure, including flexibility, scalability, and cost-effectiveness. With a large and active community of developers and users, OpenStack is constantly evolving to meet the changing needs of cloud computing. However, implementing and managing OpenStack can be complex and requires a high level of expertise. Organizations looking to adopt OpenStack should carefully consider their specific needs, resources, and capabilities before making a decision. It is also important to have a clear understanding of the potential challenges and risks associated with OpenStack, such as security issues and interoperability with other cloud platforms.

Overall, OpenStack is a powerful and versatile platform for building and managing cloud infrastructures, but it requires careful planning and expertise to implement successfully.

7.2 FUTURE SCOPE

As we know, cloud computing is getting wide in the last decade. Most of the biggest giants investing heavily to leverage the growing trend, such as Google, AWS, Microsoft, or Rackspace.

OpenStack is Open Source freely available with Apache 2.0 license and is backed by thousands of big technology companies. Over the years, it is kept on improving its popularity. Hundreds of enterprises are using it to build their own private cloud, and some are delivering support services. So, career-wise, in current business situations, I would suggest it is one of the best innovations to look for. The future calls for such a stable

environment, and as a free architecture with some excellent long-term features, OpenStack will continue its legacy.

If you're looking for something that allows you to make better use of your current hardware, doesn't lock you into costly support contracts and has a wide global engineering community (over 100,000 members across 187 countries and increasing year on year by around 33 percent), then OpenStack could be the solution for you.

One thing is to remain there for secure cloud computing, with OpenStack getting a promising future for sure in at least the next ten years. OpenStack is a bundle of numerous cutting-edge items that offer the best experience for cloud computing. So learning it and gaining experience in a few projects would not only make you a seasoned professional in the cloud but also help you develop skills in problem-solving.

REFERENCES

1. Pettey, C. and L. Goasduff. Gartner Executive Programs Worldwide Survey of More Than 2,000 CIOs Identifies Cloud Computing as Top Technology Priority for CIOs in 2011.
2. 2011[cited2012 04/06/2012]. <http://www.gartner.com/it/page.jsp?id=1526414>.
3. P., M. and T. Grance., The NIST Definition of Cloud Computing. 2011.
4. Sabahi, F. Cloud computing security threats and responses. in Communication Software and Networks (ICCSN), 2011 IEEE 3rd International Conference on. 2011.
5. Canonical. Ubuntu Cloud: Technologies for future-thinking companies. 2012.
6. <http://www.canonical.com/about-canonical/resources/white-papers/ubuntu-cloud-technologies-future-thinking-companies>.
7. <https://cloudsecurityalliance.org/research/security-guidance/>.
Eucalyptus Web Pages (OpenSource). <http://open.eucalyptus.com>.
Nimbus Project Web Page. <http://www.nimbusproject.org>
OpenNebula Web Page. <http://www.opennebula.org>.