

CLOUD COMPUTING

LABORATORY MANUALS

DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING



RAJA JAIT SINGH GOVERNMENT POLYTECHNIC, NEEMKA

INDEX

S.No.	Practical Name	Date	Sign
1	Introduction to Cloud Vendors: Amazon, Microsoft, IBM		
2	Setting up Virtualization using VirtualBox/VMWare Hypervisor		
3	Introduction to OwnCloud		
4	Installation And Configuration of OwnCloud software for Saas		
5	Accessing Microsoft AZURE cloud Services		
6	Cloud simulation Software Introduction :CloudSim		
7	Create a cluster with AWS		

PRACTICAL-1

Aim: Introduction to Cloud Vendors: Amazon, Microsoft, IBM

Cloud Service providers (CSP) offers various services such as **Software as a Service, Platform as a service, Infrastructure as a service, network services, business applications, mobile applications, and infrastructure** in the cloud. The cloud service providers host these services in a data center, and users can access these services through cloud provider companies using an Internet connection.

Amazon Web Services (AWS)

[AWS](#) (Amazon Web Services) is a **secure cloud service platform** provided by **Amazon**. It offers various services such as database storage, computing power, content delivery, Relational Database, Simple Email, Simple Queue, and other functionality to increase the organization's growth.

Features of AWS

AWS provides various powerful features for building scalable, cost-effective, enterprise applications. Some important [features of AWS](#) is given below-

- AWS is **scalable** because it has an ability to scale the computing resources up or down according to the organization's demand.
- AWS is **cost-effective** as it works on a **pay-as-you-go** pricing model.
- It provides various flexible storage options.
- It offers various **security services** such as infrastructure security, data encryption, monitoring & logging, identity & access control, penetration testing, and DDoS attacks.
- It can efficiently manage and secure Windows workloads.

2. Microsoft Azure

[Microsoft Azure](#) is also known as **Windows Azure**. It supports various operating systems, databases, programming languages, frameworks that allow **IT** professionals to easily build, deploy, and manage applications through a worldwide network. It also allows users to create different groups for related utilities.



Features of Microsoft Azure

- Microsoft Azure provides **scalable, flexible, and cost-effective**
- It allows developers to quickly manage applications and websites.
- It managed each resource individually.
- Its IaaS infrastructure allows us to launch a general-purpose virtual machine in different platforms such as Windows and Linux.
- It offers a **Content Delivery System (CDS)** for delivering the Images, videos, audios, and applications.

3. Google Cloud Platform

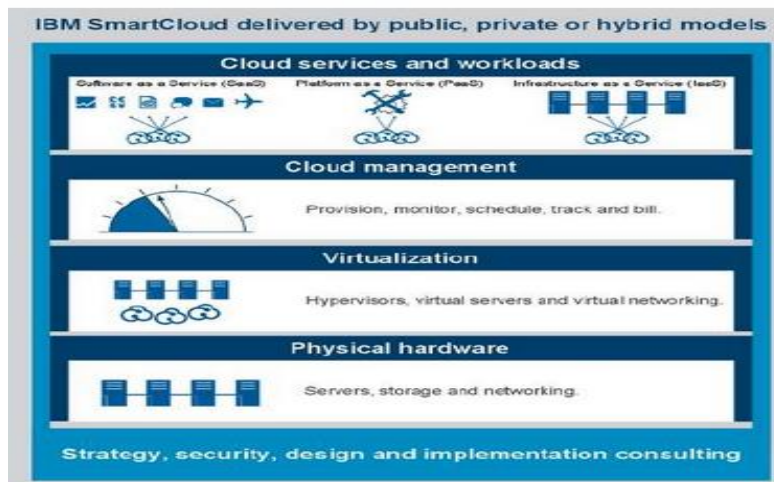
Google cloud platform is a product of **Google**. It consists of a set of physical devices, such as computers, hard disk drives, and virtual machines. It also helps organizations to simplify the migration process.

Features of Google Cloud

- Google cloud includes various **big data services** such as Google BigQuery, Google CloudDataproc, Google CloudDatalab, and Google Cloud Pub/Sub.
- It provides various services related to **networking**, including Google Virtual Private Cloud (VPC), Content Delivery Network, Google Cloud Load Balancing, Google Cloud Interconnect, and Google Cloud DNS.
- It offers various **scalable and high-performance**
- GCP provides various **serverless services** such as Messaging, Data Warehouse, Database, Compute, Storage, Data Processing, and Machine learning (ML)
- It provides a free cloud shell environment with Boost Mode.

4. IBM Cloud Services

IBM Cloud is an open-source, faster, and more reliable platform. It is built with a suite of advanced data and [AI](#) tools. It offers various services such as [Infrastructure as a service](#), [Software as a service](#), and [platform as a service](#). You can access its services like compute power, cloud data & Analytics, cloud use cases, and storage networking using internet connection.



Feature of IBM Cloud

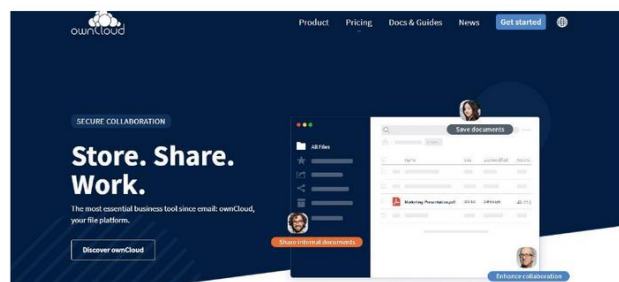
- IBM cloud improves operational efficiency.
- Its speed and agility improve the customer's satisfaction.
- It offers Infrastructure as a Service (IaaS), Platform as a Service (PaaS), as well as Software as a Service (SaaS)
- It offers various cloud communications services to our IT environment.

PRACTICAL-2

Aim: Setting up Virtualization using VirtualBox/VMWare Hypervisor

OwnCloud is an open source file sync and share software for everyone from individuals operating the free ownCloud Server edition, to large enterprises and service providers operating the ownCloud Enterprise Subscription. ownCloud provides a safe, secure, and compliant file synchronization and sharing solution on servers that you control.

You can share one or more files and folders on your computer, and synchronize them with your ownCloud server. Place files in your local shared directories, and those files are immediately synchronized to the server and to other devices using the ownCloud Desktop Sync Client, Android app, or iOS app. To learn more about the ownCloud desktop and mobile clients



Design

For desktop machines to synchronize files with their ownCloud server, desktop clients are available for PCs running Windows, macOS, FreeBSD or Linux. Mobile clients exist for iOS and Android devices. Files and other data (such as calendars, contacts or bookmarks) can also be accessed, managed, and uploaded using a web browser without any additional software. Any updates to the file system are pushed to all computers and mobile devices connected to a user's account. Encryption of files may be enforced by the server administrator.

The ownCloud server is written in the PHP and JavaScript scripting languages. For remote

access, it employs sabre/dav, an open-source WebDAV server. ownCloud is designed to work with several database management systems, including SQLite, MariaDB, MySQL, Oracle Database, and PostgreSQL.

Features

Owncloud is a software only product. owncloud does not offer off-premises storage capacity. This is in contrast to the likes of Dropbox that offers both software and off-premises storage capacity. owncloud storage capacity has to be provided on user owned devices.

OwnCloud files are stored in conventional directory structures, and can be accessed via WebDAV if necessary. User files are encrypted both at rest and during transit. ownCloud can synchronise with local clients running Windows (Windows XP, Vista, 7 and 8), macOS (10.6 or later), or various Linux distributions.

ownCloud users can manage calendars (CalDAV), contacts (CardDAV) scheduled tasks and streaming media (Ampache) from within the platform.

Enterprise Features

For Enterprise customers, ownCloud GmbH offers apps with additional functionality. They are mainly useful for large organizations with more than 500 users. An Enterprise subscription includes support services.

Commercial features include End-to-end encryption, Ransomware and Antivirus protection, Branding, Document Classification, Single-Sign-On via Shibboleth/SAML

PRACTICAL-3

AIM:- Setting up Virtualization using Virtual Box/VMware

Hypervisor.

Run VirtualBox and create a new virtual machine. Name it ESXi , select Linux as its type, and Other Linux (64-bit) as the version

Installing VirtualBox

To start with, we will download VirtualBox and install it. We should follow the steps given below for the installation.

Step 1 – To download VirtualBox, click on the following link

<https://www.virtualbox.org/wiki/Downloads> Now, depending on your OS, select which version to install. In our case, it will be the first one (Windows host).

VirtualBox

Download VirtualBox

Here, you will find links to VirtualBox binaries and its source code.

VirtualBox binaries


By downloading, you agree to the terms and conditions of the respective license.

- **VirtualBox platform packages.** The binaries are released under the terms of the GPL version 2.
 - **VirtualBox 5.1.2 for Windows hosts** ⇨ x86/amd64
 - **VirtualBox 5.1.2 for OS X hosts** ⇨ amd64
 - **VirtualBox 5.1.2 for Linux hosts**
 - **VirtualBox 5.1.2 for Solaris hosts** ⇨ amd64
- **VirtualBox 5.1.2 Oracle VM VirtualBox Extension Pack** ⇨ [All supported platforms](#)
Support for USB 2.0 and USB 3.0 devices, VirtualBox RDP and PXE boot for Intel cards. See [this chapter from the User Manual](#) for an introduction. Extension Pack binaries are released under the [VirtualBox Personal Use and Evaluation License \(PUEL\)](#).
Please install the extension pack with the same version as your installed version of VirtualBox:
If you are using **VirtualBox 5.0.26**, please download the extension pack ⇨ [here](#).
If you are using **VirtualBox 4.3.38**, please download the extension pack ⇨ [here](#).

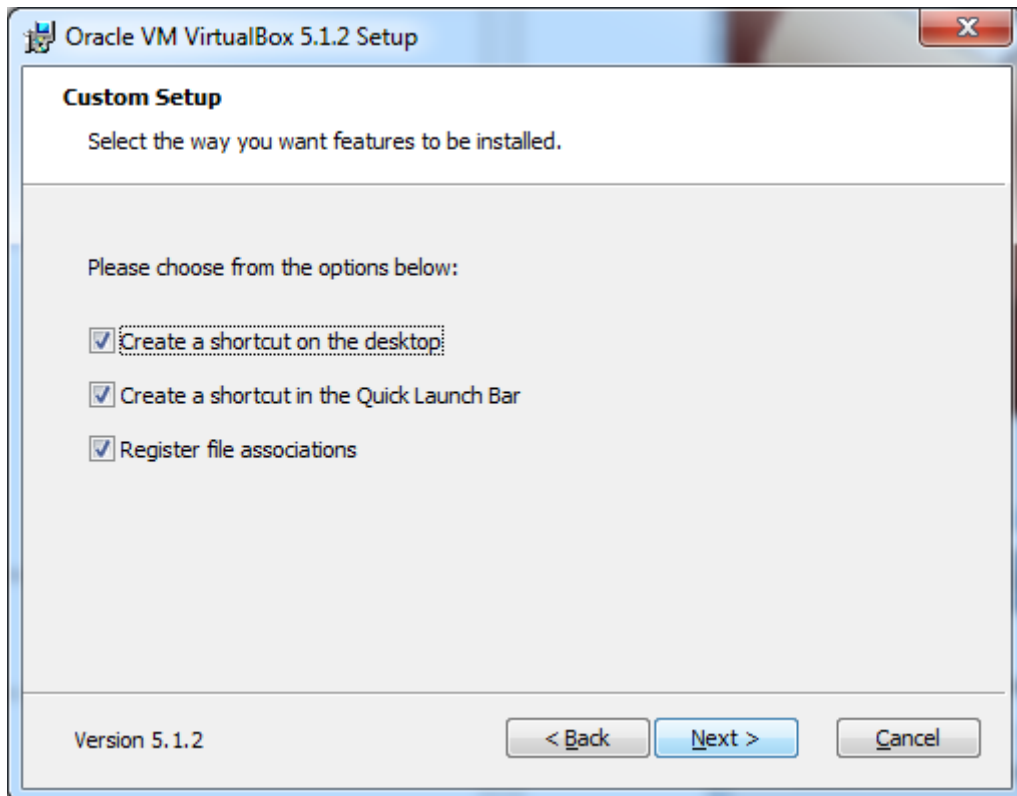
Step 2 – Once the option is selected, click on “Next”.



