



# Nutanix Community Edition

ON VMWARE ESXI 6.0

Name | Date

# Contents

- Executive summary ..... 2
- Legal Information ..... 2
- Audience ..... 3
- Objectives ..... 3
- Pre-Requisites ..... 3
- Other recommendations ..... 4
  - Community Edition Licensing Requirements ..... 4
- Access to Community edition ..... 5
  - Signup ..... 5
  - Download ..... 5
  - Deploy ..... 5
  - How to bypass memory, SSD and its size requirements ..... 6
- Getting started ..... 7
  - Step 1: Enable SSD ..... 7
  - Step 2 Enable nested virtualization ..... 7
  - Step 3 Download Nutanix ce ..... 8
  - Step 4 Enable promiscuous mode ..... 8
  - Step 5 Upload vmdk ..... 10
  - Step 6 Prepare to create VM ..... 11
  - Step 7 Create VM ..... 12
  - Step 8: Power on the VM and install nutanix ce ..... 23
- Configuring a Nutanix CE cluster ..... 28
  - Install an SSH client: ..... 28
    - SSH under Linux ..... 28
- Single-node and Multi-node clusters ..... 28
  - Create a single-node cluster ..... 29
  - Create a multi-node cluster ..... 29
- Starting and stopping services ..... 32
  - Starting ..... 32
  - Stopping ..... 32

## Executive summary

Nutanix Community Edition is a free version of Nutanix AOS, which powers the Nutanix enterprise cloud platform. The Community Edition of AOS is designed for people interested in test driving its main features on their own test hardware and infrastructure. As stated in the end user license agreement, Community Edition is intended for internal business operations and non-production use only.

Before you begin, consider that the commercially-available Nutanix NX Series systems are tuned for performance. The enterprise-ready version of AOS is bundled as the engine driving our hyper converged compute and storage platform. You should not expect to achieve comparable benchmark numbers. For example, the NX Series systems include 10 GbE network interfaces, high core count dual CPUs, large amounts of memory, and high-performance enterprise-class solid-state drives (SSDs).

Test drive and experience the technology that powers the world’s largest datacenters, at no cost. Hyper-converged Test Drive enables IT administrators and tech enthusiasts to fully experience the power of the Nutanix Enterprise Cloud Platform. Go beyond hyper-converged infrastructure and get hands-on experience with the Enterprise Cloud Platform that powers the world's largest datacenters.

## Legal Information

- The author assumes no liability whatsoever for any direct or indirect damage, loss, inconvenience or other unintended consequences caused by using, or following the information presented in this guide.
- Whilst great care and attention has been taken whilst producing this guide the author cannot guarantee the accuracy of any of the information provided in this guide.
- The guide and the information presented in this guide is based on the “beta” version of the product which will be subject to unspecified changes and modifications over time, outside of the author’s control.
- “Nutanix”, “Nutanix Community Edition”, “Nutanix CE” and other associated trademarks, logos and devices remain the sole property of their respective owners and are only used in this guide to explain a concept or illustrate a procedure.
- This guide is not officially endorsed, recognized or affiliated with “Nutanix” or any of its subsidiary’s.
- At the time of writing this guide the author is not employed by “Nutanix”, and has never been employed by “Nutanix”.
- Unless clearly stated, no other contributing authors of this guide are employees of “Nutanix”.
- The author is not officially recognized or certified by “Nutanix” for providing information or training on “Nutanix” products and services other than having passed the following Nutanix Partner Accreditations.

o “NPSR – Nutanix Platform Sales Representative” o “NPSE – Nutanix Platform Sales Engineer” o “NPP – Nutanix Platform Professional” .

- The only official source of product information and product support for “Nutanix Community Edition” is from the “Nutanix NEXT community”.

## Audience

This guide is intended primarily at IT professionals and virtualization consultants with an interest in learning and deploying “Nutanix Community Edition” in a “home-lab”, “test/dev” or “non-production” environment.

The information found in this guide should not be used or followed in a production environment or live business environment, it is provided for the sole purpose of helping individuals learn “Nutanix Community Edition” in a small training environment.

Hyper-converged infrastructure enthusiasts, System Admin who would like to gain hands on experience Nutanix platform and

- IT pros
- IT Managers
- SysAdmins
- Storage Admin

## Objectives

This document provides step by step instruction to install Nutanix CE. This primary focus of this guide is to provide the user with enough information to:

- Download and install “Nutanix Community Edition”.
- Perform initial configuration of a 1, 3 or 4 node cluster.
- Configure Storage.
- Configure Networking.
- Build a VM on the ESXi hypervisor.
- Provide a basic overview of the PRISM dashboard and management interface.
- Start, Stop and Shutdown the cluster in a controlled manner

## Pre-Requisites

To Install Nutanix CE on commodity hardware following are the requirements.

- Servers 1, 3 & 4 servers
- CPU Intel CPUs, 4 cores minimum, with VT-x support
- Memory 16GB minimum
- Storage Subsystem RAID 0 (LSI HBAs) or AHCI storage sub-systems
- Hot Tier (SSD) One SSD per server minimum, ≥ 200GB per server
- Cold Tier (HDD) One HDD per server minimum, ≥ 500GB per server
- Networking Intel NICs
- ESXi

## Other recommendations

Component	Recommendation
Cluster size	Community Edition allows you to install single-node, three-node, and four-node clusters. Depending on your available hardware, Nutanix recommends a three-node cluster to exercise the features of Community Edition most effectively.
DHCP Server	If you are using a DHCP server, ensure that static IP addresses are assigned for hypervisor hosts and Controller VMs. Nutanix recommends that you not use dynamic IP addressing for hosts or Controller VMs.
Imaging software	Linux: use the included dd utility Mac OS X: use the included dd utility Windows: <a href="#">ImageUSB</a> from PassMark Software (freeware)

## COMMUNITY EDITION LICENSING REQUIREMENTS

To use the Community Edition, Nutanix requires the following actions and access. If you fail to meet these requirements, access to your cluster will be blocked. The cluster remains operating but you will be unable to log on to it. Otherwise, there are no licensing file requirements as with the commercially-available Nutanix products.

Requirement	Enforcement
Admin account password change from default	When you first log on to the Prism web console with the admin account credentials, you must immediately change the admin user password from its default.

Requirement	Enforcement
Nutanix Next account with log on credentials and the Community Edition role enabled	After you first log on to the Prism web console and change the admin credentials, you are prompted to provide these credentials.
Outgoing TCP ports 80 or 8443 opened through your firewall	An automated feedback mechanism named Pulse that sends cluster alerts and usage statistics to Nutanix requires internet connectivity and port access.
Install any pending Community Edition updates	Nutanix make updates available occasionally. When an upgrade is available, but not yet downloaded, you must upgrade within 30 calendar days. You must install the available upgrade immediately once you have downloaded it. If you fail to meet these requirements, access to your cluster will be blocked. The cluster remains operating but you will be unable to log on to it.

## Access to Community edition

### SIGNUP

Sign up to access and download Community Edition software.

<http://www.nutanix.com/products/community-edition/register>

### DOWNLOAD

Once signed up, the link to download the software can be found in this forum (registration required):

<http://next.nutanix.com/t5/Nutanix-Community-Edition/ct-p/NutanixCommunity-Edition>

Initial Install Image [ce-2017.07.20-stable.img.gz](http://www.nutanix.com/cwm/ce-2017.07.20-stable.img.gz) 9a04da41b01db43ae6e1f31de8986057

### DEPLOY

Deploy Community Edition on up to four servers using a broad variety of hardware.

<http://www.nutanix.com/cwm/ce-hardware-table.html>

## HOW TO BYPASS MEMORY, SSD AND ITS SIZE REQUIREMENTS

How to mark a disk as SSD – In my case, I have Dell gear with PERC controllers. As some of you may know, they don't support JBOD. In my lab I just configure each drive in a single drive raid 0 with caching turned off. However, the SSD is still seen as a traditional disk. So, here is how to bypass to complete a Nutanix Install

After the image boots up, login using username "nutanix" and password "nutanix/4u"

- type "fdisk -l" and identify which disk your SSD is. In my case it was "sdb"
- type the following to identify if it is rotational or SSD
- cat /sys/block/sdb/queue/rotational (0 means SSD, 1 means HDD)

type the following to change the identified type

- echo 0 > /sys/block/sdb/queue/rotational
- type the following to confirm the change
  - cat /sys/block/sdb/queue/rotational (0 means SSD, 1 means HDD)

How to bypass the SSD size requirement – In my case, my SSD was only 120GB while there is a requirement for 200GB or more. Easily handled,

- Login as root (Password – "nutanix/4u")
- Change to directory – /home/install/phx\_iso/phoenix
- Edit minimum\_reqs.py (ex. nano minimum\_reqs.py) and change the two instances of "199" to whatever fits your system. In my case, I just set it to 100.

How to modify the memory size requirement

- Login as root (Password – "nutanix/4u")
- Change to directory – /home/install/phx\_iso/phoenix
- Edit minimum\_reqs.py

under "if os.environ.has\_key('COMMUNITY\_EDITION'):" change "MIN\_MEMORY\_GB = " to something below the current system inventory you are trying to install on is

How to modify the CPU core size requirement

- Login as root (Password – "nutanix/4u")
- Change to directory – /home/install/phx\_iso/phoenix
- Edit minimum\_reqs.py

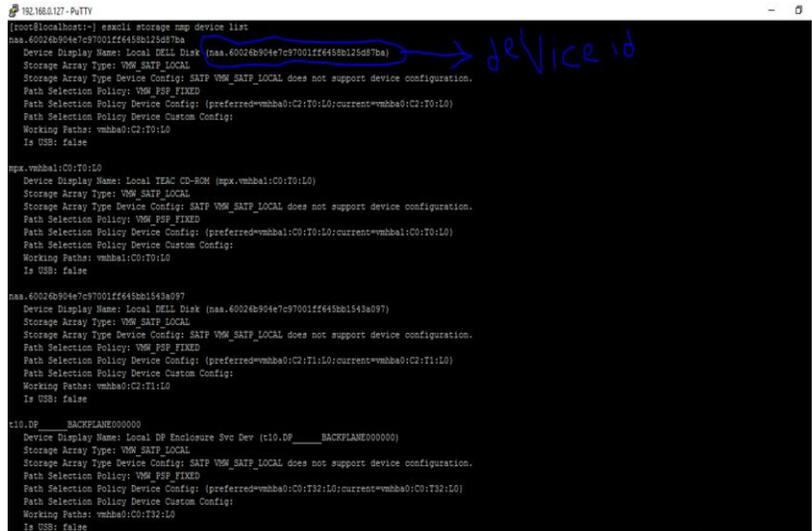
under "if os.environ.has\_key('COMMUNITY\_EDITION'):" change "MIN\_CORES = " to something below the current system inventory you are trying to install on.

# Getting started

## STEP 1: ENABLE SSD

Enable SSD in ESXi. It is required that SSD option need to be enabled if not detected by default. SSD drives are used for the following reasons.

