

INTERDISCIPLINARY PROJECT REPORT at Sathyabama Institute of Science and Technology (Deemed to be University)

Submitted in partial fulfillment of the requirements for the award of
Bachelor of Engineering Degree in
Computer Science and Engineering

By
DIVYADHARSINI
(Reg. No. 40120030)



**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 GANDHI SALAI,
CHENNAI – 600119

APRIL 2023

DEPARTMENT OF INFORMATION TECHNOLOGY



SATHYABAMA

INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)


Accredited with Grade "A" 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 DIVYADHARSINI (40120030) who carried out the project entitled "SAMPLE APPLICATION WITH APACHE AND MYSQL USING USER-DATA SCRIPT OPTION OF OPENSTACK" under my supervision from FEB 2023 to APRIL 2023.


Internal Guide

Dr. P. AJITHA, M.E., PH.D.,


Head of the Department
Dr. R. Subhashini M.E., Ph.D.

Submitted for viva voce examination held on 29-04-2023


Internal Examiner



External Examiner

DECLARATION

I, DIVYADHARSINI hereby declare that the project report entitled **Sample Application with Apache and MYSQL using User-Data script option of Openstack** done by me under the guidance of **Dr. P. AJITHA, M.E.,PH.D.**, is submitted in partial fulfillment of the requirements for the award of Bachelor of Technology Degree in Information Technology.

DATE : 29-04-2023

PLACE : Chennai


SIGNATURE OF THE CANDIDATE

ACKNOWLEDGEMENT

I am pleased to acknowledge my sincere thanks to **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, M.E., Ph.D.**, for his valuable guidance, suggestion and constant encouragement 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



रा.इ.सू.प्रौ.सं
NIELIT

Ref.No.: 71 /FSPCC/BC001/04/2023

COURSE COMPLETION LETTER

It is certified that
Divyadharsini P S
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

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

TABLE OF CONTENTS

CHAPTER NO.	TITLE	PAGE NO.
	ABSTRACT	I
	LIST OF FIGURES	II
1	INTRODUCTION	1-3
	1.1 AIM	2
	1.2 SCOPE	2
2	INSTALLATION & SETUP	3-13

	2.1 Virtual Box	3-7
	2.2 Ubuntu	8-10
	2.3 Open stack	11-13
3	METHODOLOGY AND ALGORITHMS USED	14-19
	3.1 System Requirements	14-15
	3.2 Methods & Algorithms Used	16-19
4	RESULT, PERFORMANCE ANALYSIS	20-24
	4.1 Login User Code in OpenStack	20-21
	4.2 Integrating MySQL Enterprise with Trove	22-23
	4.3 Result	24
5	SUMMARY & CONCLUSION	25-26
	5.1 Summary	25
	5.2 Conclusion	25
	5.3 Future Work	26
	REFERENCES	27

ABSTRACT

OpenStack is a free, open standard cloud computing platform. It is mostly deployed as infrastructure-as-a-service (IaaS) in both public and private clouds where virtual servers and other resources are made available to users. The software platform consists of interrelated components that control diverse, multi-vendor hardware pools of processing, storage, and networking resources throughout a data center . Users manage it either through a web-based dashboard, through command-line tools, or through restful web services.

Unlike traditional virtualisation management platforms, such as VMware vSphere or Red Hat Virtualization Manager, OpenStack is a fully functional cloud platform as

defined by the [National Institute of Standards and Technology](#). This means that OpenStack basically resembles the behaviour of public clouds. Users can request compute, network and storage resources on-demand through a self-service portal. Although they are provided in the form of virtual machines, virtual networks and virtual disks, exactly as in the traditional virtualisation management platforms, they are defined through the APIs. For businesses, the ability to define virtualised resources programmatically enables fast-paced infrastructure automation and cloudstyle operations.