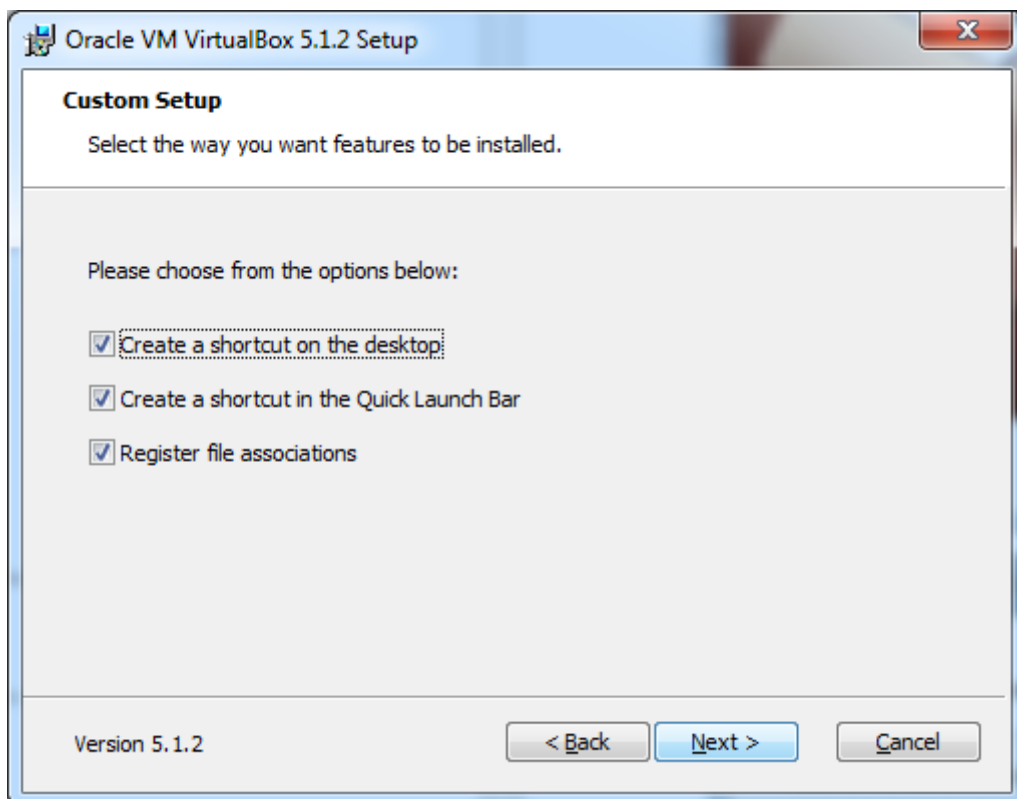
Step 3 – You have the option asking where to install the application. We can leave it as default and click on “Next”.



Step 4 – Once the options are selected as shown in the following screenshot, click on Next.




ext.



Step 5 – A dialog box will come up asking whether to proceed with the installation. Click “Yes”.



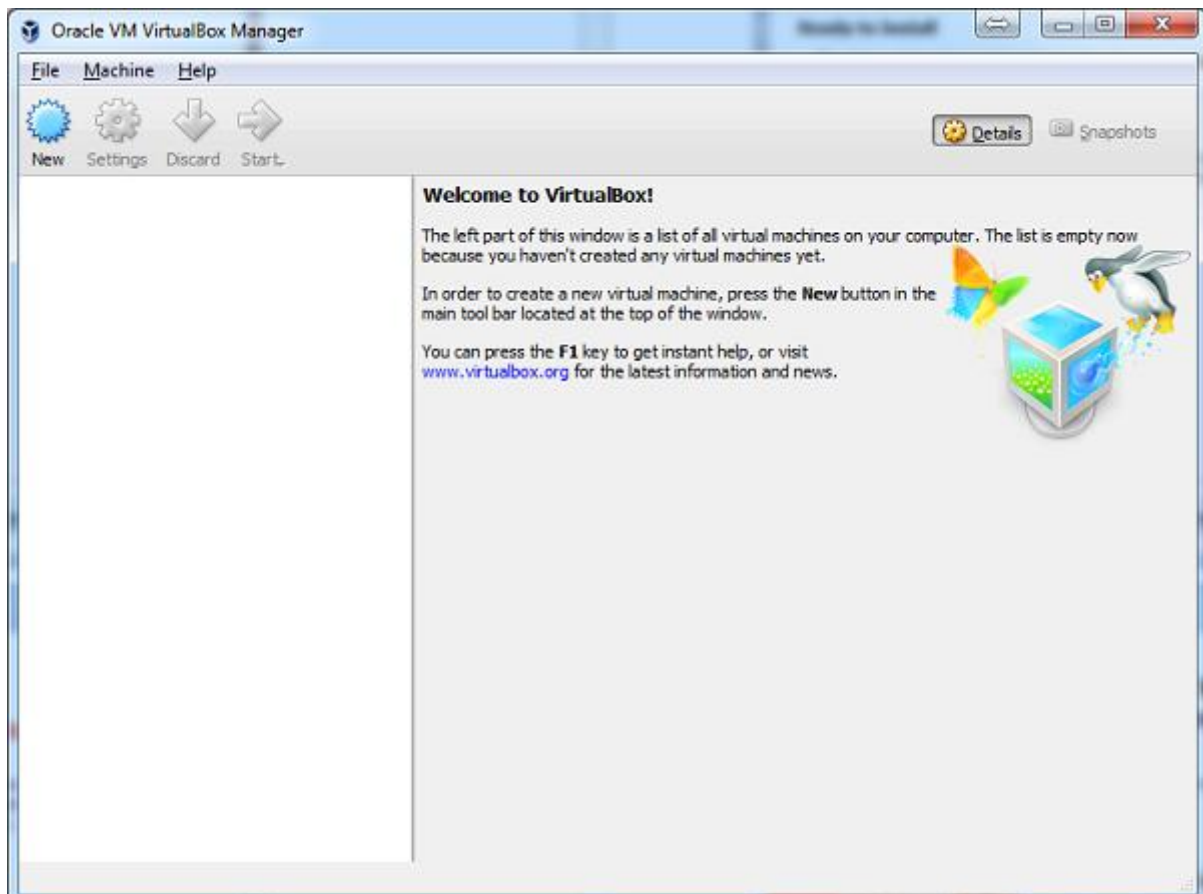
Step 6 – In the next step, click on “Install”.



Step 7 – Tick the start VirtualBox check box and click on “Finish”.




Step 8 – VirtualBox application will now open as shown in the following screenshot. Now, we are ready to install the virtual machines.



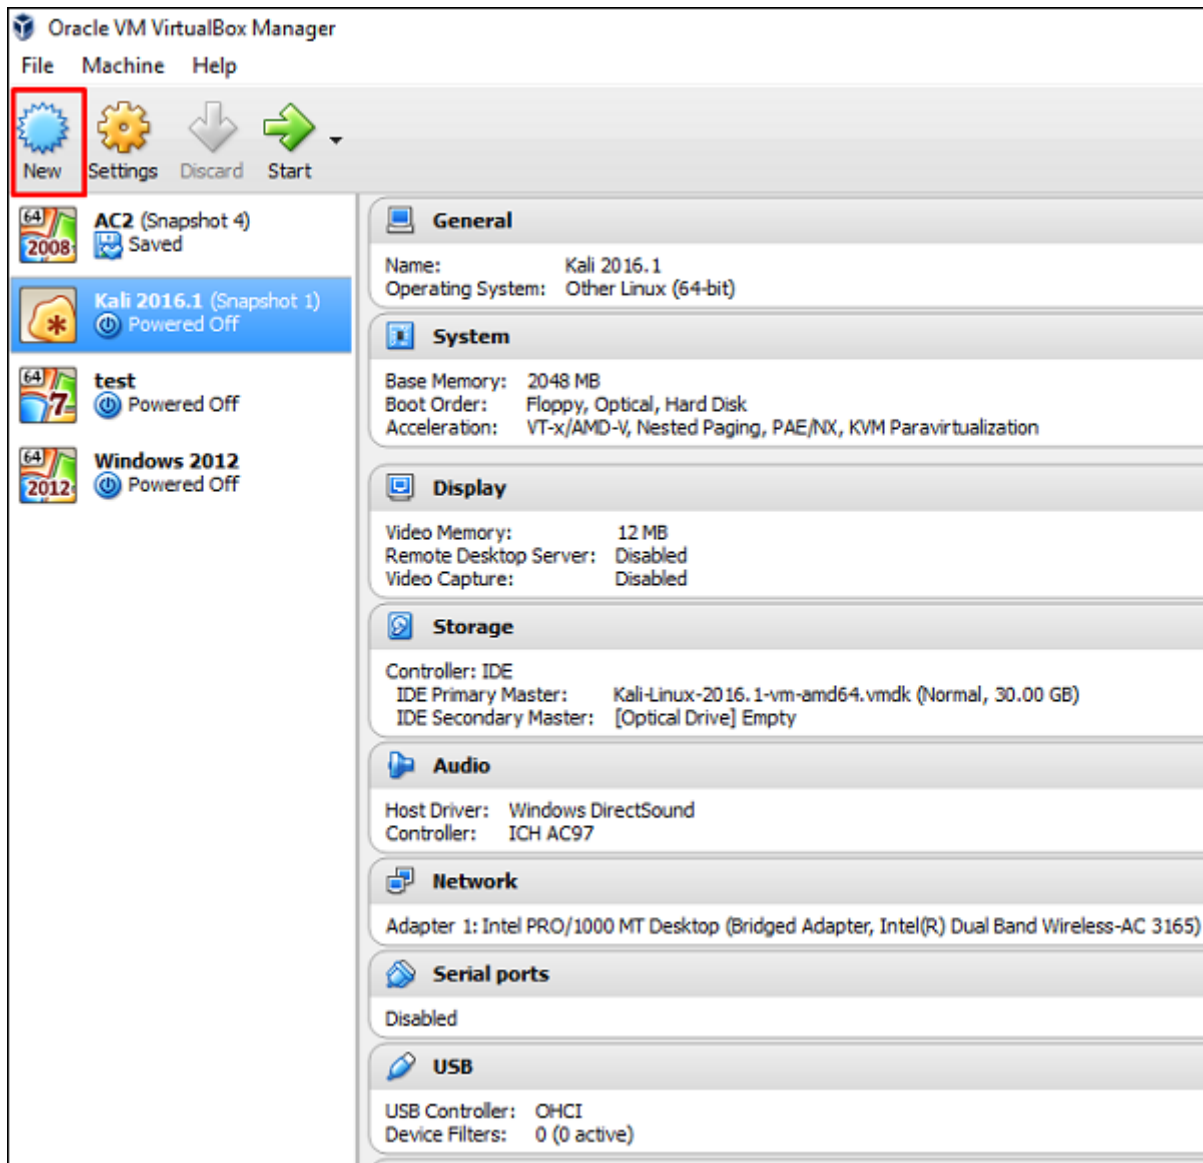
Creating a VM with VirtualBox

To create a virtual machine with Oracle VirtualBox, we should follow the steps given below.