Esxcli is used to enable SSD.  
Identify the Device name using esxcli  
# esxcli storage nmp device list  
# esxcli storage nmp satp rule add --  
satp=VMW\_SATP\_LOCAL --device  
naa.60026b904e7c97001ff645bb1543a0  
97 --option "enable\_ssd"



```
192.168.0.127 - PuTTY
[root@localhost:~]# esxcli storage nmp device list
naa.60026b904e7c97001ff645bb1543a097
Device Display Name: Local DELL Disk (naa.60026b904e7c97001ff645bb1543a097)
Storage Array Type: VMW_SATP_LOCAL
Storage Array Type Device Config: SATP VMW_SATP_LOCAL does not support device configuration.
Path Selection Policy: VMW_PSP_FIXED
Path Selection Policy Device Config: (preferred=vmmba0:C2:T0:L0;current=vmmba0:C2:T0:L0)
Path Selection Policy Device Custom Config:
Working Paths: vmmba0:C2:T0:L0
Is USB: false
mpx.vmhbal:C0:T0:L0
Device Display Name: Local TEAC CD-ROM (mpx.vmhbal:C0:T0:L0)
Storage Array Type: VMW_SATP_LOCAL
Storage Array Type Device Config: SATP VMW_SATP_LOCAL does not support device configuration.
Path Selection Policy: VMW_PSP_FIXED
Path Selection Policy Device Config: (preferred=vmhbal:C0:T0:L0;current=vmhbal:C0:T0:L0)
Path Selection Policy Device Custom Config:
Working Paths: vmhbal:C0:T0:L0
Is USB: false
naa.60026b904e7c97001ff645bb1543a097
Device Display Name: Local DELL Disk (naa.60026b904e7c97001ff645bb1543a097)
Storage Array Type: VMW_SATP_LOCAL
Storage Array Type Device Config: SATP VMW_SATP_LOCAL does not support device configuration.
Path Selection Policy: VMW_PSP_FIXED
Path Selection Policy Device Config: (preferred=vmmba0:C2:T1:L0;current=vmmba0:C2:T1:L0)
Path Selection Policy Device Custom Config:
Working Paths: vmmba0:C2:T1:L0
Is USB: false
t10.DF _____ BACKPLANE000000
Device Display Name: Local DF Enclosure 3vc Dev (t10.DF _____ BACKPLANE000000)
Storage Array Type: VMW_SATP_LOCAL
Storage Array Type Device Config: SATP VMW_SATP_LOCAL does not support device configuration.
Path Selection Policy: VMW_PSP_FIXED
Path Selection Policy Device Config: (preferred=vmmba0:C0:T32:L0;current=vmmba0:C0:T32:L0)
Path Selection Policy Device Custom Config:
Working Paths: vmmba0:C0:T32:L0
Is USB: false
```

## STEP 2 ENABLE NESTED VIRTUALIZATION

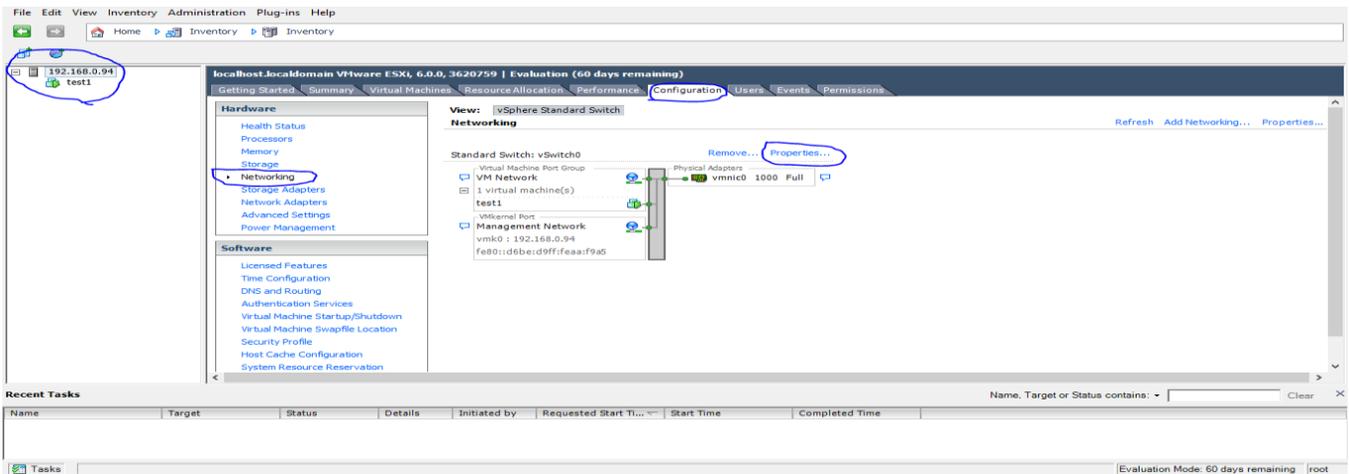
Enable Login to vmware esxi 6.0 using putty and navigate to cd /etc/vmware/config.

We have to add line at the end of file “vhv.enable = “true” to enable the nested virtualization as shown in the figure below enable nested virtualization as we are installing on top of VMware esxi

What is nested virtualization:

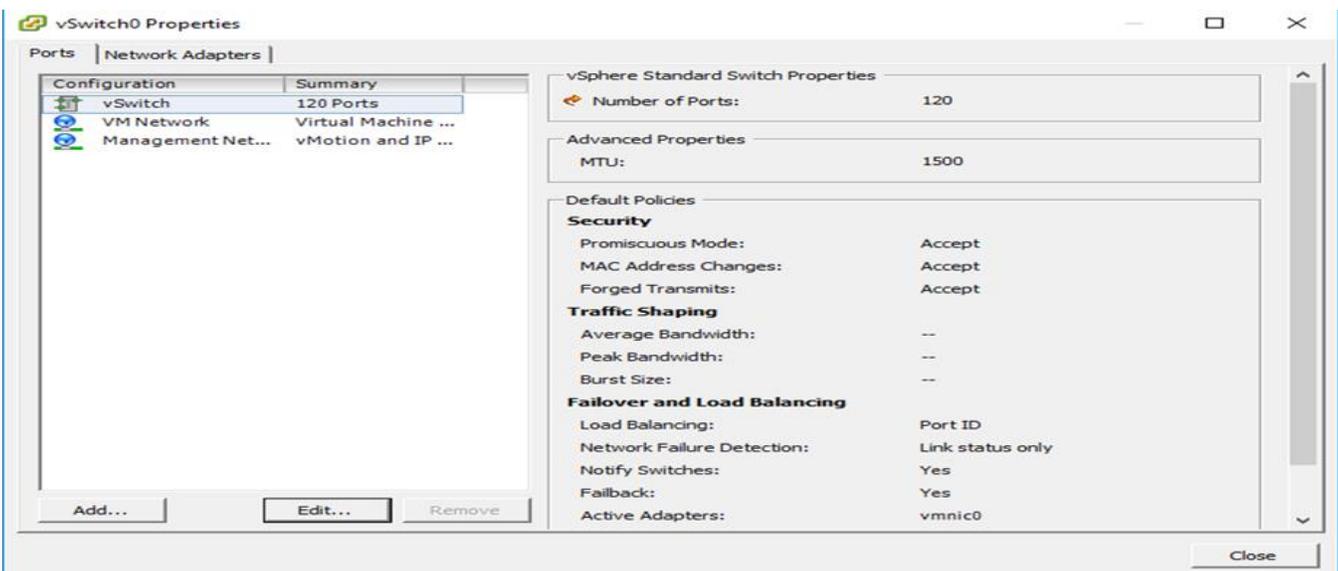
Nested virtualization is a virtual machine contained within another virtual machine. The real benefit here is that you can run multiple nested nodes on one physical box.



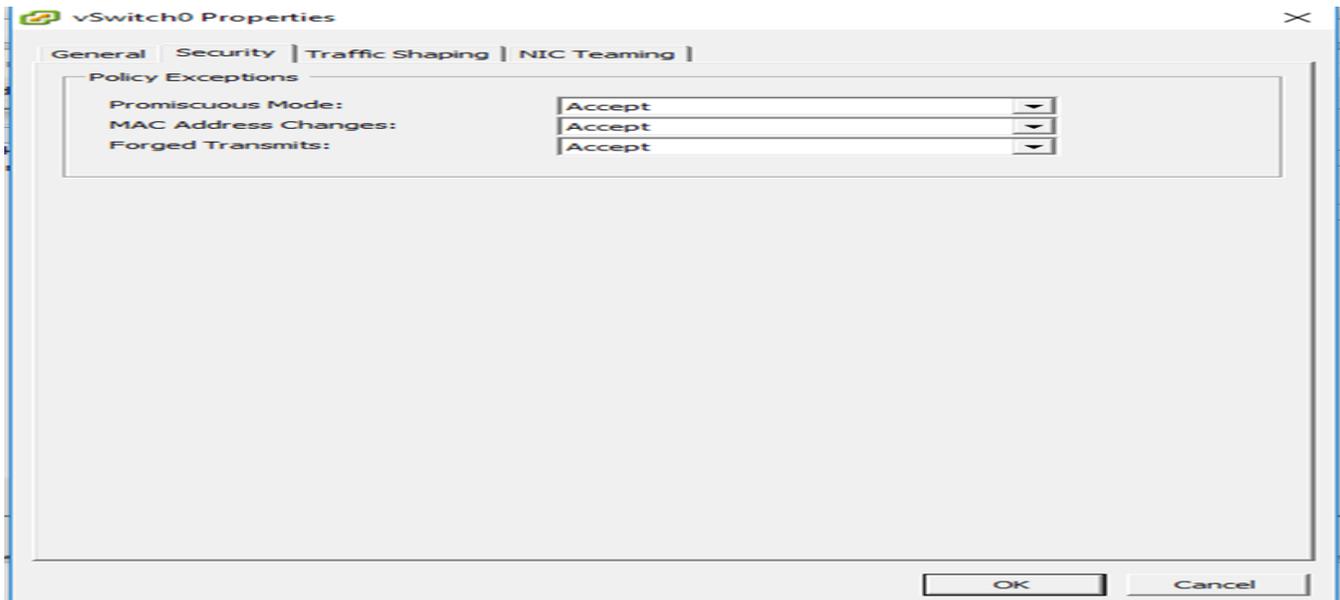


- Select the vSwitch or any of the below options and click on edit settings

Note: Lets understand what promiscuous mode: In the realm of computer networking, promiscuous mode refers to the special mode of Ethernet hardware, in particular network interface cards (NICs), that allows a NIC to receive all traffic on the network, even if it is not addressed to this NIC. If “promiscuous mode” on the vSwitch or VM network where Nutanix-ce VM is connected is NOT enabled the Controller VM (CVM) is not accessible.