OpenStack is an open-source cloud platform designed to manage distributed compute, network and storage resources in the data centre. In principle, OpenStack aggregates physical resources into one big pool and allocates virtual resources out of this pool to users who can request them on-demand through a self-service portal or application programming interfaces (APIs). But OpenStack itself does not handle virtualisation. Instead, it leverages the existing virtualisation technologies. Therefore, OpenStack is more like a wrapper around traditional virtualisation tools, enabling cloud-native capabilities.

I

LIST OF FIGURES

FIGURE NO.	FIGURE NAME	PAGE NO.
1.1	Download Virtual Box	3
1.2	Install Virtual Box	4
1.3	Install Ubuntu OS	10
1.4	Install OpenStack	11
2.1	Methods & Algorithm Used	10
4.1	Login User Code in OpenStack	20

4.2	MySQL Enterprise & Trove	23
5.3	Future Work	26

II

CHAPTER 1 INTRODUCTION

The work mentioned in this report has been done under the course CDAC-Cloud Computing, internship provided by Nasscom. OpenStack is a versatile cloud computing platform that can be used to build a wide range of applications. When combined with the popular open-source database management system MySQL and the web server Apache, OpenStack becomes a powerful tool for building web applications.

MySQL is an open-source relational database management system that provides a robust and scalable solution for storing and retrieving data. It is widely used in web applications and is an excellent choice for OpenStack deployments that require a reliable and scalable database.

Apache is a popular web server that is commonly used for hosting web applications. It provides a flexible and modular architecture that allows for the easy integration of different components, such as PHP, Python, and Ruby.

When combined with OpenStack, MySQL and Apache provide a powerful platform for building scalable and flexible web applications. OpenStack provides the infrastructure layer, allowing you to easily deploy and manage virtual machines, storage, and networking resources, while MySQL and Apache provide the application layer, enabling you to build web applications that can handle large amounts of data and traffic.

Overall, OpenStack with MySQL and Apache is an excellent choice for building web applications that require scalability, reliability, and flexibility.

1

1.1 AIM

To build a sample application in Ubuntu with Apache and MYSQL using data script option. The open stack mission is to produce a ubiquitous open source cloud computing platform that will meet the needs of public and private clouds regardless of size, by being simple to implement and massively scalable. OpenStack is an open source platform that uses pooled virtual resources to build and manage private and public clouds. The tools that comprise the OpenStack platform, called "projects," handle the core cloud-computing services of compute, networking, storage, identity, and image services.

1.2 SCOPE

Cloud: “ Single Most Disruptive Technology”

Since 2000, 52% of the companies in the Fortune 500 have either gone bankrupt, been acquired, ceased to exist, or dropped off the list. Digitalization of business is a key factor in this accelerated pace of change. Information flows faster.

Cloud is the foundation for digital transformation:

- Ubiquity and ease of adoption
- Unlimited and dynamic capacity
- Helps you innovate faster

2

CHAPTER-2 INSTALLATION & SETUP

1. VIRTUAL BOX

If we want to install the VirtualBox in windows, we can use the following steps:

Step 1: First, we need to go to the **VirtualBox Website** through the <https://www.virtualbox.org/> in the system's internet browser. It is the link where we will download the VirtualBox setup file.

Step 2: Click on the **Download VirtualBox** button, which appears on a blue button on the page.



Step 3: Now, Click on the **Windows hosts** link. We will see it below the "VirtualBox 6.1.16 platform packages" option.





Fig 1.1 : Download Virtual Box

Step 4: Now, open the **VirtualBox EXE** file from where we have downloaded this file in the system. After that, the VirtualBox installation window will open.




Step 5: Click on the **Next** button in the Welcome window.



Fig 1.2 : Install Virtual Box

Step 6: Choose the **installation folder** and click on the **Next** button.



Step 7: Now, we will choose the features that we want to install and then click on the **Next** button.




Fig 1.2 : Install Virtual Box

Step 8: Click on the **Yes** button to install the Oracle VirtualBox interfaces.