Step 1 – To begin with, click on the “Oracle VM VirtualBox” icon on the desktop as shown in the screenshot below.



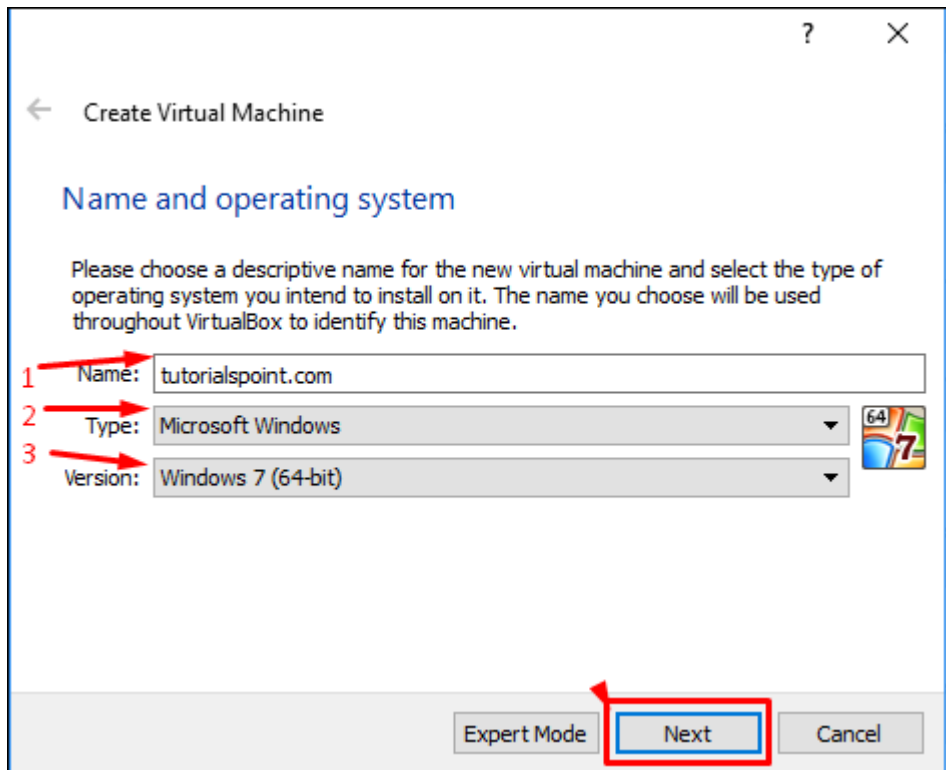
Step 2 – The next step is to click on “New” button, which is in the top left hand side of the screen.



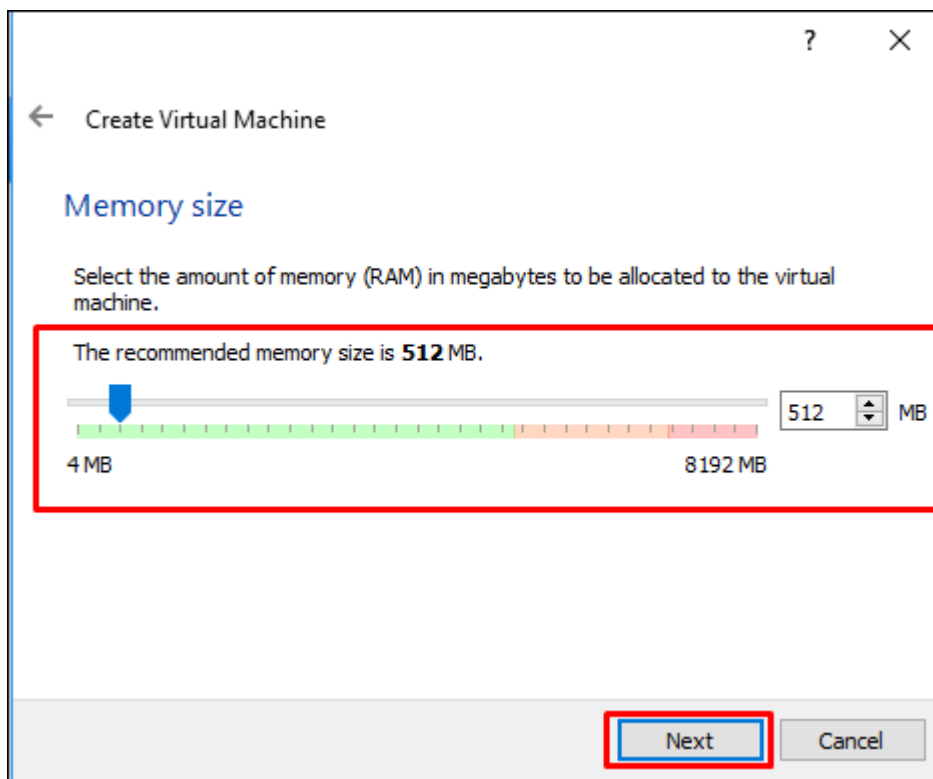
Step 3 – A table will pop-up requesting you the parameters for the virtual machine. These will be –

- **Name** – We have to put a friendly name for this Virtual Machine.
- **Type** – Enter the OS that is going to be installed on it.
- **Version** – Enter the specific version for that OS, which we have selected earlier.

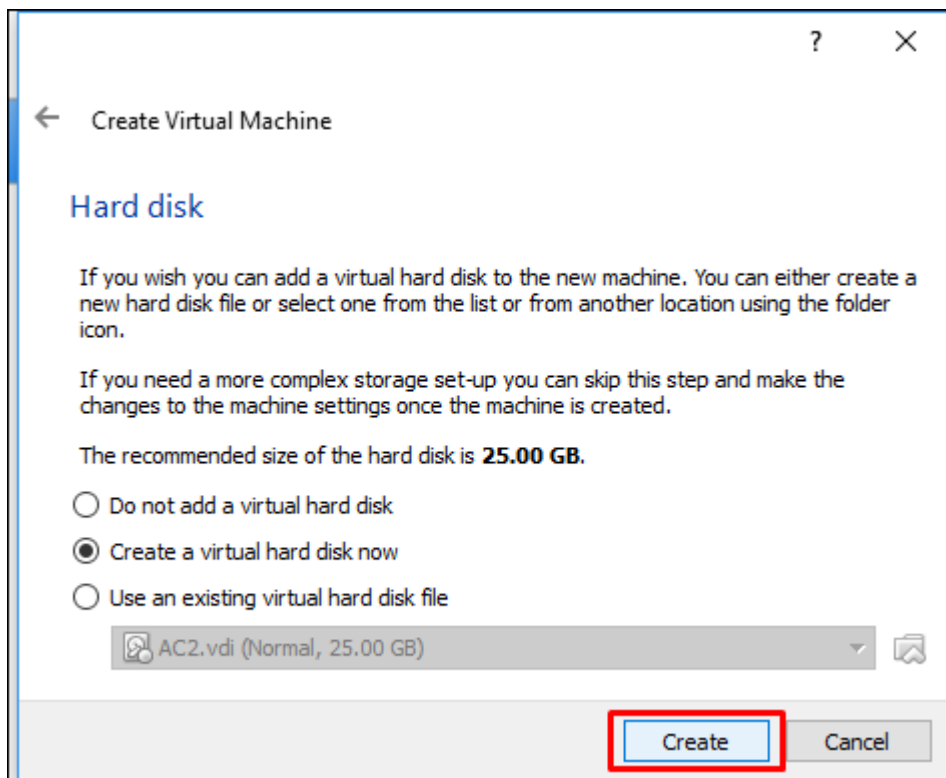
Once all the above parameters are filled, click on “Next”.



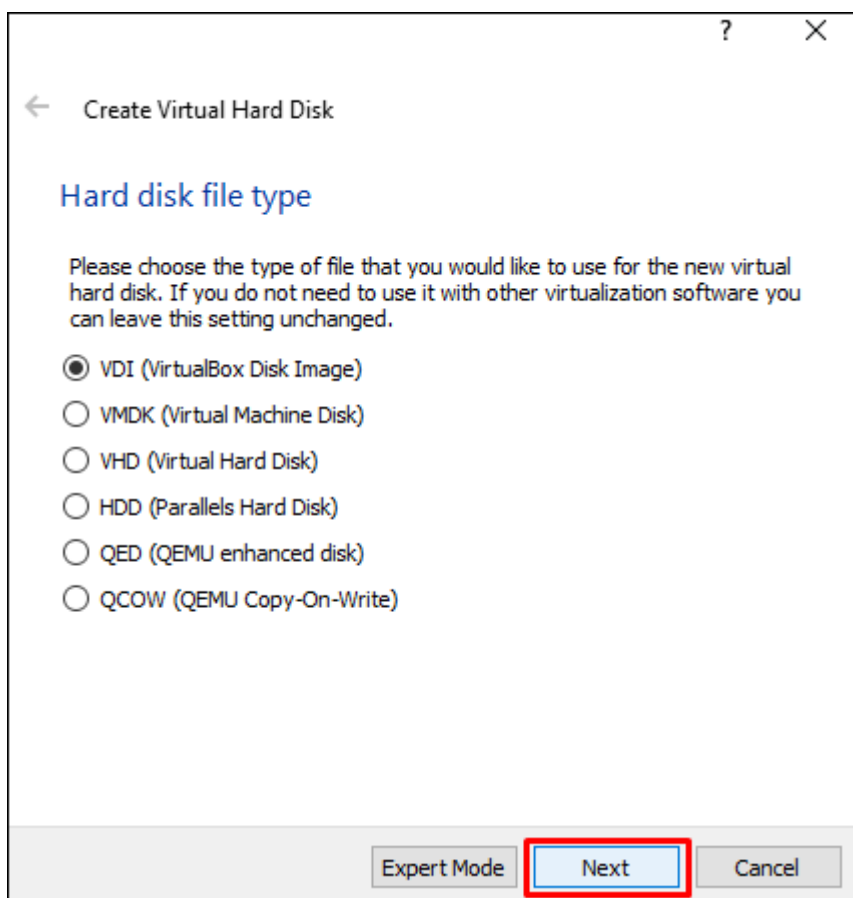
Step 4 – Select the amount of memory that you need to allocate in this VM → Click on “Next”.



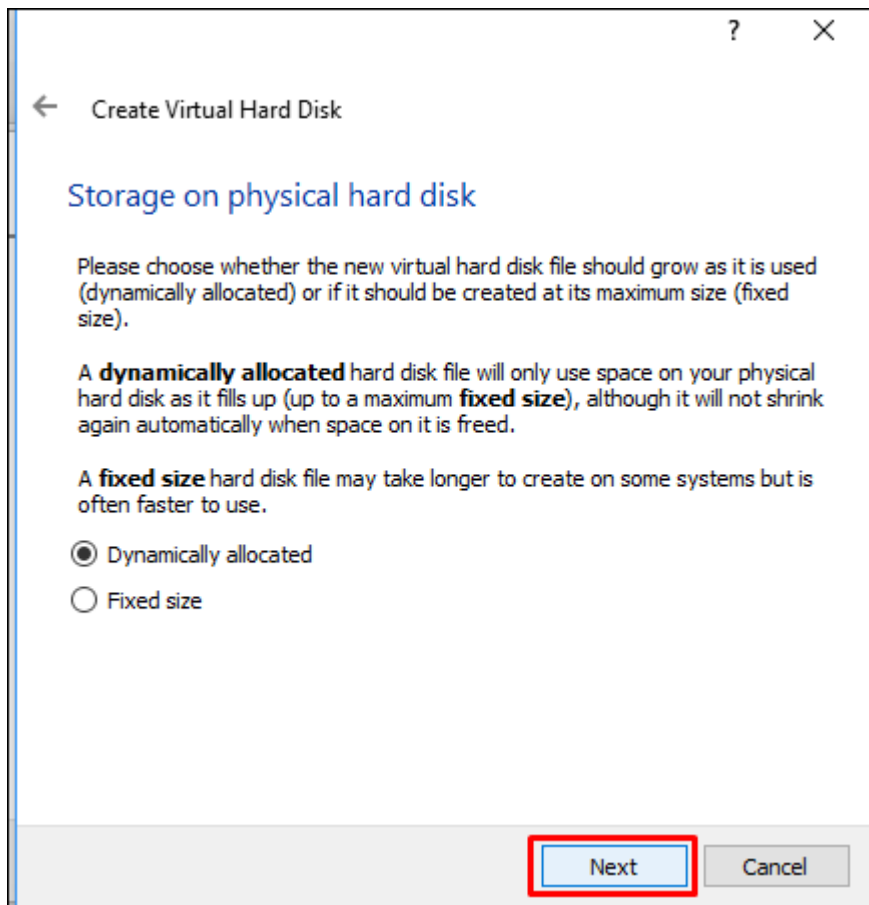
Step 5 – Check one of the three options for the HDD and click on “Create”.