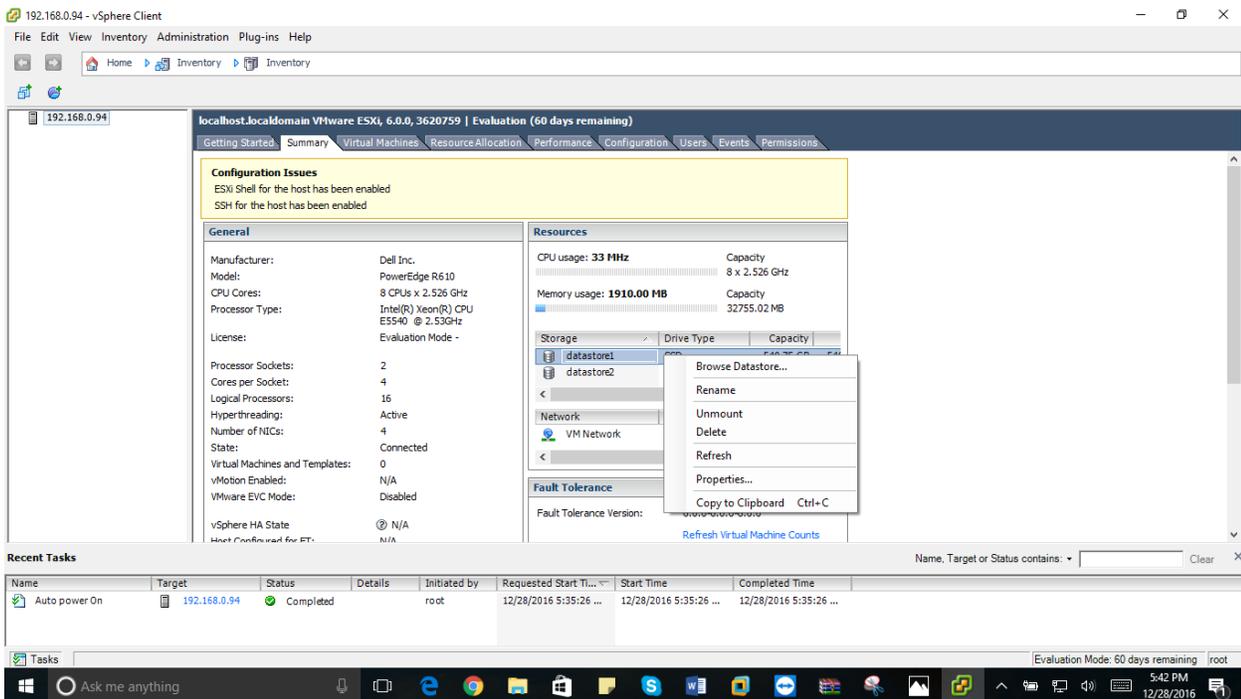
- Go the security tab and check the promiscuous mode as accept and forged transmits as accept as show in the figure below.
- You may also need to enable “forged transmits”



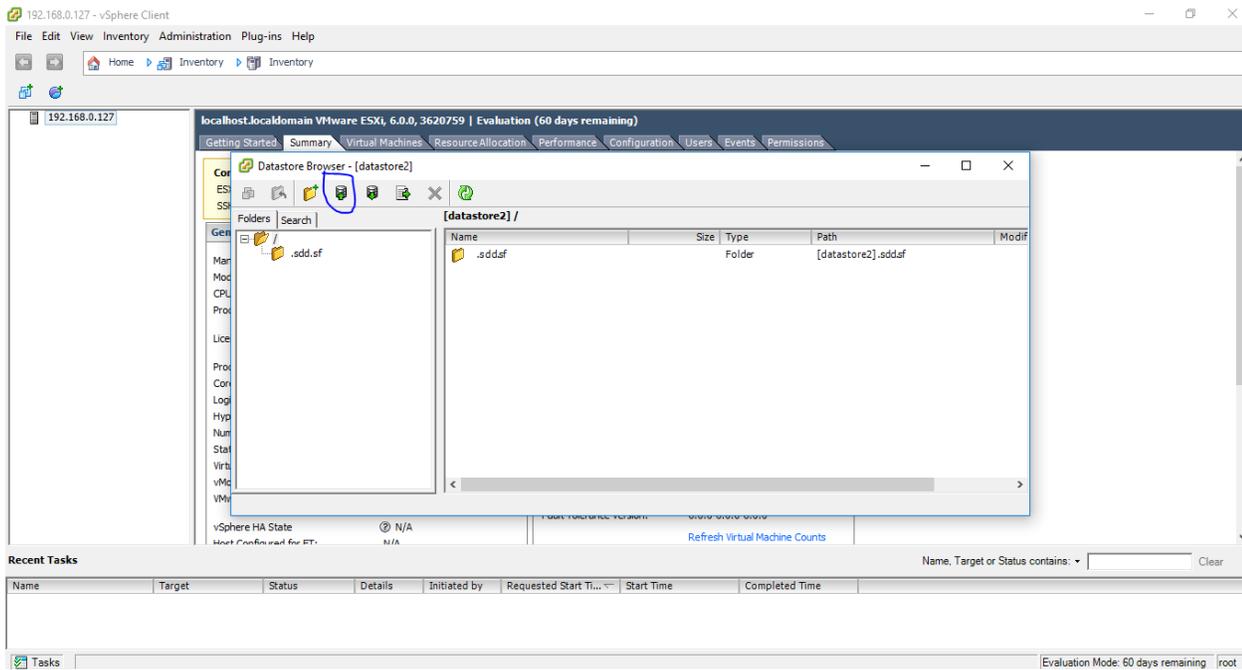
## STEP 5 UPLOAD VMDK

Upload ce-flat.vmdk and cd.vmdk into datastore.

- Right click on the datastore and select browse datastore.



- Click on the upload button on the top as shown in the picture and select both files and upload.



## STEP 6 PREPARE TO CREATE VM

Login to the VMware ESXi 6.0 using putty

Browse to /vmfs/volumes/datastore2 and change the ce-flat.vmdk.img to ce-flat.vmdk and ce.vmdk.txt to ce.vmdk using below commands

- `mv ce-flat.vmdk.img ce-flat.vmdk`
- `mv ce.vmdk.txt ce.vmdk`

```

192.168.0.127 - PuTTY
login as: root
Using keyboard-interactive authentication.
Password:
The time and date of this login have been sent to the system logs.

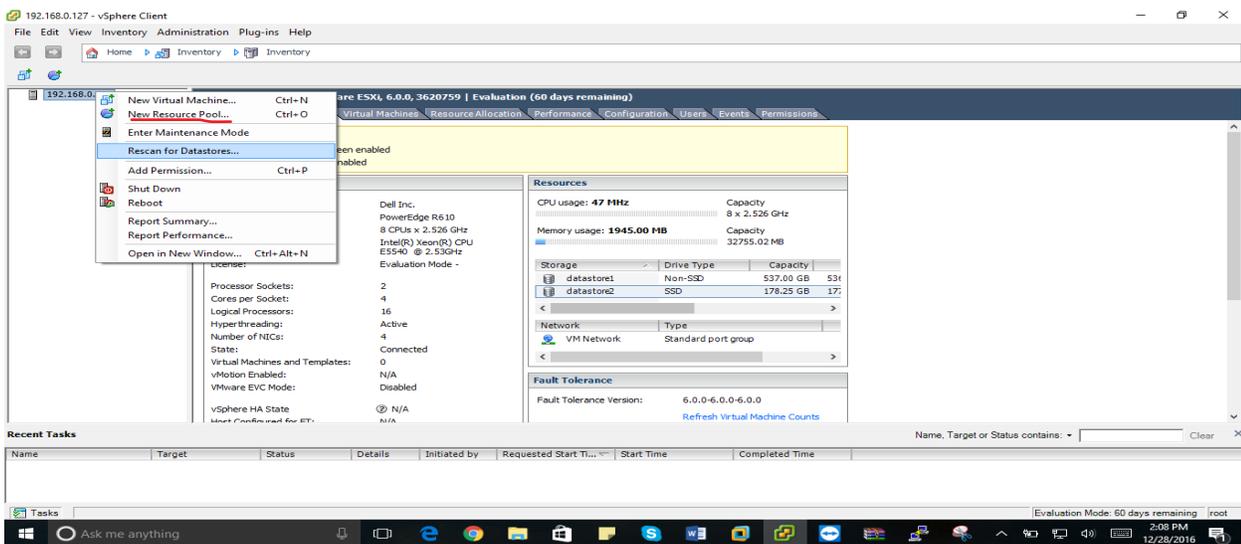
VMware offers supported, powerful system administration tools. Please
see www.vmware.com/go/sysadmintools for details.

The ESXi Shell can be disabled by an administrative user. See the
vSphere Security documentation for more information.
[root@localhost:~] cd vmfs
[root@localhost:vmfs] ls
devices volumes
[root@localhost:vmfs] cd volumes/
[root@localhost:vmfs/volumes] ls
0d540570-eea4f0a3-ecaa-819295c48950 586396a3-85e7172e-fbf5-180373f2d834
37bf687a-ba2b9397-f14d-ebba87719a43 586396a3-d2337d3e-b995-180373f2d834
58638c6b-4c298559-a3a4-180373f2d834 8a4aad78-c7bf21e1-949f-aa0d1d7535ff
58638c76-c8f116b6-5b73-180373f2d834 cb2b107e-789f038e-c727-cbc0dda54c23
58638c7c-3d656b71-c2af-180373f2d834 datastore1
586396a1-6a8b75b8-4047-180373f2d834 datastore2
[root@localhost:vmfs/volumes] cd datastore2
[root@localhost:vmfs/volumes/586396a3-85e7172e-fbf5-180373f2d834] ls
ce-flat.vmdk.img ce.vmdk.txt
[root@localhost:vmfs/volumes/586396a3-85e7172e-fbf5-180373f2d834] cd
[root@localhost:~] cd /vmfs/volumes/
[root@localhost:vmfs/volumes] ls
0d540570-eea4f0a3-ecaa-819295c48950 58638c76-c8f116b6-5b73-180373f2d834 586396a3-85e7172e-fbf5-180373f2d834 cb2b107e-789f038e-c727-cbc0dda54c23
37bf687a-ba2b9397-f14d-ebba87719a43 58638c7c-3d656b71-c2af-180373f2d834 586396a3-d2337d3e-b995-180373f2d834
58638c6b-4c298559-a3a4-180373f2d834 586396a1-6a8b75b8-4047-180373f2d834 8a4aad78-c7bf21e1-949f-aa0d1d7535ff datastore1
[root@localhost:vmfs/volumes] cd datastore2
[root@localhost:vmfs/volumes/586396a3-85e7172e-fbf5-180373f2d834] ls
ce-flat.vmdk.img ce.vmdk.txt
[root@localhost:vmfs/volumes/586396a3-85e7172e-fbf5-180373f2d834] mv ce-flat.vmdk.img ce-flat.vmdk
[root@localhost:vmfs/volumes/586396a3-85e7172e-fbf5-180373f2d834] mv ce.vmdk.txt ce.vmdk
[root@localhost:vmfs/volumes/586396a3-85e7172e-fbf5-180373f2d834] ls
ce-flat.vmdk ce.vmdk
[root@localhost:vmfs/volumes/586396a3-85e7172e-fbf5-180373f2d834]

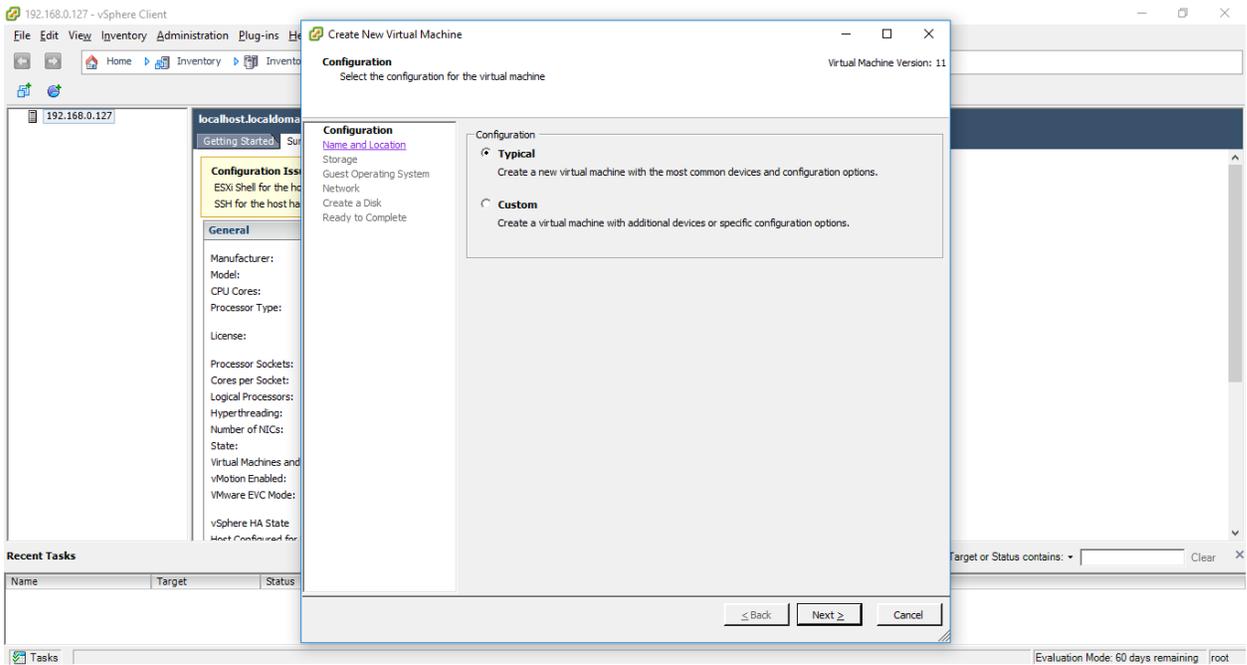
```