Step 9: Click on the **Install** button when it prompted. After that, it will start installing the VirtualBox in our system.




Fig 1.2 : Install Virtual Box

Step 10: After completing all these processes, click on the **Finish** button. When we do this, the installation tab will be closed, and VirtualBox will be opened. Now, we can create a virtual machine to run any OS on the system.



Fig 1.2 : Install Virtual Box

2. UBUNTU

This chapter provides an overview of how to install Ubuntu Server Edition. You can also refer to this guide on [how to operate the installer](#) for more information on using the installer, and to this [screen-by-screen reference guide](#) for more information about each of the installer screens.

Preparing to install

Ubuntu Server Edition provides a common, minimalist base for a variety of server applications, such as file/print services, web hosting, email hosting, etc. This version supports four 64-bit architectures:

amd64 (Intel/AMD 64-bit) arm64
(64-bit ARM) ppc64el (POWER8
and POWER9) s390x (IBM Z and
LinuxONE)

The recommended system requirements are:

CPU: 1 gigahertz or better

RAM: 1 gigabyte or more

Disk: a minimum of 2.5 gigabytes

Perform a system back up

Before installing Ubuntu Server Edition you should make sure all data on the system is backed up. If this is not the first time an operating system has been installed on your computer, it is likely you will need to re-partition your disk to make room for Ubuntu.

Any time you partition your disk, you should be prepared to lose everything on the disk should you make a mistake or something goes wrong during partitioning. The programs used in installation are quite reliable, most have seen years of use, but they also perform destructive actions.

Download the server ISO

You can obtain the amd64 server download from <https://releases.ubuntu.com/>.

Select the version you wish to install and select the “server install image” download. Note that the server download includes the installer.

There are platform specific how-to guides for installations on:

[s390x LPAR](#)

[z/VM](#)

[Ppc64el](#)

Create a bootable USB

There are many ways to boot the installer but the simplest and most common way is to [create a bootable USB stick](#) to boot the system to be installed with ([tutorials for other operating systems](#) are also available).

Boot the installer

Plug the USB stick into the system to be installed and start it.

Most computers will automatically boot from USB or DVD, though in some cases this is disabled to improve boot times. If you don't see the boot message and the "Welcome" screen which should appear after it, you will need to set your computer to boot from the install media.

There should be an on-screen message when the computer starts telling you what key to press for settings or a boot menu. Depending on the manufacturer, this could be Escape, F2, F10 or F12. Simply restart your computer and hold down this key until the boot menu appears, then select the drive with the Ubuntu install media.

If you are still having problems, check out the [Ubuntu Community documentation on booting from CD/DVD](#).

After a few moments, the installer will start in its language selection screen.




Fig 1.3 : Install Ubutnu OS

Using the installer

The installer is designed to be easy to use and have sensible defaults so for a first install you can mostly just accept the defaults for the most straightforward install:

- Choose your language
- Update the installer (if offered)

- Select your keyboard layout
- Do not configure networking (the installer attempts to configure wired network interfaces via DHCP, but you can continue without networking if this fails)
- Do not configure a proxy or custom mirror unless you have to in your network
- For storage, leave “ use an entire disk” checked, and choose a disk to install to, then select “ Done” on the configuration screen and confirm the install
- Enter a username, hostname and password
- On the SSH and snap screens, select “ Done”
- You will now see log messages as the install is completed
- Select restart when this is complete, and log in using the username and password provided

3. OPENSTACK

Add user

We need to add a user to install DevStack.

adduser stack

```
root@Openstack:~# adduser stack
Adding user `stack' ...
Adding new group `stack' (1001) ...
Adding new user `stack' (1001) with group `stack' ...
Creating home directory `/home/stack' ...
Copying files from `/etc/skel' ...
Enter new UNIX password:
Retype new UNIX password:
passwd: password updated successfully
Changing the user information for stack
Enter the new value, or press ENTER for the default
  Full Name []: Openstack
  Room Number []:
  Work Phone []:
  Home Phone []:
  Other []:
Is the information correct? [Y/n] y
root@Openstack:~#
```