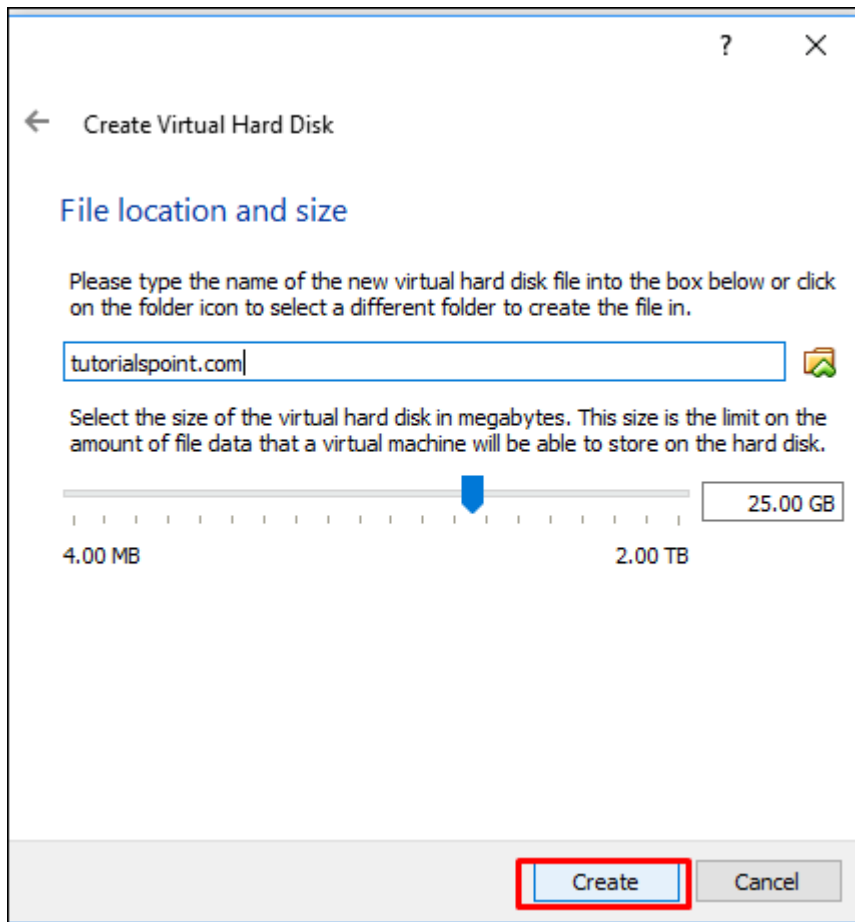
Step 6 – Select a file extension for your virtual HDD (It is recommended to use a common file extension that most of the hypervisors use like VHD) → click on “Next”.



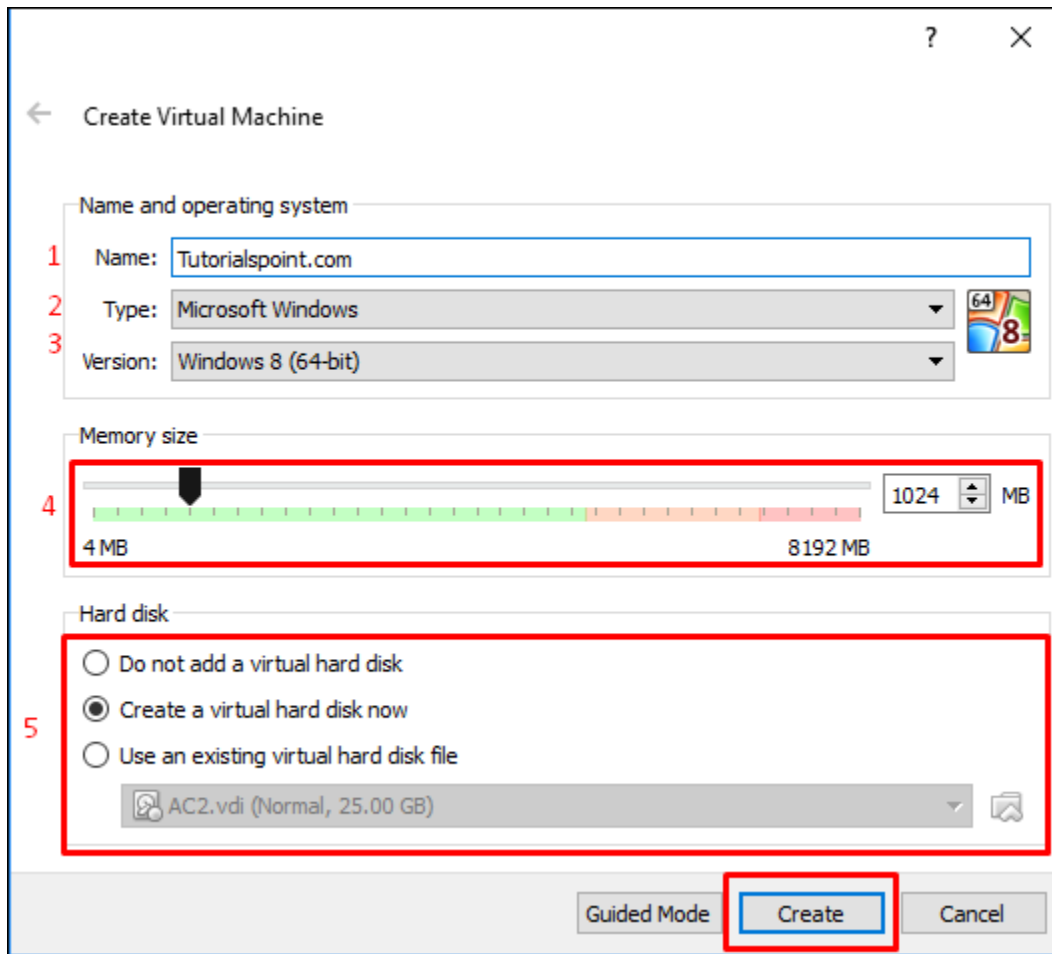
Step 7 – Choose whether you want the Virtual HDD as dynamic or fixed. This is based on your needs → Click on “Next”.



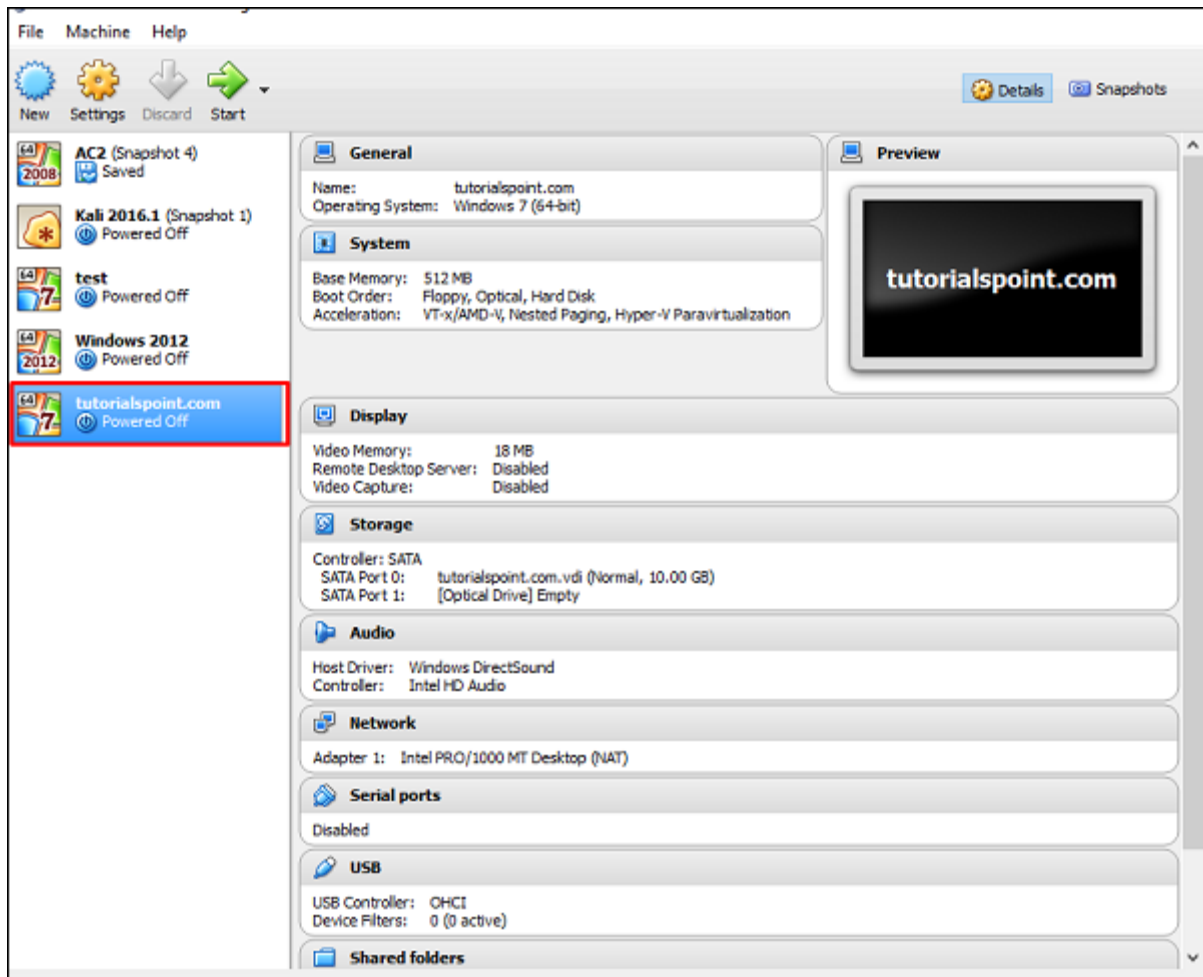
Step 8 – Put a name for your virtual HDD file and select the disk size for your VM → Click on “Create”.



All the above steps can be done in one shot by selecting the “Expert mode”.



The virtual machine created will be as shown in the screenshot below.