## STEP 7 CREATE VM

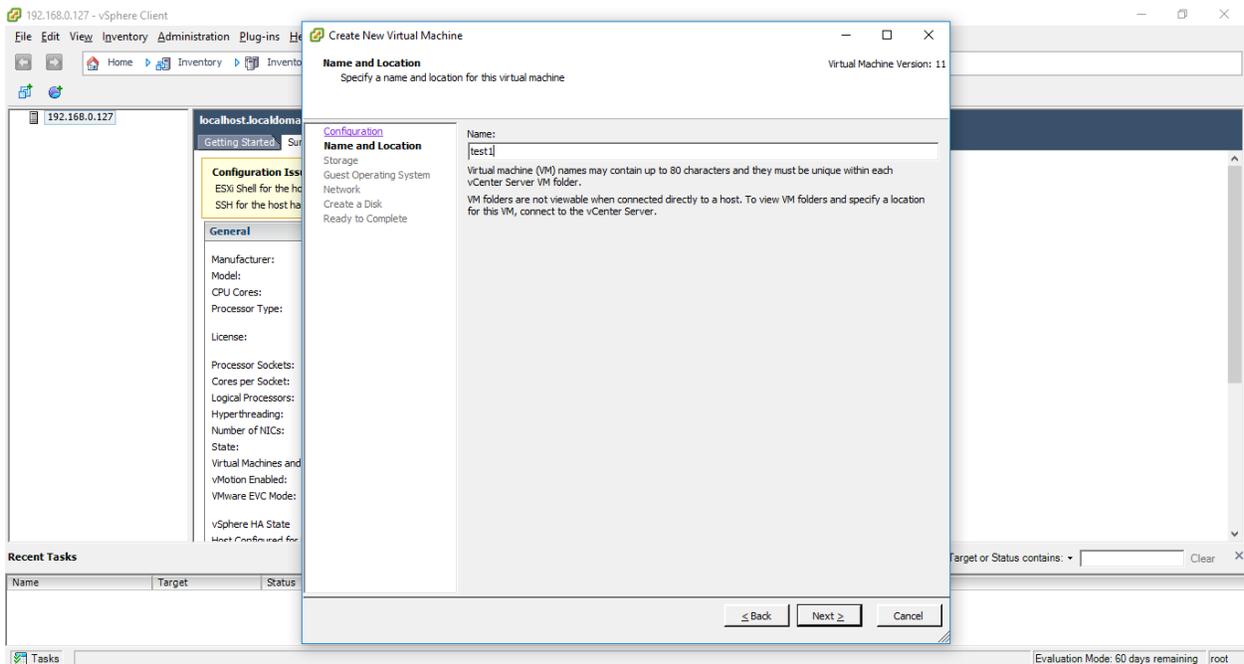
- Right click on the server and create a new virtual machine.



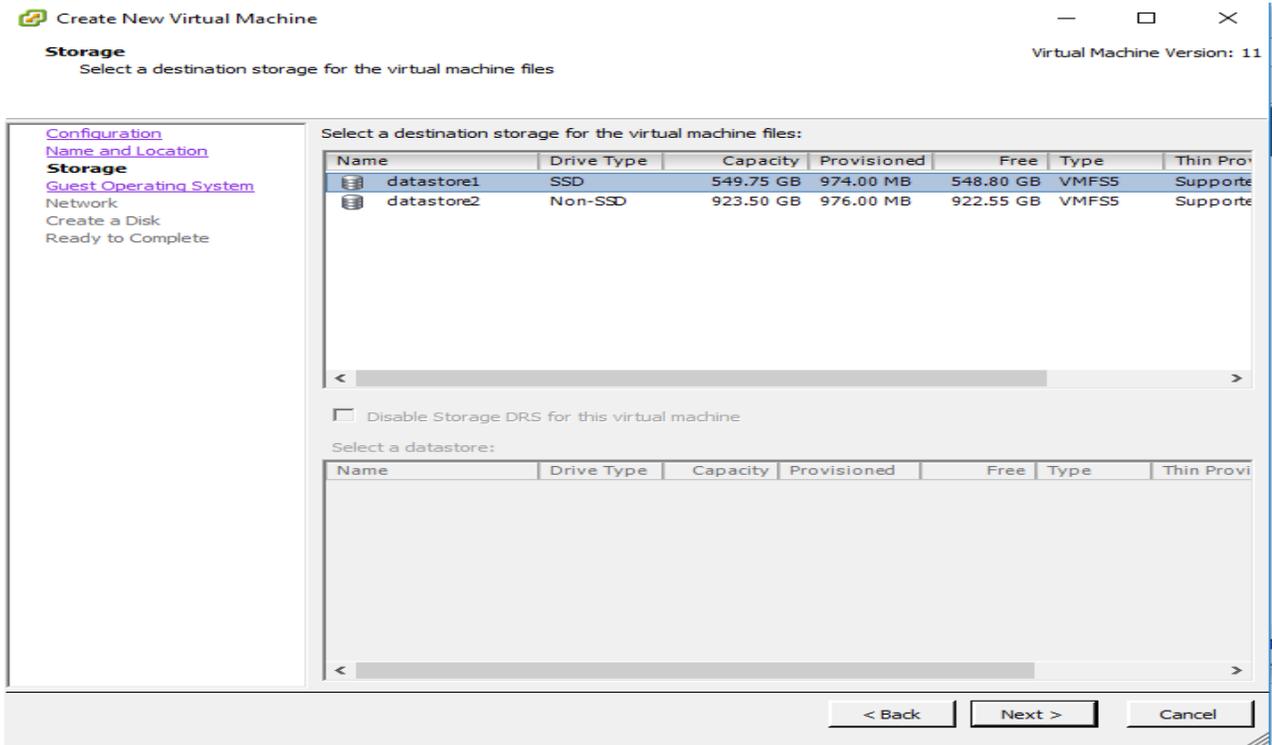
- Select typical configuration and click next



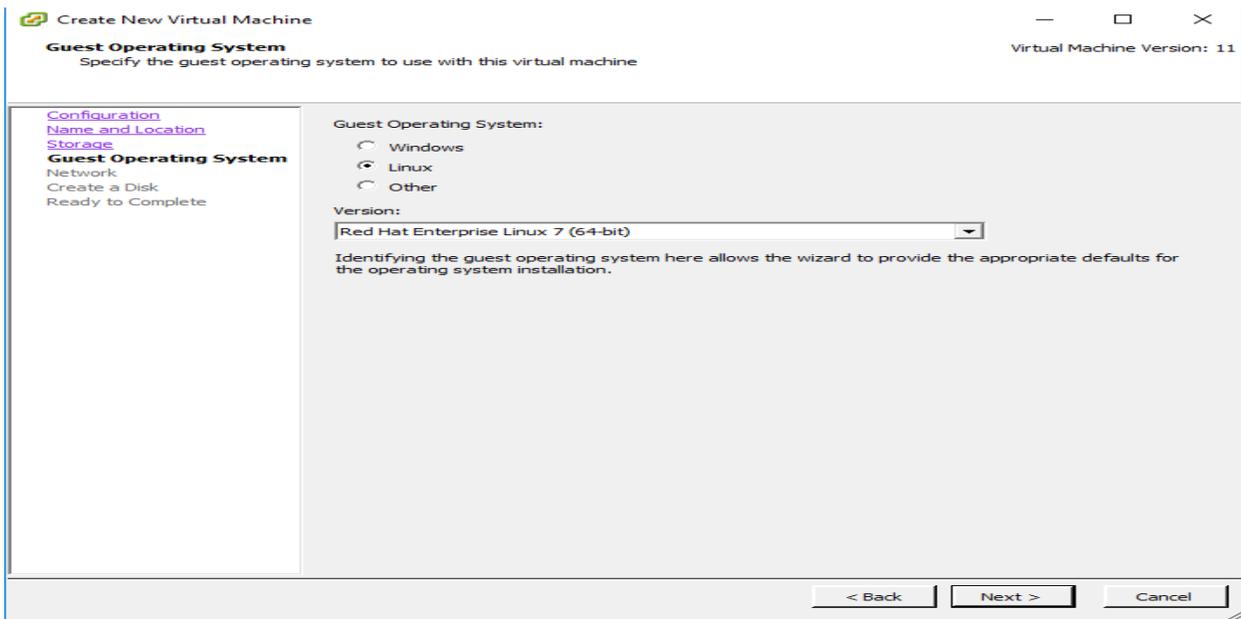
➤ Give name for the virtual machine and click next