Fig 1.4 : Install OpenStack

Since this user will be making many changes to your system, it will need to have sudo privileges:

echo "stack ALL=(ALL) NOPASSWD: ALL">> /etc/sudoers

```
root@Openstack:~# echo "stack ALL=(ALL) NOPASSWD: ALL" >> /etc/sudoers
root@Openstack:~#
```

Login in as stack user

su - stack **Download**

DevStack

git clone https://opendev.org/openstack/devstack -b stable/ussuri cd devstack

A terminal window with a dark background and light-colored text. The prompt is 'cloud@openstack:~\$'. The command 'git clone https://opendev.org/openstack/devstack -b stable/ussuri' has been executed. The output shows the cloning progress: 'Cloning into 'devstack'...', 'remote: Enumerating objects: 46244, done.', 'remote: Counting objects: 100% (46244/46244), done.', 'remote: Compressing objects: 100% (21131/21131), done.', 'remote: Total 46244 (delta 32691), reused 37526 (delta 24408)', 'Receiving objects: 100% (46244/46244), 9.51 MiB | 1.41 MiB/s, done.', and 'Resolving deltas: 100% (32691/32691), done.'. The prompt is now 'cloud@openstack:~\$' with a cursor.

```
cloud@openstack:~$ git clone https://opendev.org/openstack/devstack -b stable/ussuri
Cloning into 'devstack'...
remote: Enumerating objects: 46244, done.
remote: Counting objects: 100% (46244/46244), done.
remote: Compressing objects: 100% (21131/21131), done.
remote: Total 46244 (delta 32691), reused 37526 (delta 24408)
Receiving objects: 100% (46244/46244), 9.51 MiB | 1.41 MiB/s, done.
Resolving deltas: 100% (32691/32691), done.
cloud@openstack:~$
```

Fig 1.4 : Install OpenStack

Run DevStack:

./stack.sh

```
++ declare -r -g _DEVSTACK_FUNCTIONS=1
++++ dirname /home/stack/devstack/functions
+++ cd /home/stack/devstack
+++ pwd
++ FUNC_DIR=/home/stack/devstack
++ source /home/stack/devstack/functions-common
++++ set +o
++++ grep xtrace
+++ _XTRACE_FUNCTIONS_COMMON='set -o xtrace'
+++ set +o xtrace
++ source /home/stack/devstack/inc/ini-config
++++ set +o
++++ grep xtrace
+++ INC_CONF_TRACE='set -o xtrace'
+++ set +o xtrace
++ source /home/stack/devstack/inc/meta-config
++++ set +o
++++ grep xtrace
+++ _XTRACE_INC_META='set -o xtrace'
+++ set +o xtrace
```

```
=====
DevStack Component Timing
=====
Total runtime          4321

run_process            146
test_with_retry        12
apt-get-update         62
pip_install            734
restart_apache_server  28
wait_for_service       82
git_timed              240
apt-get                29
=====
```

In browser give your machines ip `http://<ip Address>`




Fig 1.4 : Install OpenStack

Openstack Dashboard

In browser give your machines ip `http://<ip Address>` or `localhost`



3.1 SYSTEM REQUIREMENTS

Software Requirements for Openstack

Supported operating systems:

- Ubuntu Server 14.04 LTS
- RHEL/CentOS
- Debian

System packages for Ubuntu:

- gcc
- python-pip
- python-dev
- libxml2-dev
- libxslt-dev
- libffi-dev
- libpq-dev
- python-openssl
- mysql-client

System packages for CentOS:

- gcc
- python-pip
- python-devel
- libxml2-devel
- libxslt-devel
- libffi-devel
- postgresql-devel
- pyOpenSSL
- mysql