Setting up Networking with VirtualBox

There are two types of networking modes in VirtualBox, which are –

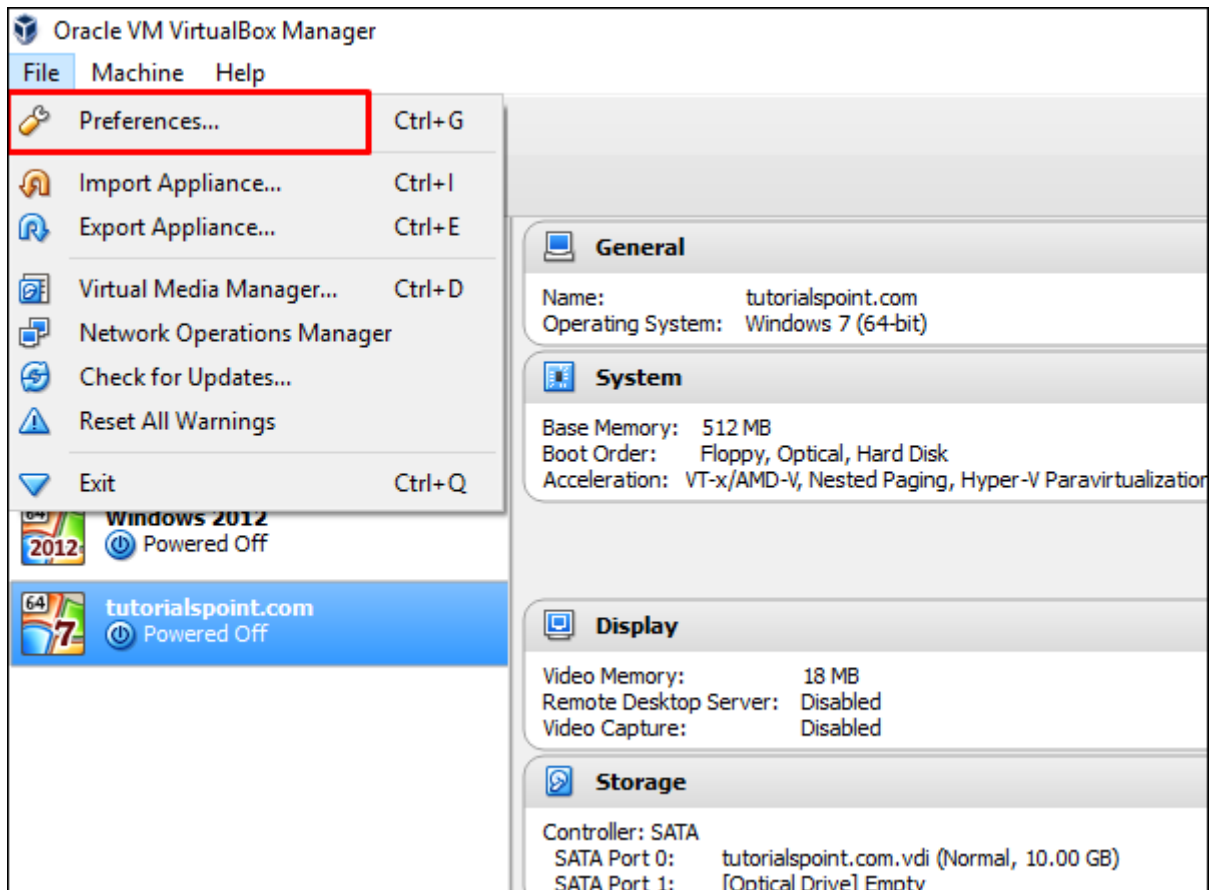
- Nat Networks and
- Host-only Networks.

Both of these are explained in detail below.

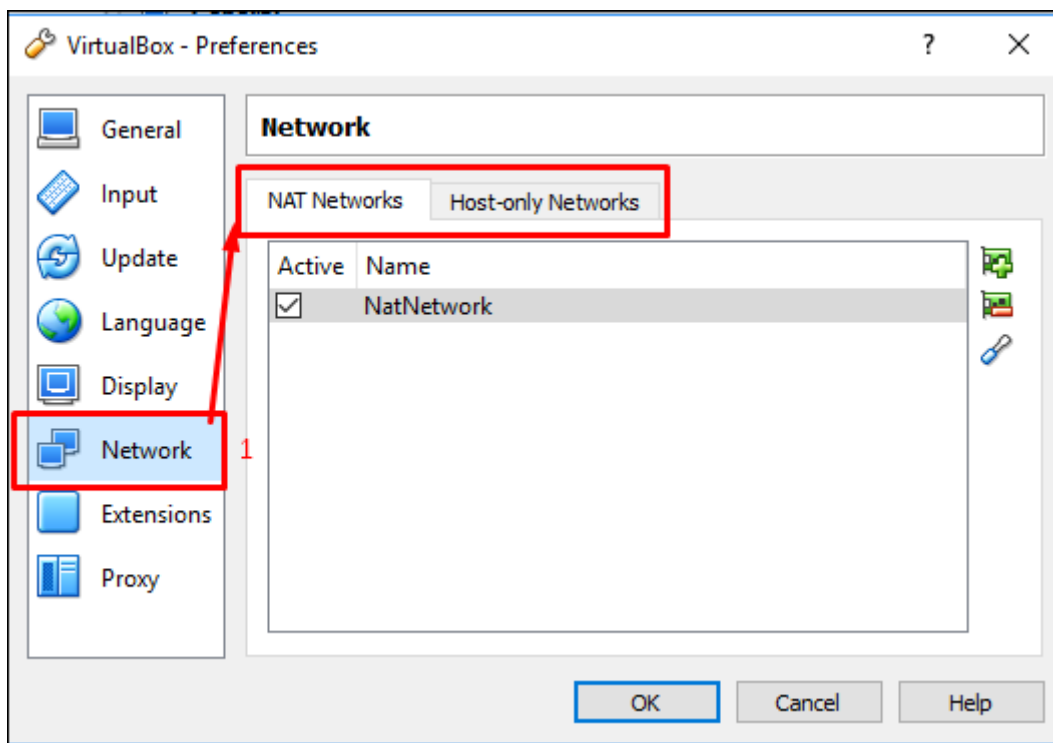
Nat Networks

For setting up Nat Networks, we should follow the steps given below.

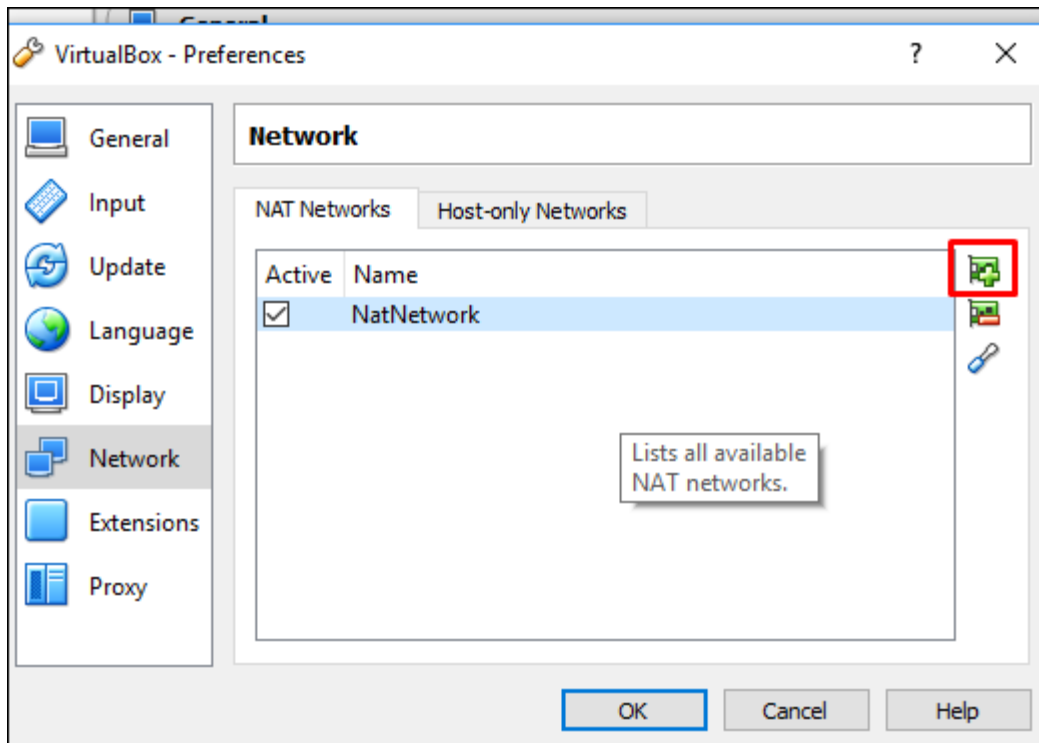
Step 1 – Go to Oracle VM VirtualBox Manager → Click on “Preferences...”



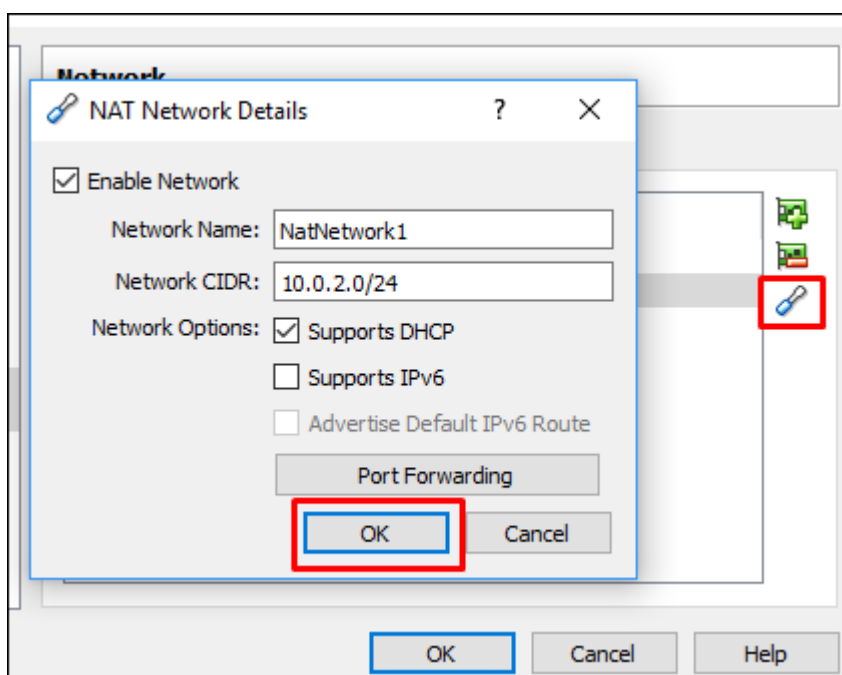
Step 2 – Click on “Network” and then on the left panel click on the “NAT Networks” tab.



Step 3 – Click on the “+” button, which is highlighted in the screenshot below.