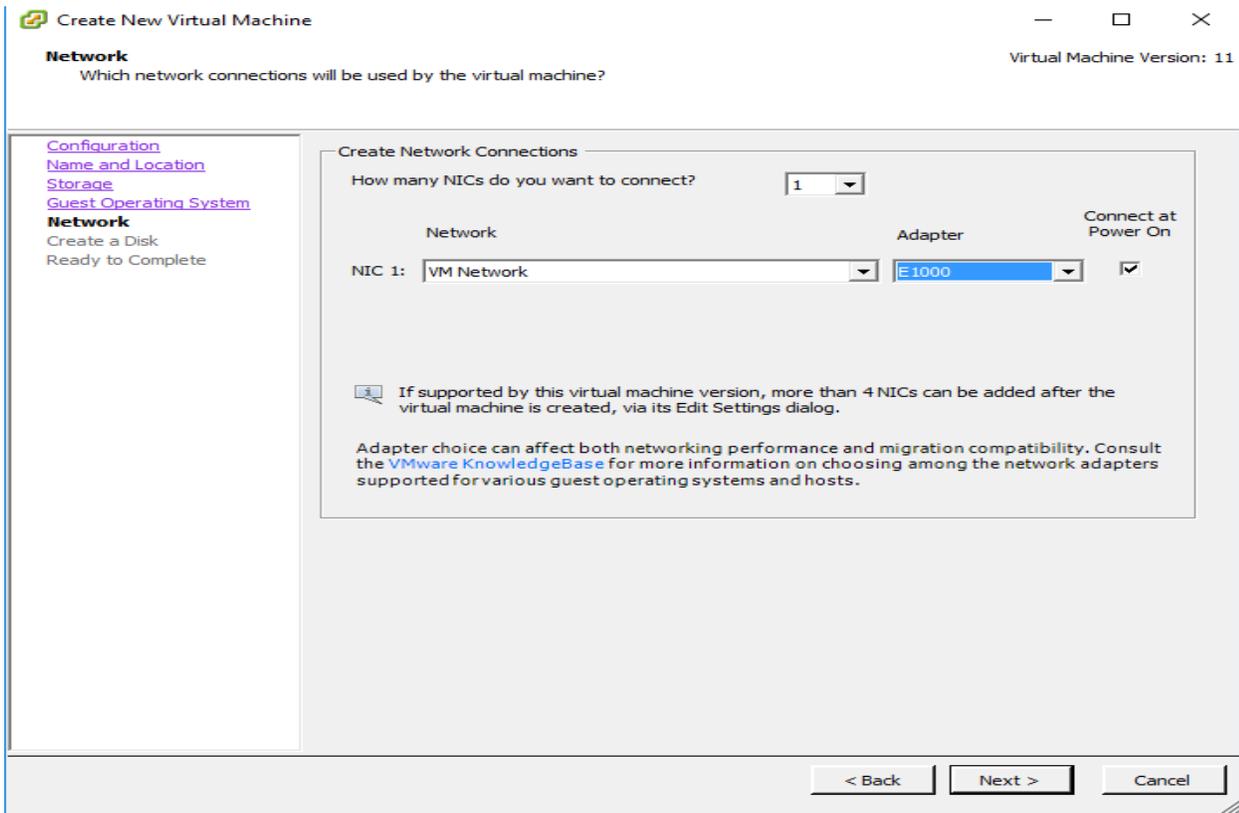
➤ Click on the datastore1 and click next, here in the below image Non-SSD is nothing but your HDD.



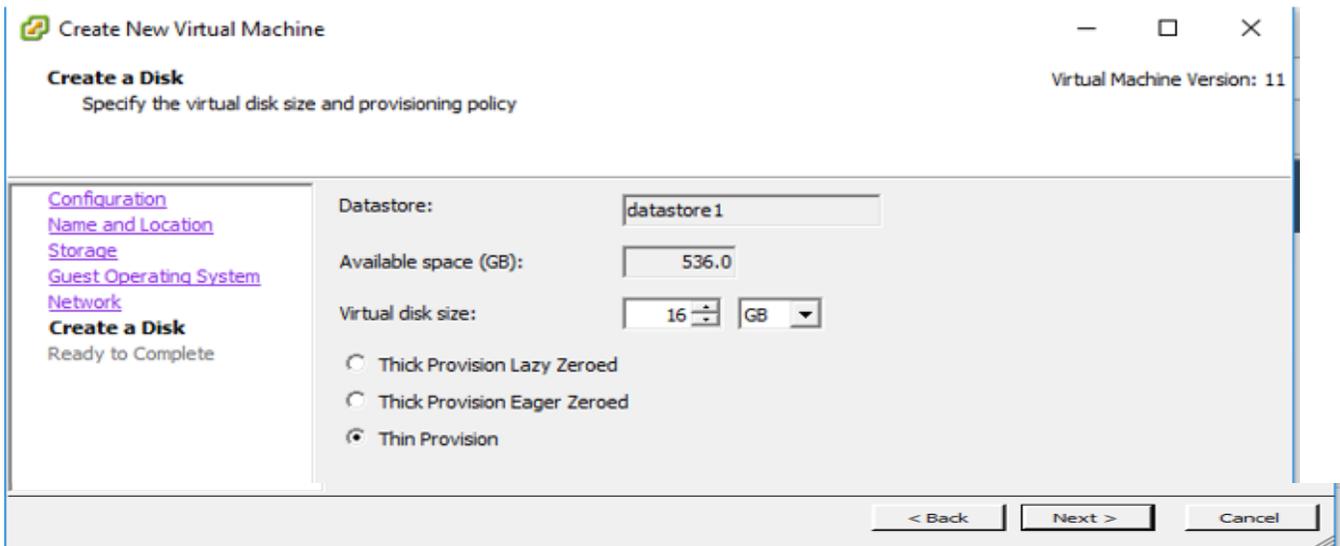
- Select guest operating system as linux and click next

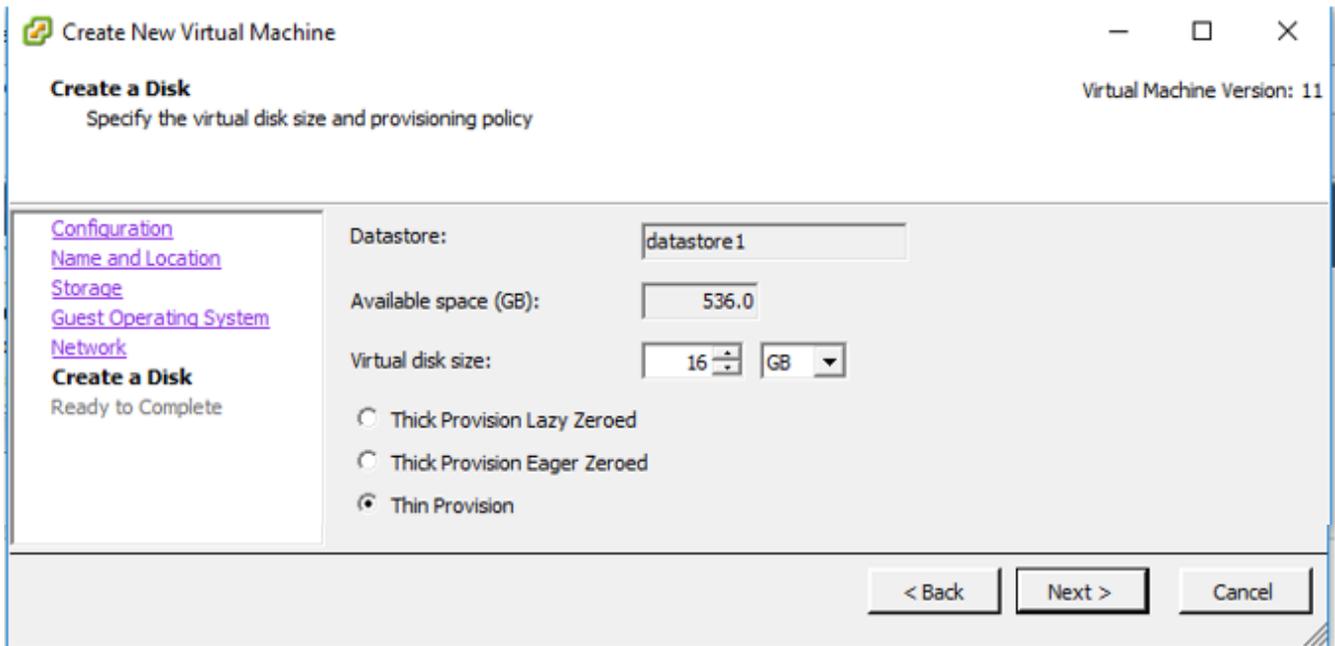


- Change the adapter to E1000 and click next

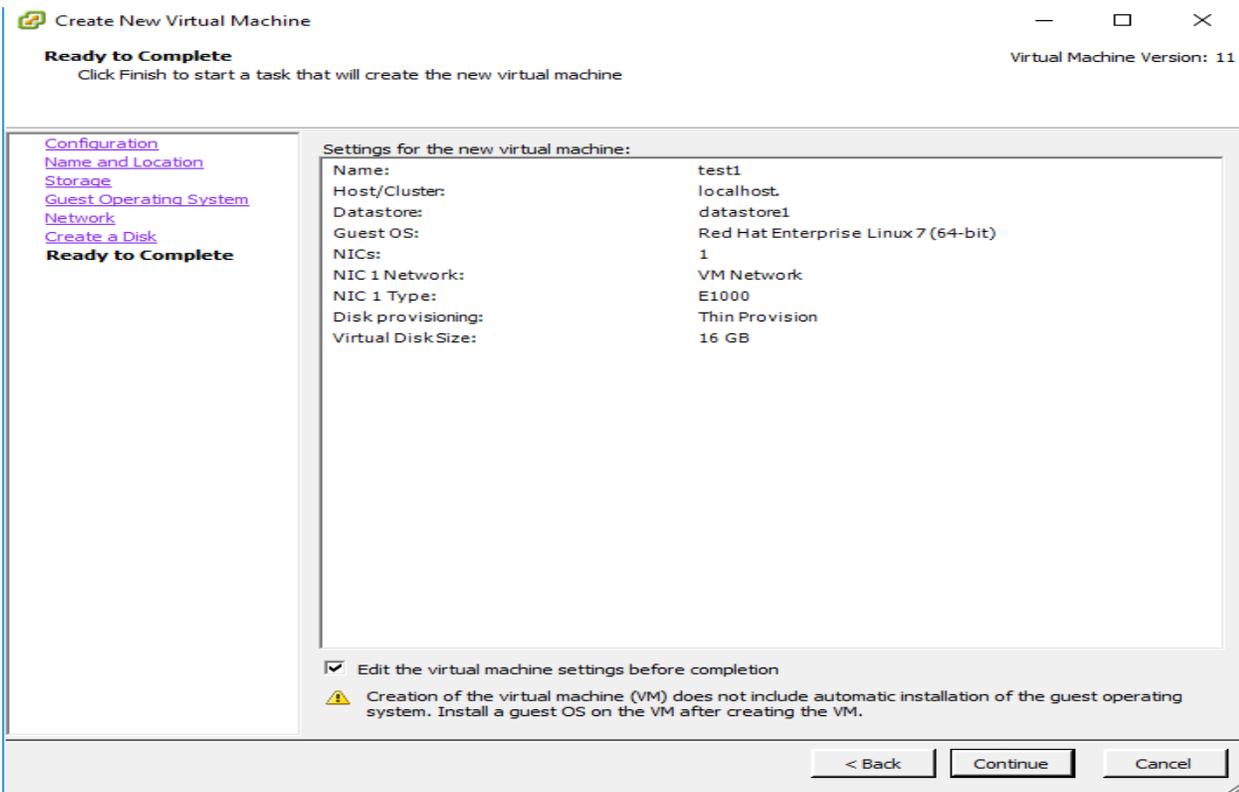


- Give virtual disk size as 16gb and select thin provision and click next.