Hardware Requirements for Openstack

We recommend that your system meets the following hardware requirements:

<u>Criteria</u>	<u>Minimal</u>	<u>Recommended</u>
CPU	4 core @ 2.4 GHz	24 core @ 2.67 GHz
RAM	8 GB	24 GB or more
HDD	2 x 500 GB (7200 rpm)	4 x 500 GB (7200 rpm)
RAID	Software RAID-1 (use mdadm as it improves the read performance almost twice)	Hardware RAID-10

Other possible storage configurations:

- 1x SSD 500+ GB
- 1x HDD (7200 rpm) 500+ GB and 1x SSD 250+ GB (install the system onto the HDD and mount the SSD drive to the directory where the virtual machines images are stored)
- 1x HDD (15000 rpm) 500+ GB

3.2 METHODS AND ALGORITHM USED

Uploading Images

A sample image is placed in /root filename is:

Project -> compute -> images Click Create Image Button

openstack

admin

Project

API Access

Compute

Overview

Instances

Images

Key Pairs

Server Groups

Volumes

Network

Project / Compute / Images

Images

Click here for filters or full text search.

Create Image

Delete Images

Displaying 1 item

	Owner	Name	Type	Status	Visibility	Protected	
☐	admin	cirros-0.4.0-x86_64-disk	Image	Active	Public	No	Launch

Displaying 1 item

Create Image

Image Details

Metadata

Image Details

Specify an image to upload to the Image Service.

Image Name

Cirros-0.3.2

Image Description

Cirros-0.3.2 Image

Image Source

File

Browse...

cirros-0.3.2-i386-disk.img

Format

QCOW2 - QEMU Emulator

Image Requirements

Kernel

Choose an image

Ramdisk

Choose an image

Architecture

Minimum Disk (GB)

1

Minimum RAM (MB)

128

Image Sharing

Visibility

Private

Shared

Community

Public

Protected

Yes

No

Cancel

Back

Next

Create Image

Fig 2.1 : Methods & Algo Used

16

Images

Q

Click here for filters or full text search.

x

+ Create Image

Delete Images

Displaying 2 items

☐	Owner	Name ^	Type	Status	Visibility	Protected	
☐	> admin	Cirros-0.3.2	Image	Active	Shared	No	Launch ▾
☐	> admin	cirros-0.4.0-x86_64-disk	Image	Active	Public	No	Launch ▾

Displaying 2 items

Creating Network

Project->Network-> Network Topology, Click Create Network Button.

In Network tab, give the Network Name and Admin State is UP and Click next

Create Network

Network

Subnet

Subnet Details

Network Name

Cloud1

Create a new network. In addition, a subnet associated with the network can be created in the following steps of this wizard.

☒ Enable Admin State ⓘ

☐ Shared

☒ Create Subnet

Availability Zone Hints ⓘ

nova

Cancel

« Back

Next »

In Subnet tab, Give the Subnet Name, Network Address, Select the IP Version as IPV4 and give the gateway IP and click Next.

Create Network

Network

Subnet

Subnet Details

Subnet Name

Cloud1_sub

Network Address Source

Enter Network Address manually ▾

Network Address ⓘ

192.168.100.0/24

IP Version

IPv4 ▾

Gateway IP ⓘ

192.168.100.1

☐ Disable Gateway

Creates a subnet associated with the network. You need to enter a valid "Network Address" and "Gateway IP". If you did not enter the "Gateway IP", the first value of a network will be assigned by default. If you do not want gateway please check the "Disable Gateway" checkbox. Advanced configuration is available by clicking on the "Subnet Details" tab.