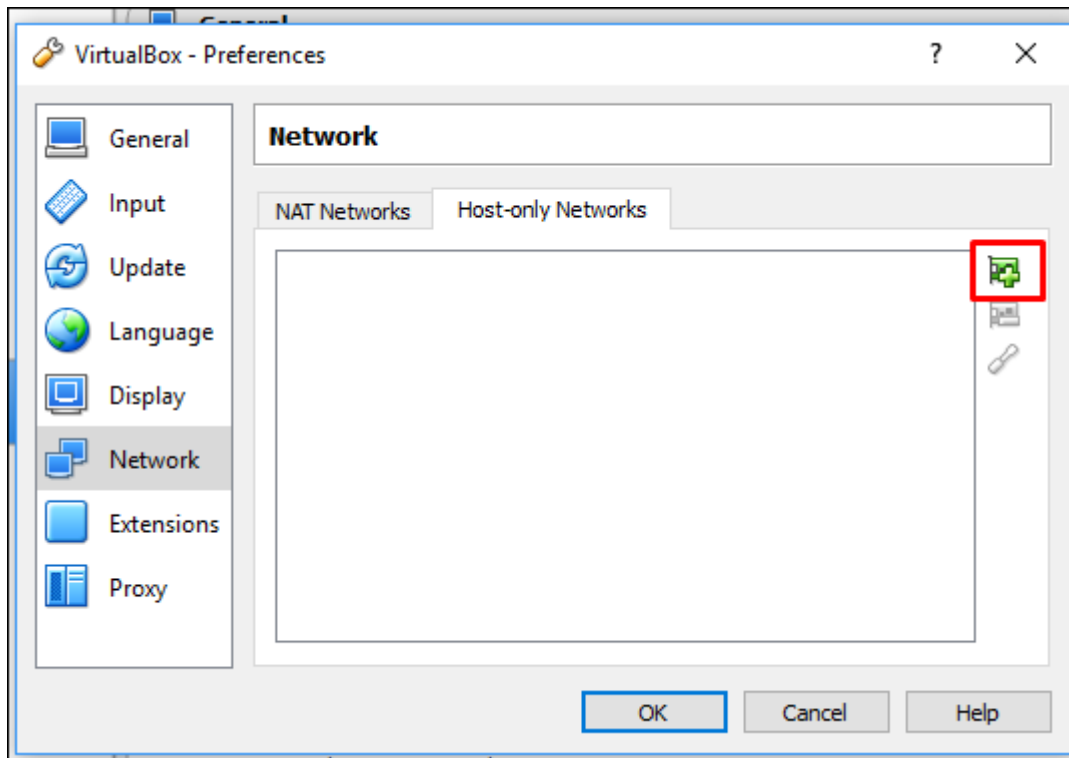
Step 4 – Here, we have to put the “Network Name” and the IP range for this network that will be NAT-ed, in order to have access to internet and to other networks.



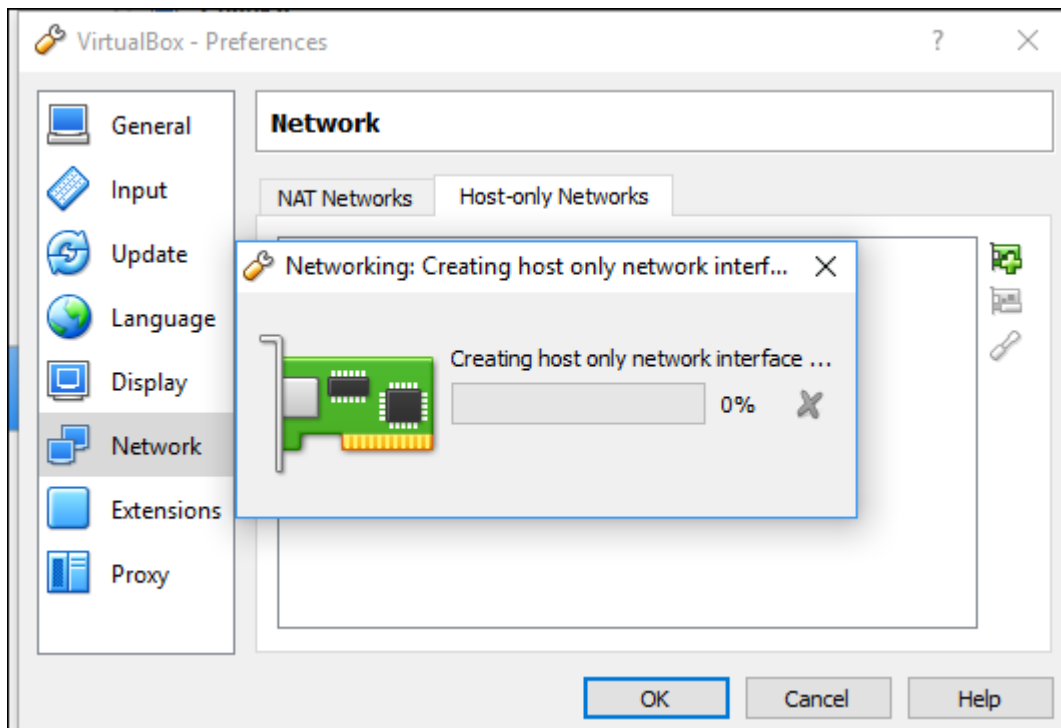
Host-only Networks

For setting up Host-only Networks, we should follow the steps given below.


Step 1 – If you click on the “Host-only Networks” tab, you can create networks that are isolated from the other networks. However, VM hosts communicate with each other and the Hypervisor machine. Click on the “+” sign.



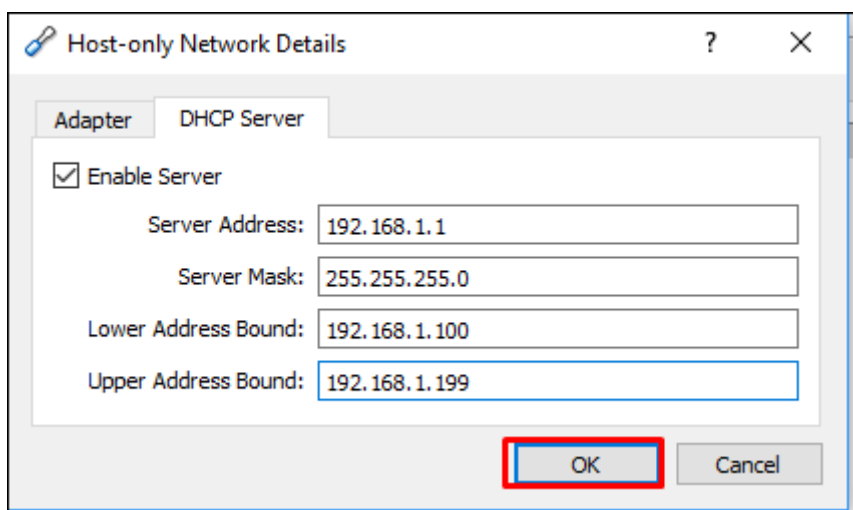
Step 2 – The host interface will continue to be created as shown in the screenshot below.



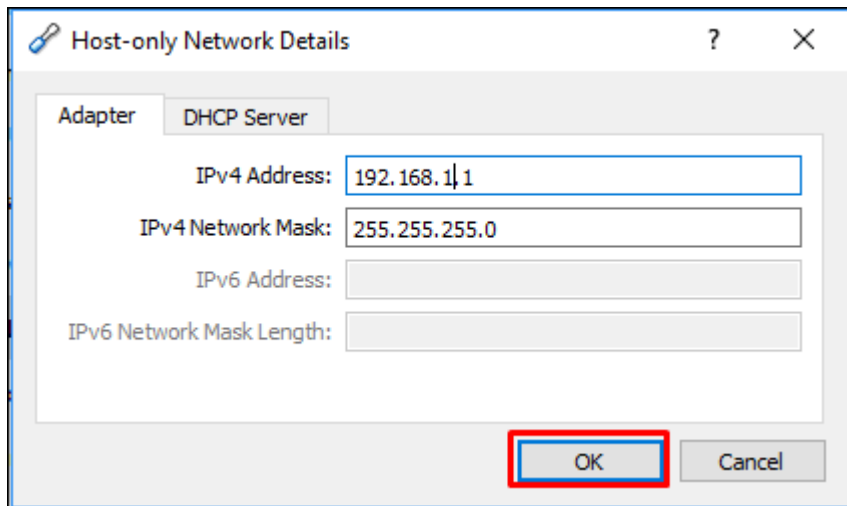
Step 3 – If you click on  button, you can edit the settings.



Step 4 – If you want your host machines to take “DHCP IP”, click on the “DHCP Server” tab and check the box “Enable Server” → Click “OK”.

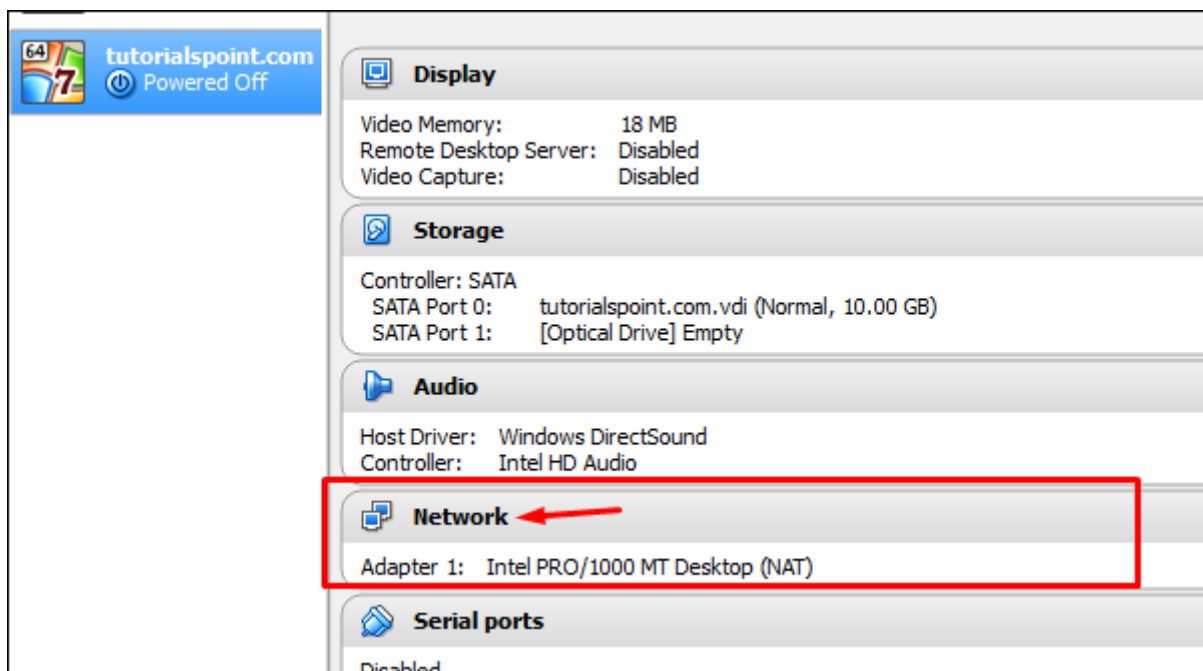


Step 5 – In the “Adapter” tab, put the IP of the hypervisor.



After all these preparations for setting up the network modes is complete. It is now time to assign a network to our VMs.

To do this, Click on the VMs on the left side of the panel, then right click on the “Network” option and a table will be open.