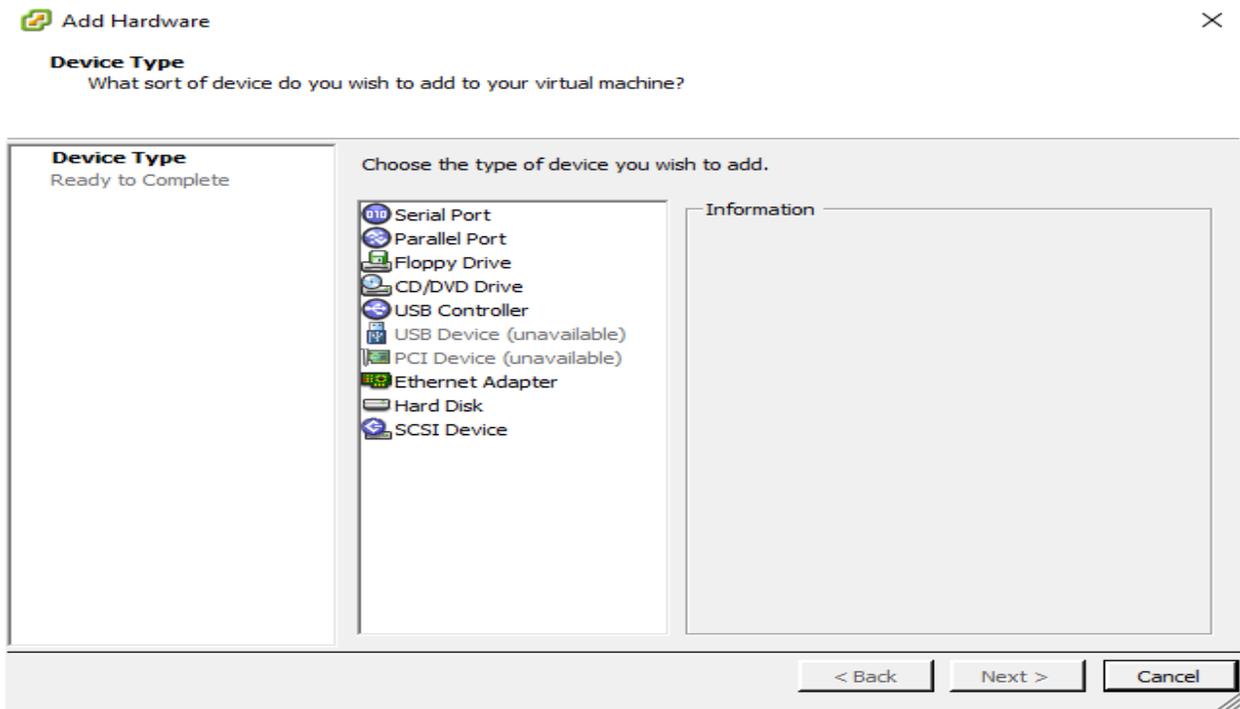




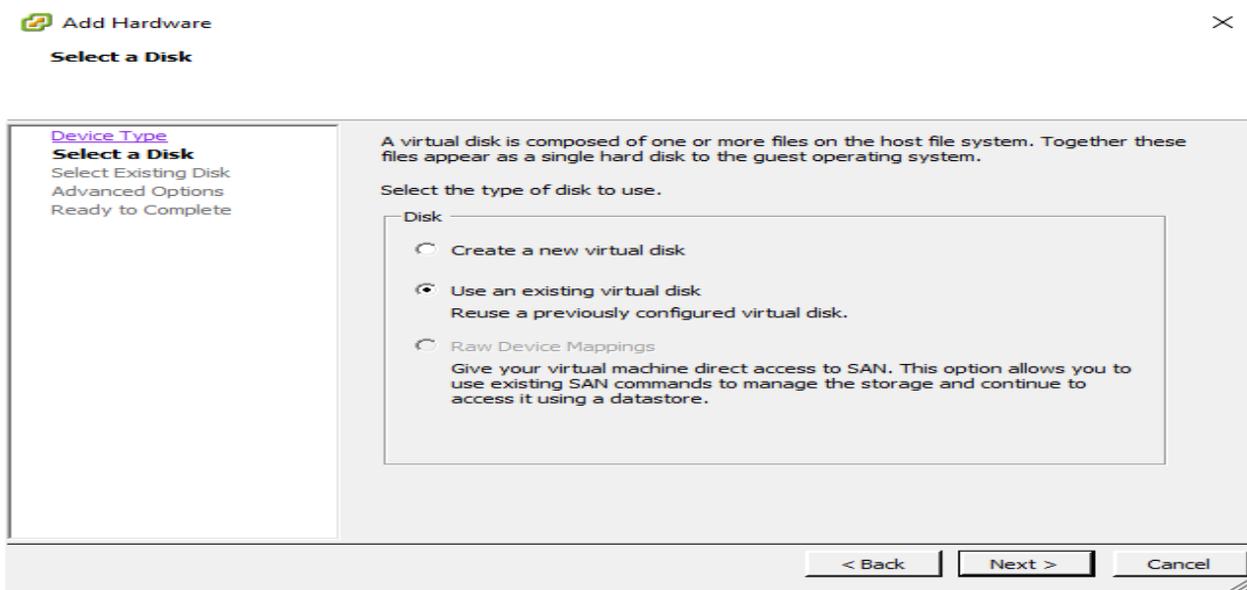
- Check the option edit virtual machine settings before completion and click next



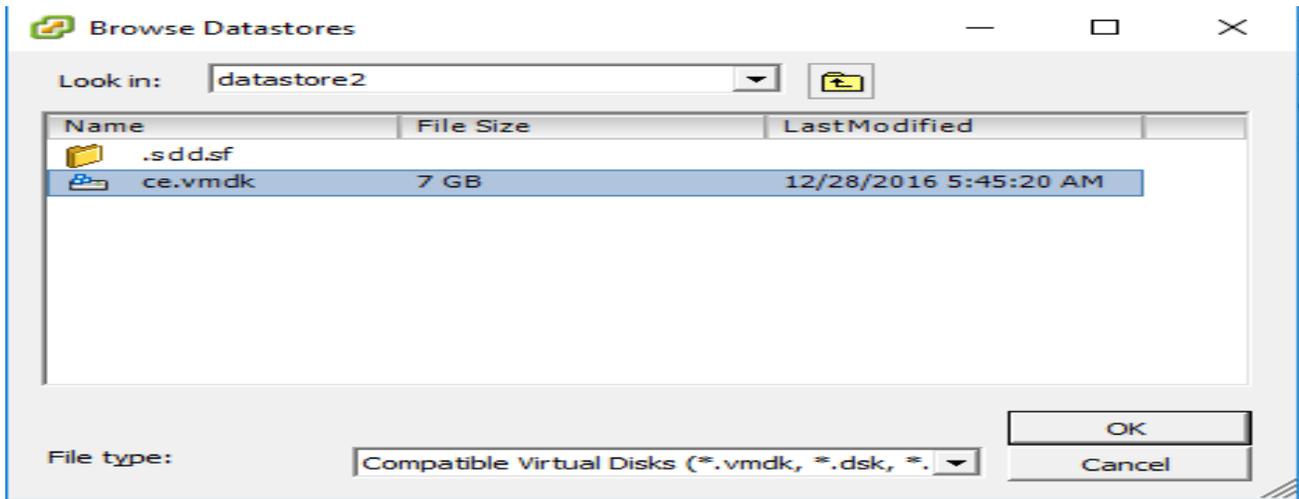
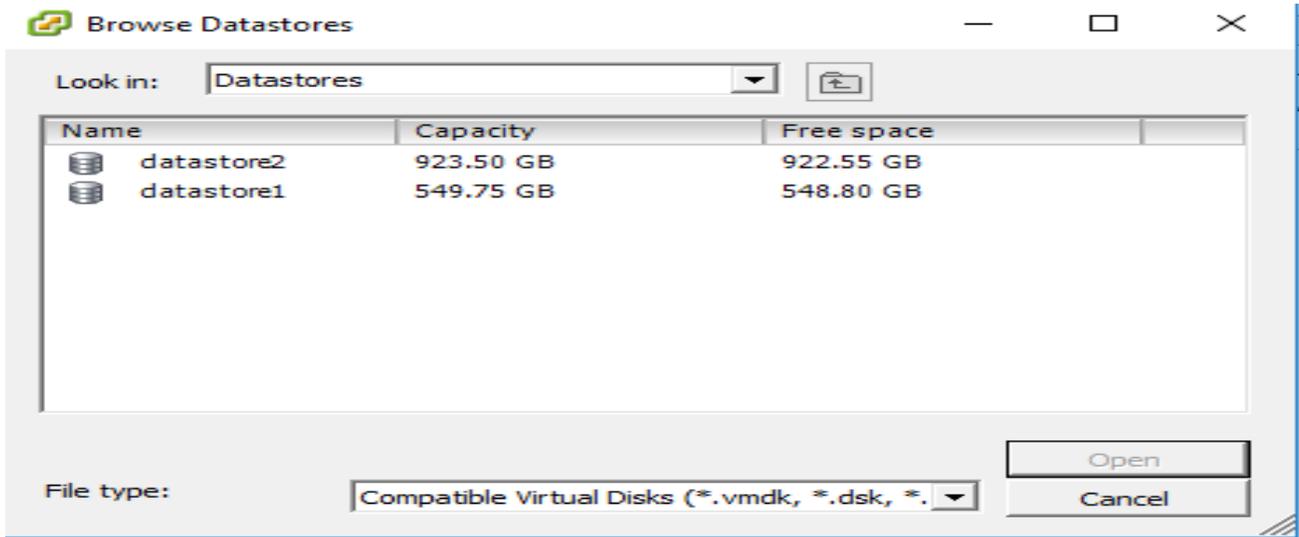
- Remove floppy drive, cd/dvd, new hard disk and next click on and select hard disk and click next.