Cancel

« Back

Next »

Fig 2.1 : Methods & Algo Used

Network Topology



Creating Instance

openstack

admin

Project

API Access

Compute

Overview

Instances

Images

Key Pairs

Server Groups

Volumes

Network

Admin

Identity

Project / Compute / Instances

Instances

Instance ID = Filter [Launch Instance](#)

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

Launch Instance

Details

Source

Flavor

Networks

Network Ports

Security Groups

Key Pair

Configuration

Server Groups

Scheduler Hints

Metadata

Please provide the initial hostname for the instance, the availability zone where it will be deployed, and the instance count. Increase the Count to create multiple instances with the same settings.

Instance Name *

Cloud

Description

Availability Zone

nova

Count *

1

Total Instances (10 Max)

10%

0 Current Usage

1 Added

9 Remaining

Cancel

Back

Next

Launch Instance

Fig 2.1 : Methods & Algo Used

Select the Network in which instance to be launched and Click Launch Instance

Launch Instance

Details

Source

Flavor

Networks

Network Ports

Security Groups

Key Pair

Configuration

Server Groups

Scheduler Hints

Metadata

Networks provide the communication channels for instances in the cloud.

▼ Allocated 1

Select networks from those listed below.

	Network	Shared	Admin State	Status	
1	Cloud	No	Up	Active	↓

▼ Available 3

Select at least one network

Q

Click here for filters or full text search.

X

	Network	Shared	Admin State	Status	
	Cloud1	No	Up	Active	↑
	public	No	Up	Active	↑
	shared	Yes	Up	Active	↑

X Cancel

< Back

Next >

Launch Instance

Instances

Instance ID =

Filter

Launch Instance

Delete Instances

More Actions

Displaying 1 item

	Instance Name	Image Name	IP Address	Flavor	Key Pair	Status	Availability Zone	Task	Power State	Age	Actions
☐	Cloud	cirros-0.4.0-x86_64-disk	192.168.100.100	m1.nano	-	Active	nova	None	Running	0 minutes	Create Snapshot

Displaying 1 item

Fig 2.1 : Methods & Algo Used

CHAPTER-4

RESULT, PERFORMANCE ANALYSIS

4.1 LOGIN USER CODE IN OPENSTACK

Login as cirros user, then give `sudo -i`.

```

Press Control_L+Alt_L to release pointer. Instance-00000003 on QEMU/KVM
File Virtual Machine View Send Key

[ 4.175202] Write protecting the kernel read-only data: 14336k
[ 4.179241] Freeing unused kernel memory: 1860K (ffff88000182f000 - ffff880001a00000)
[ 4.184259] Freeing unused kernel memory: 168K (ffff880001dd6000 - ffff880001e00000)
[ 4.238203] clocksource: Switched to clocksource tsc

further output written to /dev/ttyS0
[ 9.602568] random: dd urandom read with 11 bits of entropy available

login as 'cirros' user. default password: 'gocubsgo'. use 'sudo' for root.
cloud login: [ 963.327282] random: nonblocking pool is initialized

login as 'cirros' user. default password: 'gocubsgo'. use 'sudo' for root.
cloud login: cirros
Password:
$ sudo -i
# df -h
Filesystem                Size      Used Available Use% Mounted on
/dev                        51.2M          0    51.2M   0% /dev
/dev/vda1                  978.9M    24.0M    914.1M   3% /
tmpfs                      55.2M          0    55.2M   0% /dev/shm
tmpfs                      55.2M    88.0K    55.1M   0% /run
# _

```

Fig 4.1 :Login User Code in Openstack

```

Press Control_L+Alt_L to release pointer. Instance-00000003 on QEMU/KVM
File Virtual Machine View Send Key

Disk /dev/vda: 1 GiB, 1073741824 bytes, 2097152 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: gpt
Disk identifier: 127BB530-4FDD-4B55-B653-77C39F7AE9C4

# mkdir /media/Backup
# mkfs.ext4 /dev/vdb
mke2fs 1.42.12 (29-Aug-2014)
/dev/vdb contains a ext4 file system
created on Sat Aug  8 09:42:08 2020
Proceed anyway? (y,n) y
Creating filesystem with 262144 4k blocks and 65536 inodes
Filesystem UUID: ef12c405-c415-4974-9c1b-4cc8df826e2c
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376

Allocating group tables: done
Writing inode tables: done
Creating journal (8192 blocks): done
Writing superblocks and filesystem accounting information: done

#
#