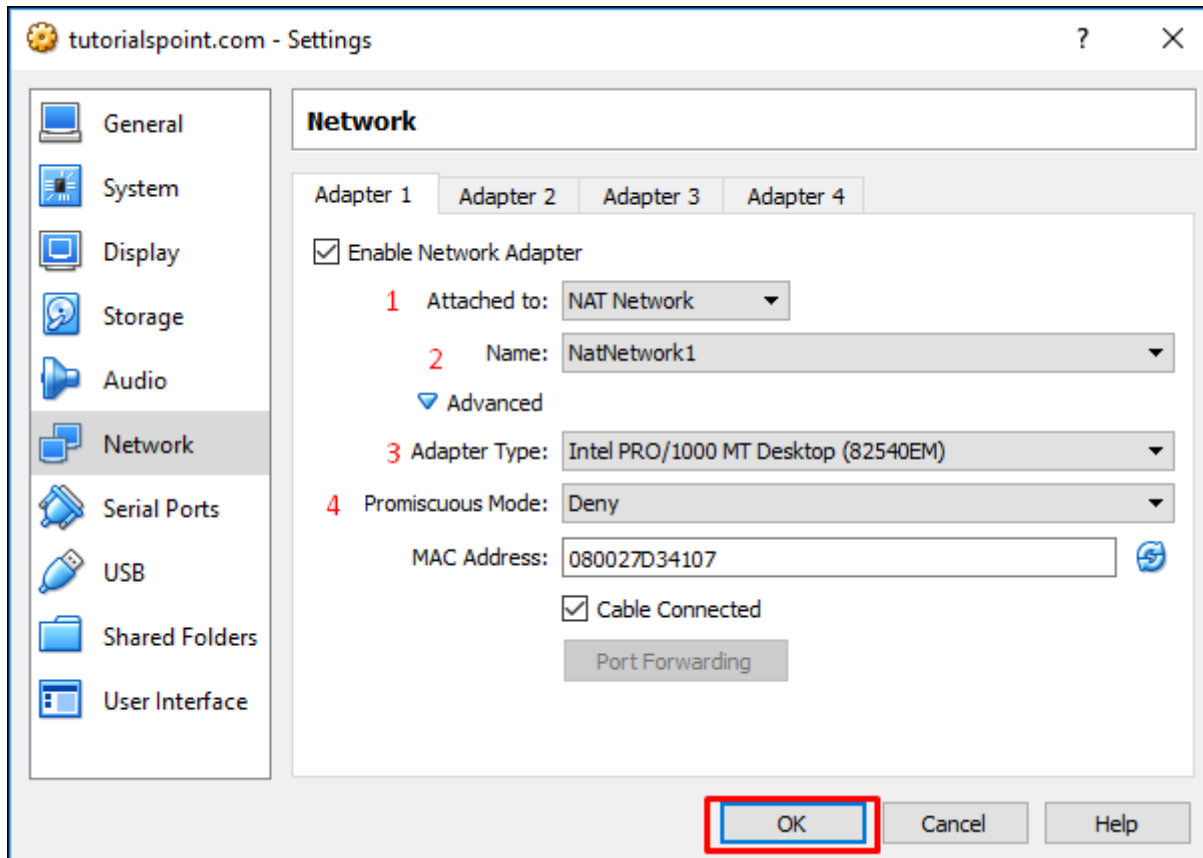


You can have up to four Ethernet adaptors per machine. The following image has four sections highlighted, which are explained below.

- Check the box “Enable Network Adapter” to enable the vNIC on the VM and attach it to one network.
- You can have many networks created, so we have to select one of them in the “Name” dropdown box.
- In the adapter type dropdown-box, we have to select a physical NIC that the hypervisor has.

- Promiscuous Mode: Here, we can select “Deny”, if we do not want the VMs to communicate with each other.

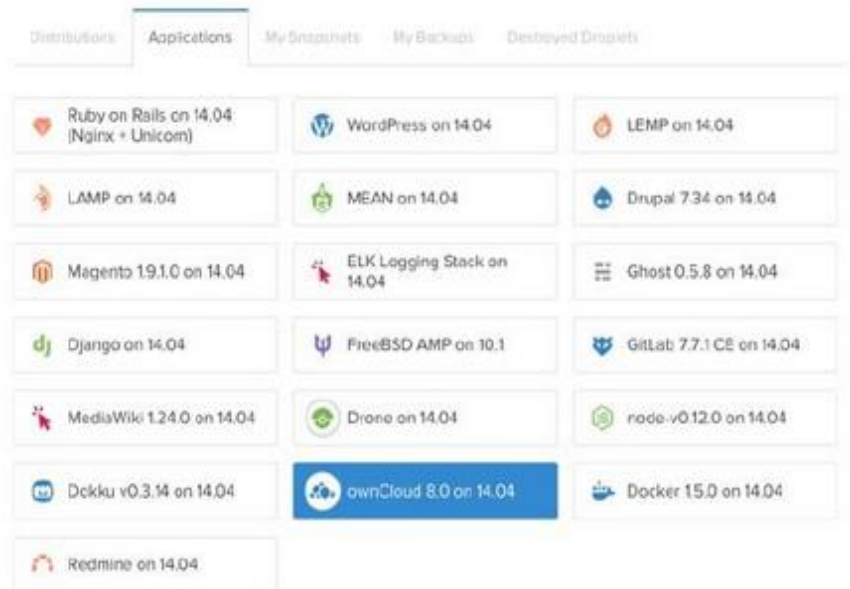
Once all the above parameters are completed. Click on “OK”.



PRACTICAL-4

AIM:- Installation and configuration of OwnCloud software for SaaS.

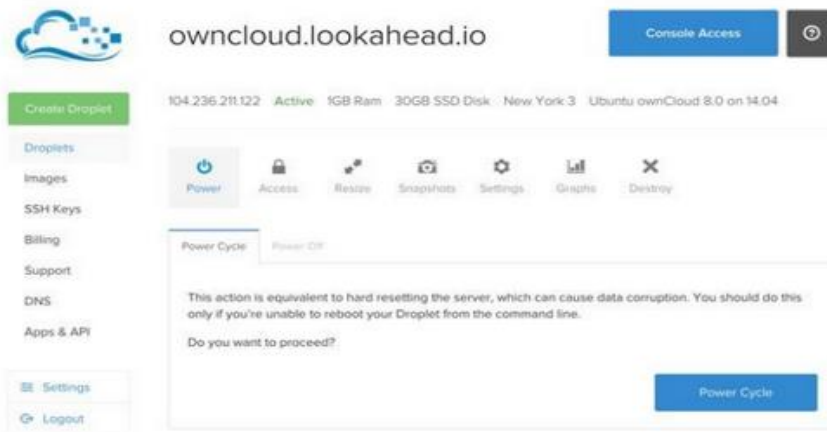
Click on the Applications tab and click the ownCloud 8.x image button



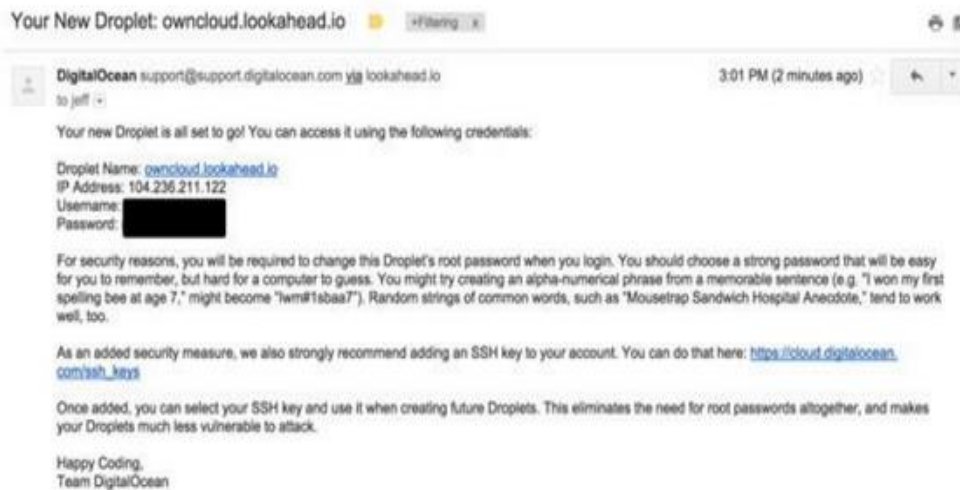

Submit the form and create droplet progress screen:



When it's complete, see droplet on the dashboard:



Now, configure domain name DNS to point to the droplet, i.e. set an A record for the subdomain or domain name to the IP address



Once your DNS changes propagate visit your site—but you'll get this warning. OwnCloud requires you to add this domain to its listed of trusted domains.



When you click the grey button, you should see the general login screen:

Log In to Your OwnCloud Dashboard

However, you may have noticed we don't have our dashboard credentials yet:



To obtain them, use the SSH login credentials and a terminal application to log in to your droplet server

```
Admins-MBP:~ Jeff$ ssh root@owncloud.lookahead.io
The authenticity of host 'owncloud.lookahead.io (104.236.211.122)' can't be established.
RSA key fingerprint is [REDACTED]
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'owncloud.lookahead.io,104.236.211.122' (RSA) to the list of known hosts.
root@owncloud.lookahead.io's password:
You are required to change your password immediately (root enforced)
-----
Thank you for using DigitalOcean's OwnCloud Application.

Your OwnCloud instance can be accessed at http://104.236.211.122/
Your OwnCloud login credentials are:
Username: admin
Password: [REDACTED]
-----
To delete this message of the day: rm -rf /etc/motd.tail
Changing password for root.
(current) UNIX password:
Enter new UNIX password:
Retype new UNIX password:
root@owncloud:~#
```

Then, return to the browser and log in to your OwnCloud.

104.236.211.122/index.php/apps/files/

Files admin

All files New

Size Modified

32 minutes ago
32 minutes ago
32 minutes ago

Welcome to ownCloud

Your personal web services. All your files, contacts, calendar and more, in one place.

Get the apps to sync your files

Desktop app
Available on Google play
Available on the App Store

Connect your desktop apps to ownCloud


Connect your Calendar Connect your Contacts Access files via WebDAV

There's more information in the [documentation](#) and on our [website](#).
If you like ownCloud, **recommend it to your friends** and **spread the word!**

PRACTICAL-5

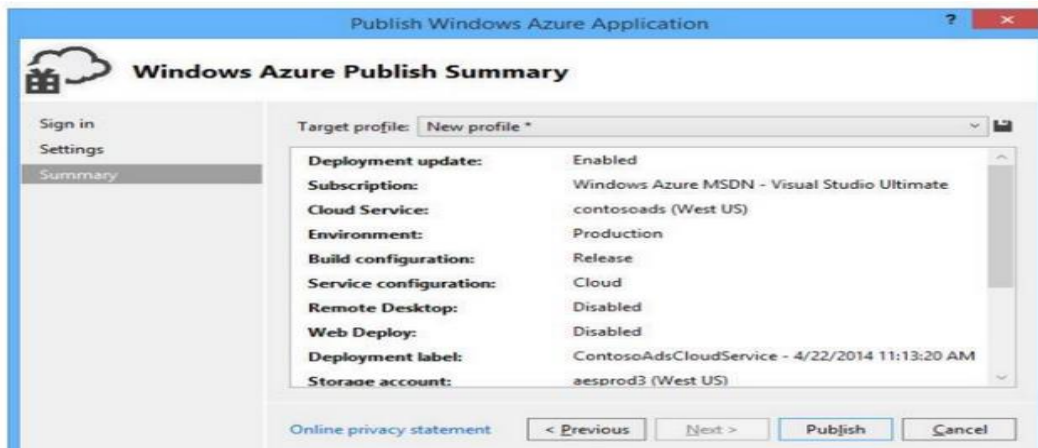
AIM:- Accessing Microsoft AZURE cloud-services.

In Solution Explorer, right-click the ContosoAdsCloudService cloud project and then select Publish.




In the Sign in step of the Publish Azure Application wizard, click Next





When the deployment status is complete, click the Web app URL to start the application.



PRACTICAL-6

AIM:- Cloud Simulation Software Introduction: CloudSim.

CloudSim is a simulation toolkit that supports the modeling and simulation of the core functionality of cloud, like job/task queue, processing of events, creation of cloud entities(datacenter, datacenter brokers, etc), communication between different entities, implementation of broker policies, etc. This toolkit allows to:

- Test application services in a repeatable and controllable environment.
- Tune the system bottlenecks before deploying apps in an actual cloud.
- Experiment with different workload mix and resource performance scenarios on simulated infrastructure for developing and testing adaptive application provisioning techniques

Core features of CloudSim are:

- The Support of modeling and simulation of large scale computing environment as federated cloud data centers, virtualized server hosts, with customizable policies for provisioning host resources to virtual machines and energy-aware computational resources
- It is a self-contained platform for modeling cloud's service brokers, provisioning, and allocation policies.
- It supports the simulation of network connections among simulated system elements.
- Support for simulation of federated cloud environment, that inter-networks resources from both private and public domains.
- Availability of a virtualization engine that aids in the creation and management of multiple independent and co-hosted virtual services on a data center node.
- Flexibility to switch between space shared and time shared allocation of processing cores to virtualized services.