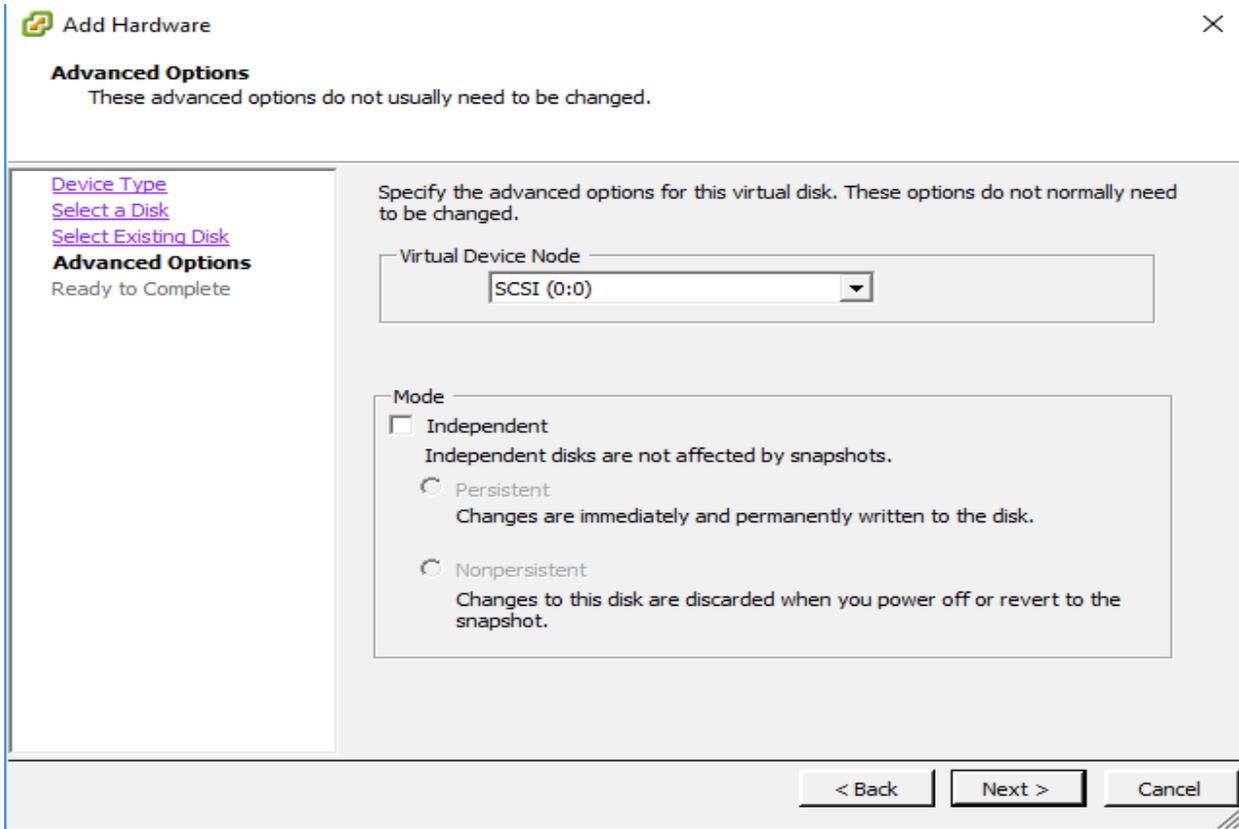
➤ Select the use an existing virtual disk and click next.



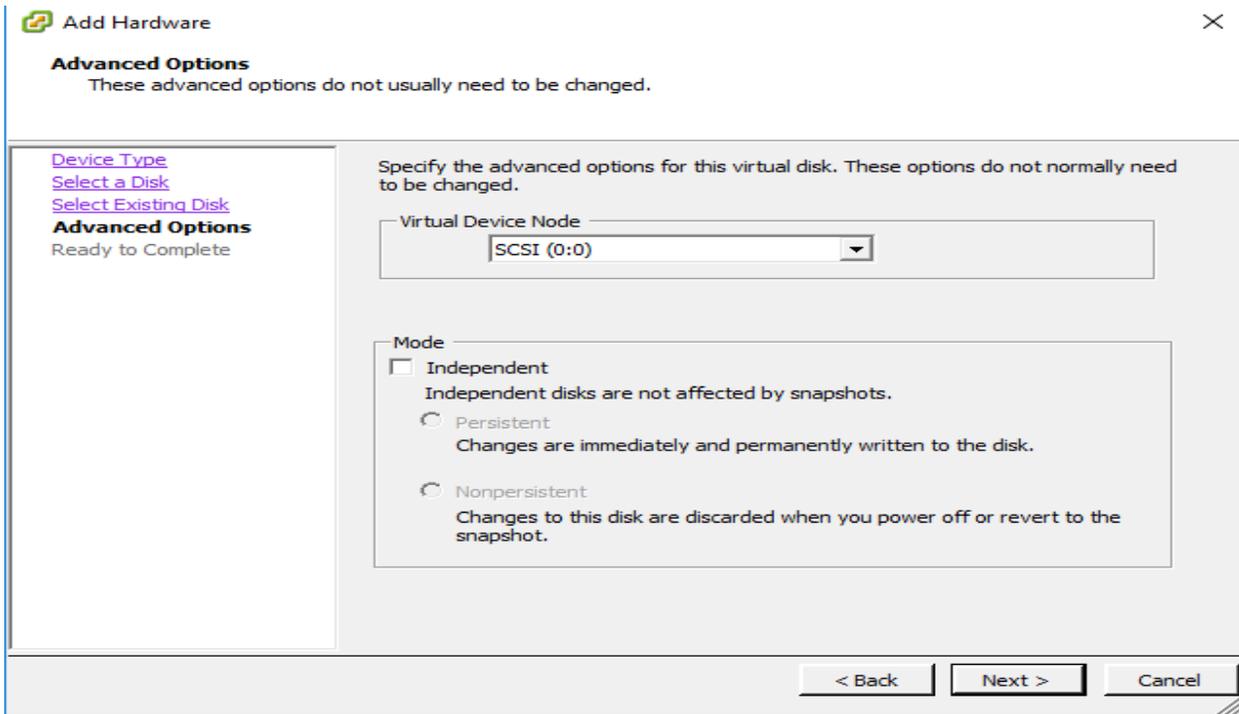
Select the datastore where you have stored the ce.vmdk file and select the ce.vmdk and click ok.



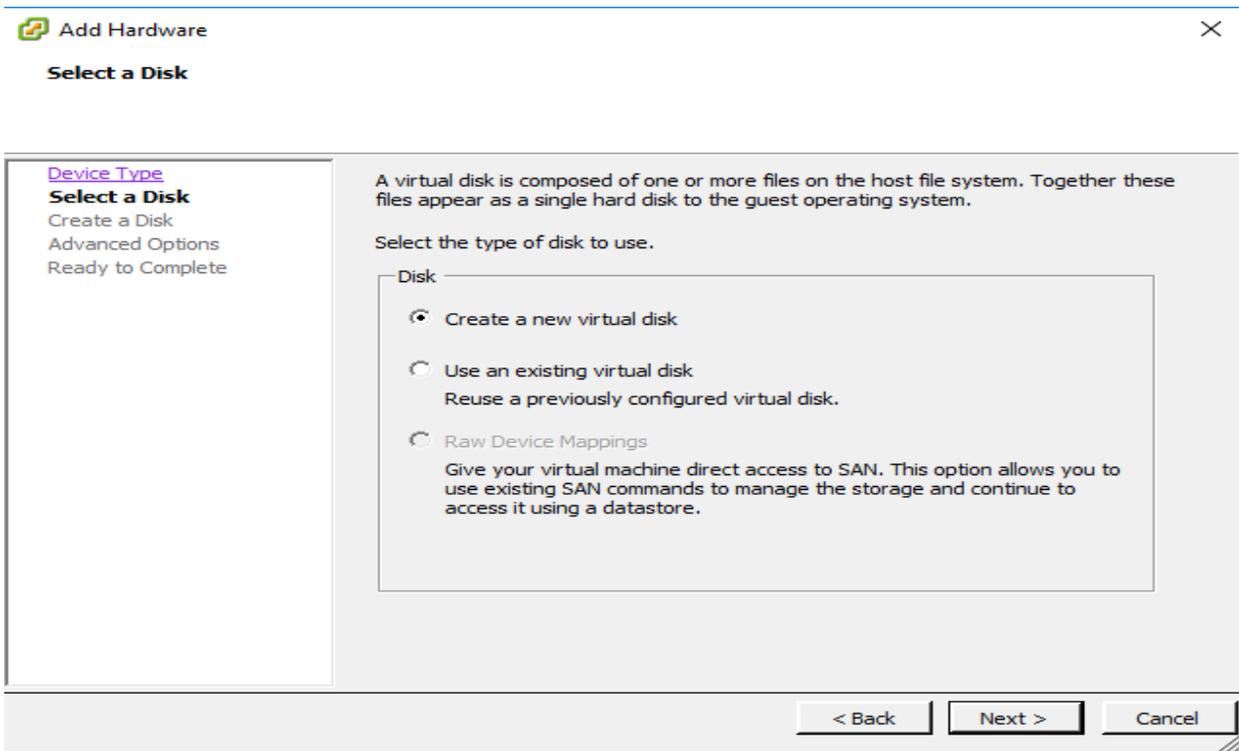
- Leave the options to default and click next



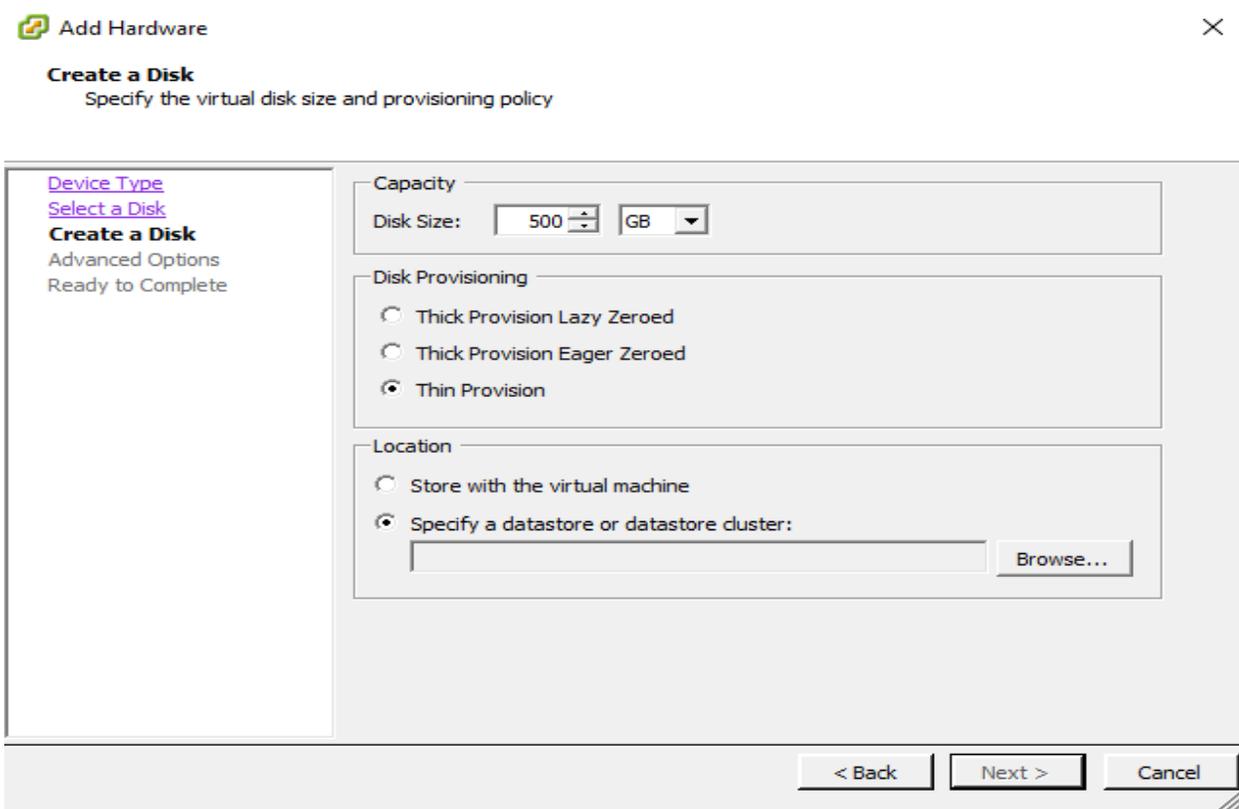
- Click finish.