```

```

Press Control_L+Alt_L to release pointer. Instance-00000003 on QEMU/KVM
File Virtual Machine View Send Key


# df -h
Filesystem                Size      Used Available Use% Mounted on
/dev                      51.2M         0      51.2M   0% /dev
/dev/vda1                 978.9M    24.0M    914.1M   3% /
tmpfs                     55.2M         0      55.2M   0% /dev/shm
tmpfs                     55.2M    88.0K    55.1M   0% /run
#
# mount /dev/vdb /media/Backup/
[ 2100.907364] EXT4-fs (vdb): couldn't mount as ext3 due to feature incompatibilities
[ 2100.924509] EXT4-fs (vdb): couldn't mount as ext2 due to feature incompatibilities
# df -h
Filesystem                Size      Used Available Use% Mounted on
/dev                      51.2M         0      51.2M   0% /dev
/dev/vda1                 978.9M    24.0M    914.1M   3% /
tmpfs                     55.2M         0      55.2M   0% /dev/shm
tmpfs                     55.2M    88.0K    55.1M   0% /run
/dev/vdb                  975.9M     1.3M    907.4M   0% /media/Backup
# _

```

Fig 4.1 :Login User Code in Openstack

OpenStack Momentum

- Launched in 2010, OpenStack has grown to become a leader. Some numbers from <http://www.stackalytics.com/> as of May 2015:
 - Over 400 companies contributed code to OpenStack (Kilo)
 - Almost 2700 individuals contributed code to the project (Kilo)



Getting Started with OpenStack Trove

- As a Trove user
 - OpenStack distribution that includes Trove (Eg. Oracle Openstack)
 - <http://www.oracle.com/us/technologies/linux/openstack/overview/index.htm>
 - !
 - Tesora DBaaS platform, a standalone Trove-based DBaaS product
 - <http://tesora.com/download-tesora-dbaas-platform-community-edition>

4.2 INTEGRATING MySQL ENTERPRISE WITH TROVE

MySQL Enterprise & Trove

- Create a MySQL Enterprise machine image for use with Trove — Pull a standard image from Glance and add/install: MySQL.
- MySQL Enterprise Server
- Default MySQL config (/etc/mysql/my.cnf) that enables all Enterprise Plugins -
 - Can later use a Trove configuration group to sync and manage options across all instances.
- MySQL Enterprise Backup
- MySQL Enterprise Monitor Agent
 - Configured to know how to reach out to the MySQL Enterprise Monitor
- Trove Guest Agent
 - Configured to know how to reach the Trove Server (API, Taskmanager)
 - Customized backup strategy definitions for MySQL Enterprise Backup
- trove guest agent/strategies/backup/mysql impl.py
- See:<http://dev.mysql.com/doc/mysql-enterprisebackup/3.12/en/mysalbackup.usage.html>
- Store customized MySQL Enterprise Trove image in Glance
 - Or skip all of the image related steps
 - Official certified Oracle Linux + MySQL Enterprise Trove images
 - Tesora DBaaS Enterprise Certified MySQL Enterprise Trove images
- Create necessary entries in Keystone
 - User, role, tenant, service, endpoint
- Create MySQL datastore withing Trove
 - Use trove-manage client
 - Sync Trove information to OpenStack' s MySQL data stroe(trove manage)
 - Create MySQL datastre using new Trove machine image in Glance
- Configure MySQL Enterprise Backups
 - Specify MySQL Enterprise Backup as tool for backups(/etc/trove/trove.conf)

- Instances can then be created with trove client or via Horizon web UI

MySQL Enterprise & Trove(cont.)