PRACTICAL-7

Aim : Create a cluster with AWS

Step 1: Create a cluster

If you already have a cluster that you want to use, you can skip this step.

For the exercises in this tutorial, use a four-node cluster.

To create a cluster

1. Sign in to the AWS Management Console and open the Amazon Redshift console at <https://console.aws.amazon.com/redshift/>

Using the navigation menu, choose the **Provisioned clusters dashboard**.

- *Important*

Make sure that you have the necessary permissions to perform the cluster operations. For information on granting the necessary permissions, see [Authorizing Amazon Redshift to access AWS services](#).

- At top right, choose the AWS Region in which you want to create the cluster. For the purposes of this tutorial, choose **US West (Oregon)**.
- On the navigation menu, choose **Clusters**, then choose **Create cluster**. The **Create cluster** page appears.
- On the **Create cluster** page enter parameters for your cluster. Choose your own values for the parameters, except change the following values:
 - Choose `dc2.large` for the node type.
 - Choose `4` for the **Number of nodes**.
 - In the **Cluster permissions** section, choose an IAM role from **Available IAM roles**. This role should be one that you previously created and that has access to Amazon S3. Then choose **Associate IAM role** to add it to the list of **Associated IAM roles** for the cluster.
- Choose **Create cluster**.

Step 2: Download the data files

In this step, you download a set of sample data files to your computer. In the next step, you upload the files to an Amazon S3 bucket.

To download the data files

1. Download the zipped file: [LoadingDataSampleFiles.zip](#).
2. Extract the files to a folder on your computer.
3. Verify that your folder contains the following files.

```
customer-fw-manifest
customer-fw.tbl-000
customer-fw.tbl-000.bak
customer-fw.tbl-001
customer-fw.tbl-002
customer-fw.tbl-003
customer-fw.tbl-004
customer-fw.tbl-005
customer-fw.tbl-006
customer-fw.tbl-007
customer-fw.tbl.log
dwdate-tab.tbl-000
dwdate-tab.tbl-001
dwdate-tab.tbl-002
dwdate-tab.tbl-003
dwdate-tab.tbl-004
dwdate-tab.tbl-005
dwdate-tab.tbl-006
dwdate-tab.tbl-007
part-csv.tbl-000
part-csv.tbl-001
part-csv.tbl-002
part-csv.tbl-003
part-csv.tbl-004
part-csv.tbl-005
part-csv.tbl-006
part-csv.tbl-007
```

Step 3: Upload the files to an Amazon S3 bucket

In this step, you create an Amazon S3 bucket and upload the data files to the bucket.

To upload the files to an Amazon S3 bucket

1. Create a bucket in Amazon S3.
 1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
1. Click **Create Bucket**.
2. In the **Bucket Name** box of the **Create a Bucket** dialog box, type a bucket name.

The bucket name you choose must be unique among all existing bucket names in Amazon S3. One way to help ensure uniqueness is to prefix your bucket names with the name of your organization. Bucket names must comply with certain rules. For more information, go to [Bucket restrictions and limitations](#) in the *Amazon Simple Storage Service User Guide*.

3. Select a Region.

Create the bucket in the same Region as your cluster. If your cluster is in the US West (Oregon) Region, choose **US West (Oregon) Region (us-west-2)**.

4. Choose **Create**.

When Amazon S3 successfully creates your bucket, the console displays your empty bucket in the **Buckets** panel.

- Create a folder.

1. Choose the name of the new bucket.
2. Choose the **Actions** button, and click **Create Folder** in the drop-down list.
3. Name the new folder `load`.
4. Upload the data files to the new Amazon S3 bucket.
 1. Choose the name of the data folder.
 2. In the Upload - Select Files wizard, choose **Add Files**.

Follow the Amazon S3 console instructions to upload all of the files you downloaded and extracted,

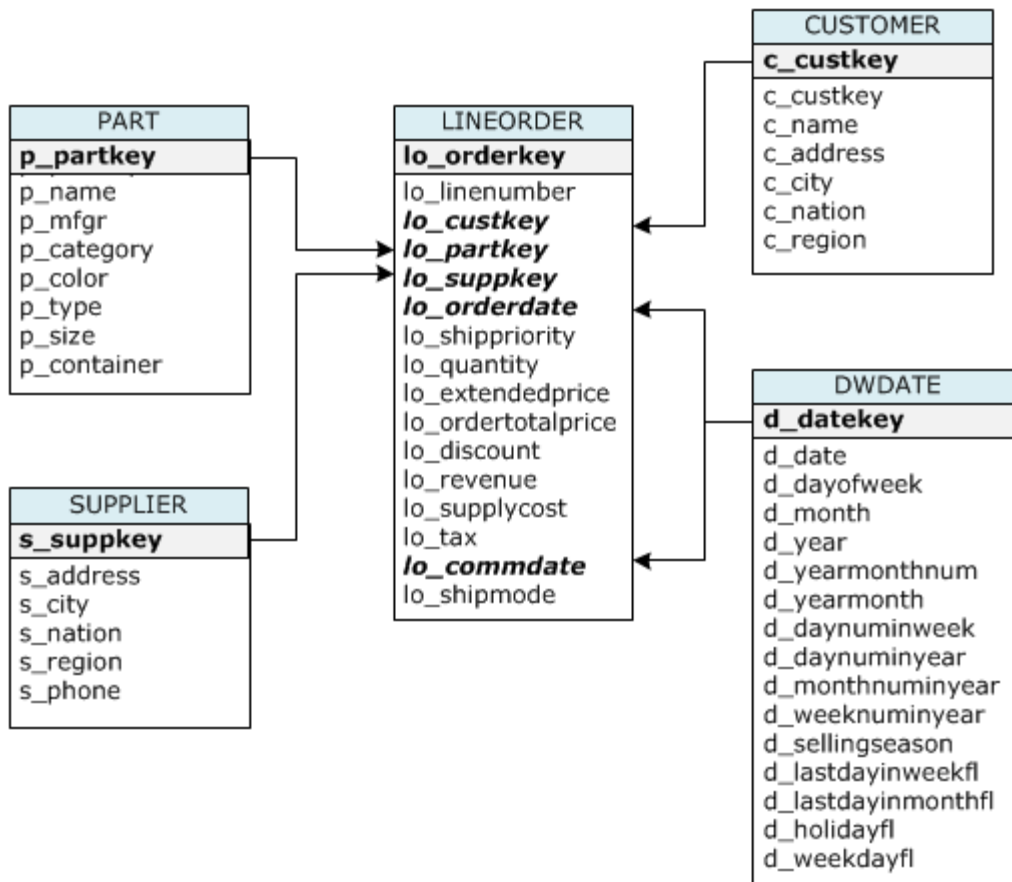
3. Choose **Start Upload**.

User Credentials

The Amazon Redshift COPY command must have access to read the file objects in the Amazon S3 bucket. If you use the same user credentials to create the Amazon S3 bucket and to run the Amazon Redshift COPY command, the COPY command has all necessary permissions. If you want to use different user credentials, you can grant access by using the Amazon S3 access controls. The Amazon Redshift COPY command requires at least ListBucket and GetObject permissions to access the file objects in the Amazon S3 bucket. For more information about controlling access to Amazon S3 resources, go to [Managing access permissions to your Amazon S3 resources](#).

Step 4: Create the sample tables

For this tutorial, you use a set of five tables based on the Star Schema Benchmark (SSB) schema. The following diagram shows the SSB data model.



The SSB tables might already exist in the current database. If so, drop the tables to remove them from the database before you create them using the CREATE TABLE commands in the next step. The tables used in this tutorial might have different attributes than the existing tables.

To create the sample tables

1. To drop the SSB tables, run the following commands in your SQL client.

```

• drop table part cascade;
drop table supplier;
drop table customer;
drop table dwdate;
drop table lineorder;
  
```

• Run the following CREATE TABLE commands in your SQL client.

```

CREATE TABLE part
(
  p_partkey      INTEGER NOT NULL,
  p_name         VARCHAR(22) NOT NULL,
  p_mfgr        VARCHAR(6),
  p_category     VARCHAR(7) NOT NULL,
  p_brand1      VARCHAR(9) NOT NULL,
  p_color       VARCHAR(11) NOT NULL,
  p_type        VARCHAR(25) NOT NULL,
  p_size        INTEGER NOT NULL,
  p_container    VARCHAR(10) NOT NULL
  
```

```
);
```

```
CREATE TABLE supplier
```

```
(  
  s_suppkey    INTEGER NOT NULL,  
  s_name       VARCHAR(25) NOT NULL,  
  s_address    VARCHAR(25) NOT NULL,  
  s_city       VARCHAR(10) NOT NULL,  
  s_nation     VARCHAR(15) NOT NULL,  
  s_region     VARCHAR(12) NOT NULL,  
  s_phone      VARCHAR(15) NOT NULL  
);
```

```
CREATE TABLE customer
```

```
(  
  c_custkey    INTEGER NOT NULL,  
  c_name       VARCHAR(25) NOT NULL,  
  c_address    VARCHAR(25) NOT NULL,  
  c_city       VARCHAR(10) NOT NULL,  
  c_nation     VARCHAR(15) NOT NULL,  
  c_region     VARCHAR(12) NOT NULL,  
  c_phone      VARCHAR(15) NOT NULL,  
  c_mktsegment VARCHAR(10) NOT NULL  
);
```

```
CREATE TABLE dwdate
```

```
(  
  d_datekey          INTEGER NOT NULL,  
  d_date             VARCHAR(19) NOT NULL,  
  d_dayofweek        VARCHAR(10) NOT NULL,  
  d_month            VARCHAR(10) NOT NULL,  
  d_year             INTEGER NOT NULL,  
  d_yearmonthnum     INTEGER NOT NULL,  
  d_yearmonth        VARCHAR(8) NOT NULL,  
  d_daynuminweek     INTEGER NOT NULL,  
  d_daynuminmonth    INTEGER NOT NULL,  
  d_daynuminyear     INTEGER NOT NULL,  
  d_monthnuminyear   INTEGER NOT NULL,  
  d_weeknuminyear    INTEGER NOT NULL,  
  d_sellingseason    VARCHAR(13) NOT NULL,  
  d_lastdayinweekfl  VARCHAR(1) NOT NULL,  
  d_lastdayinmonthfl VARCHAR(1) NOT NULL,  
  d_holidayfl        VARCHAR(1) NOT NULL,  
  d_weekdayfl        VARCHAR(1) NOT NULL  
);
```

```
CREATE TABLE lineorder
```

```
(  
  lo_orderkey        INTEGER NOT NULL,  
  lo_linenumbers     INTEGER NOT NULL,  
  lo_custkey         INTEGER NOT NULL,  
  lo_partkey         INTEGER NOT NULL,  
  lo_suppkey         INTEGER NOT NULL,  
  lo_orderdate       INTEGER NOT NULL,  
  lo_orderpriority   VARCHAR(15) NOT NULL,  
  lo_shippriority    VARCHAR(1) NOT NULL,  
  lo_quantity        INTEGER NOT NULL,  
  lo_extendedprice   INTEGER NOT NULL,  
  lo_ordertotalprice INTEGER NOT NULL,  
  lo_discount         INTEGER NOT NULL,  
  lo_revenue         INTEGER NOT NULL,  
  lo_supplycost      INTEGER NOT NULL,  
);
```

```
lo_tax          INTEGER NOT NULL,  
lo_commitdate  INTEGER NOT NULL,  
lo_shipmode    VARCHAR(10) NOT NULL  
);
```

Step 5: Run the COPY commands

You run COPY commands to load each of the tables in the SSB schema. The COPY command examples demonstrate loading from different file formats, using several COPY command options, and troubleshooting load errors.