- Launch MySQL EE instance —Flavor from Nova
 - Cinder volume for the datadir
- Initialize the instance
 - Specify default accounts
 - Specify default schemas
- Optionally restore from a backup.




Fig 4.2 :MySQL Enterprise & Trove

4.3 RESULT

The results and discussion of the Sathyabama budget and expenses database project can be focused on how well the database met the objectives outlined in the project plan.

The database provides valuable insights into the college's financial history and trends, which can help in long-term financial planning. The reporting and analysis module provides data visualizations and analysis tools, making it easier for users to understand financial trends and plan accordingly.

the Sathyabama budget and expenses database project has successfully achieved its objectives of providing efficient budget management, real-time financial information, transparency, cost reduction, compliance, accountability, and future planning capabilities.

CHAPTER-5 SUMMARY & CONCLUSION

5.1 SUMMARY

This project on OpenStack which is an open source platform for cloud computing is built using MySQL and Apache server. It is a sample application which hosts the database in MySQL which is hosted using Apache Server using OpenStack. Ubuntu is used as the host Operating System which is an open source linux based OS. The application helps in maintaining the database over cloud that makes it easily manageable and globally accessible with security and almost 100 percent up time.

5.2 CONCLUSION

Finally, the report presents the results of testing and evaluation of the application, including the identification of any issues or areas for improvement. The project was successful in achieving its objectives and resulted in the development of a fully functional OpenStack cloud application that meets the needs of its users while also being secure and reliable.

In conclusion, this project has demonstrated the importance of designing and developing a secure and user-friendly file sharing web application in today's digital age. The knowledge and experience gained from this project can be applied to future projects in the field of web application development

5.3 FUTURE WORK

Plans for OpenStack

- Liberty (October 2015) and beyond ...
- Queue Service (Zagar)
- Shared File System(Manilla)
- DNS Service (Designate)
- Key management (Barbican)
- Container Service (Docker and Kubernetes: Magnum)
- Docker deployment and management tools (Kolla)

Plans for MySQL Enterprise

- MySQL 5.7...
- Improved scalability, performance, security, and ease-of-use
- MySQL Group Replication
- Easy HA with Active/Active clusters
- Replication & HA Team Blog
- MySQL Router : Load Balancing, transparent failover, aggregation •
- MySQL Fabric : Fabric 1.6 supports multi-node clusters




Fig 5.3 :Future Work

REFERENCES

1. MySQL Enterprise Resources, Manual – <https://dev.mysql.com/doc/refman/5.7/en/>
2. MySQL Enterprise Resources, Additional Info – <https://www.mysql.com/products/enterprise/>
3. MySQL Enterprise Resources, Community forums – <http://forums.mysql.com/>
4. MySQL Enterprise Resources, Server Team Blog – <http://www.mysqlservertime.com>
5. MySQL Enterprise Resources, Replication and HA Team Blog – <http://mysqlhighavailability.com>
6. Oracle OpenStack Resources, Product web page(data sheet, white paper, videos, FAQ) – <http://www.oracle.com/us/tehnologies/linux/openstack/overview/index.html>
7. Oracle OpenStack Resources, Documentaion: Installation and User Guide, HOL – <http://www.oracle.com/technetwork/serverstorage/openstack/linux/documentatio/index.html>
8. Oracle OpenStack Resources, Downloads – <http://www.oracle.com/technetwork/serverstorage/openstack/linux/downloads/index.html>
9. Trove Resources, Docs – <http://docs.openstack.org/developer/trove/>
10. Trove Resources, Wiki – <https://www.wiki.openstack.org/wili/Trove>
11. Trove Resources, Launchpad – <https://Launchpad.net/trove>
12. Trove Resources, Community help - #openstack trove on Freenode 13. Trove Resources, Events – <http://www.tesora.com/troveday/>