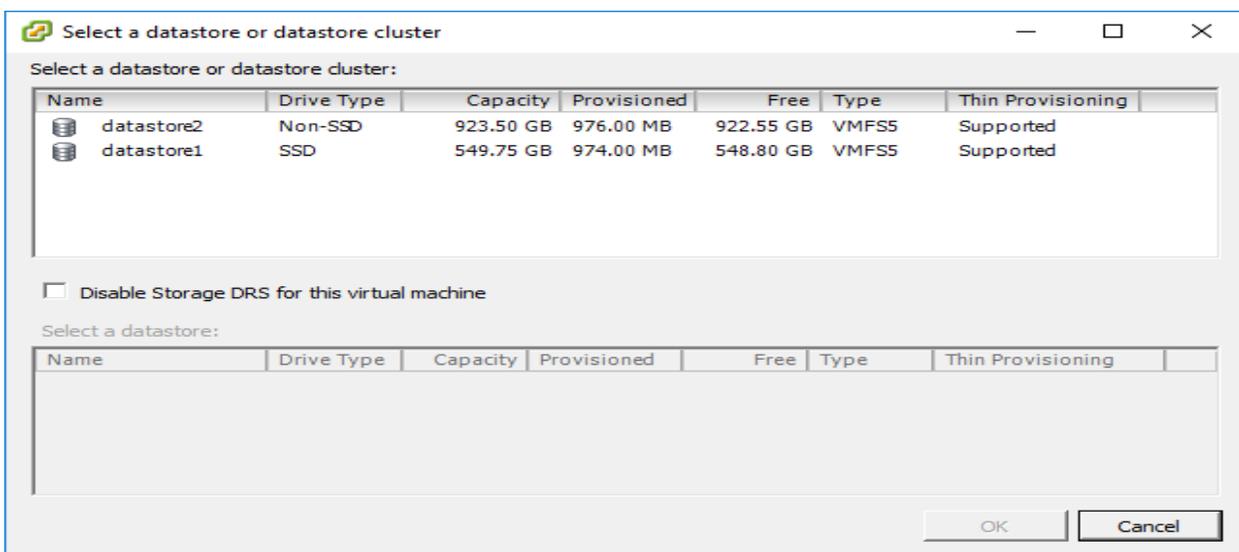
- Click on the add and select hard disk and click next and select an option to create a new virtual disk.



- Specify the disk size as 500GB and select thin provision and select “ Specify a datastore or datastore cluster ” and click browse.



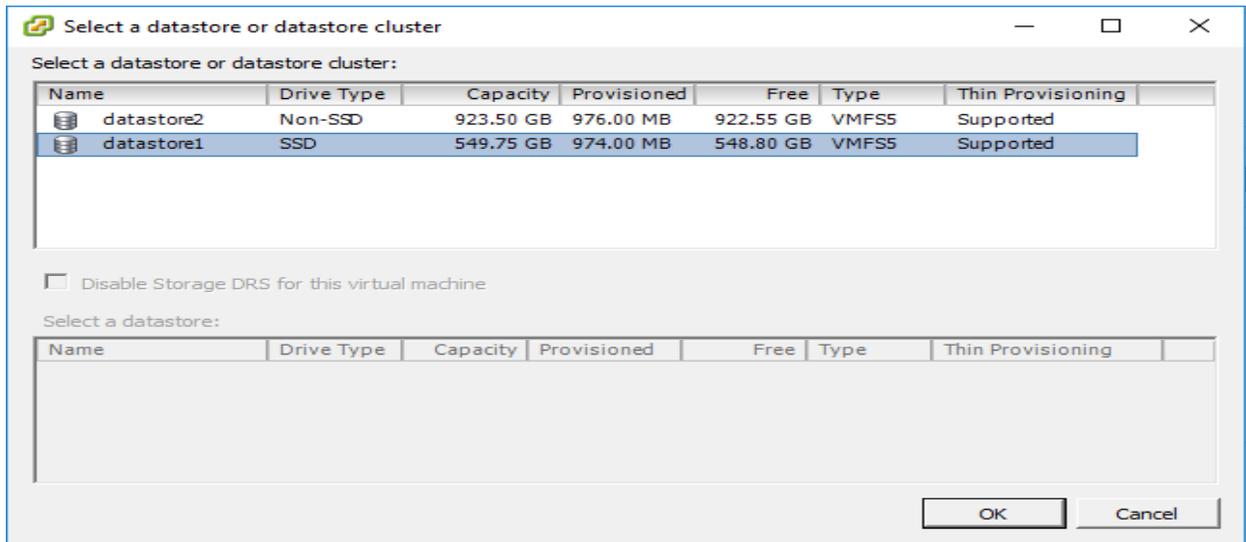
- Select a non-ssd datastore for the 500GB virtual disk and click ok.



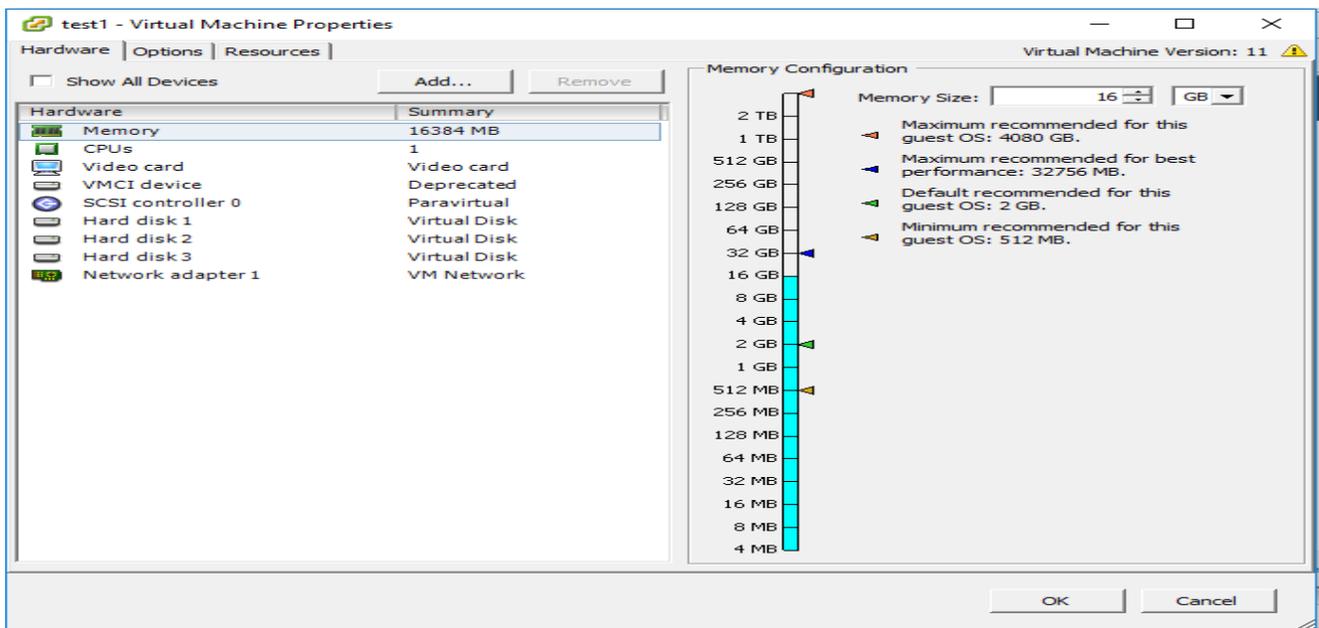
- Follow the same steps and create another virtual disk using following option
  - Disk size : 200GB

- Disk provisioning: thin provision
- Location: specify a datastore or a datastore cluster

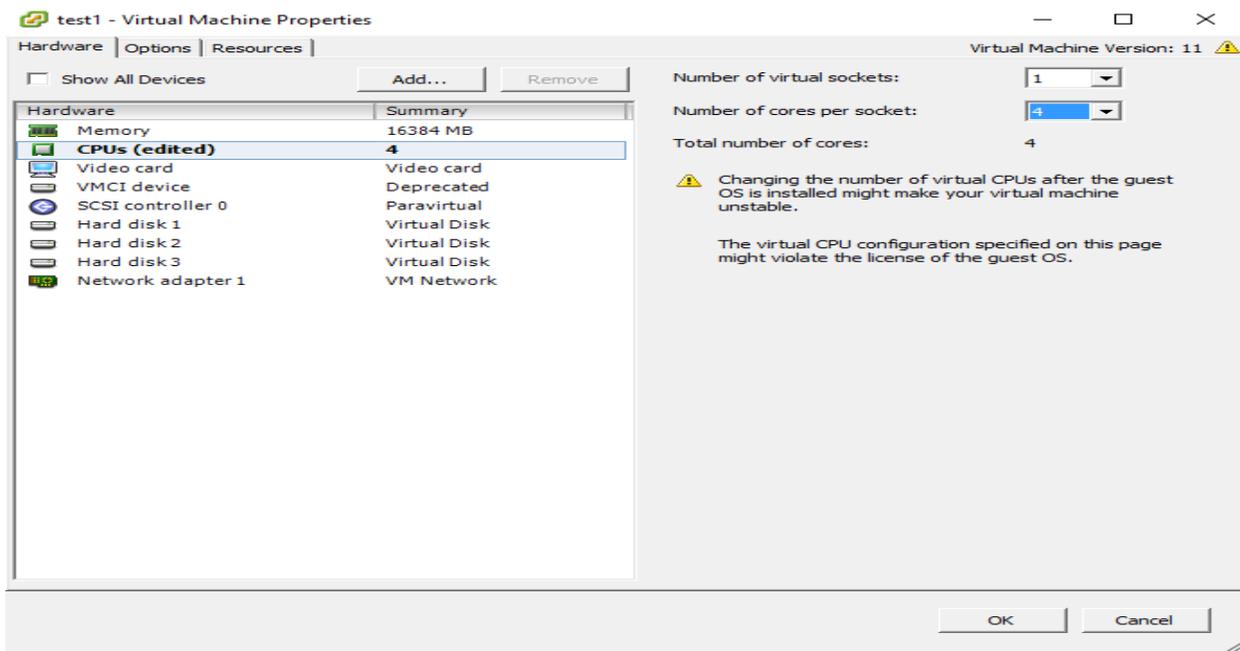
And select ssd as a datastore for the 200GB virtual disk and click ok and finish.



- Increase the memory to 16GB as show in the figure below.

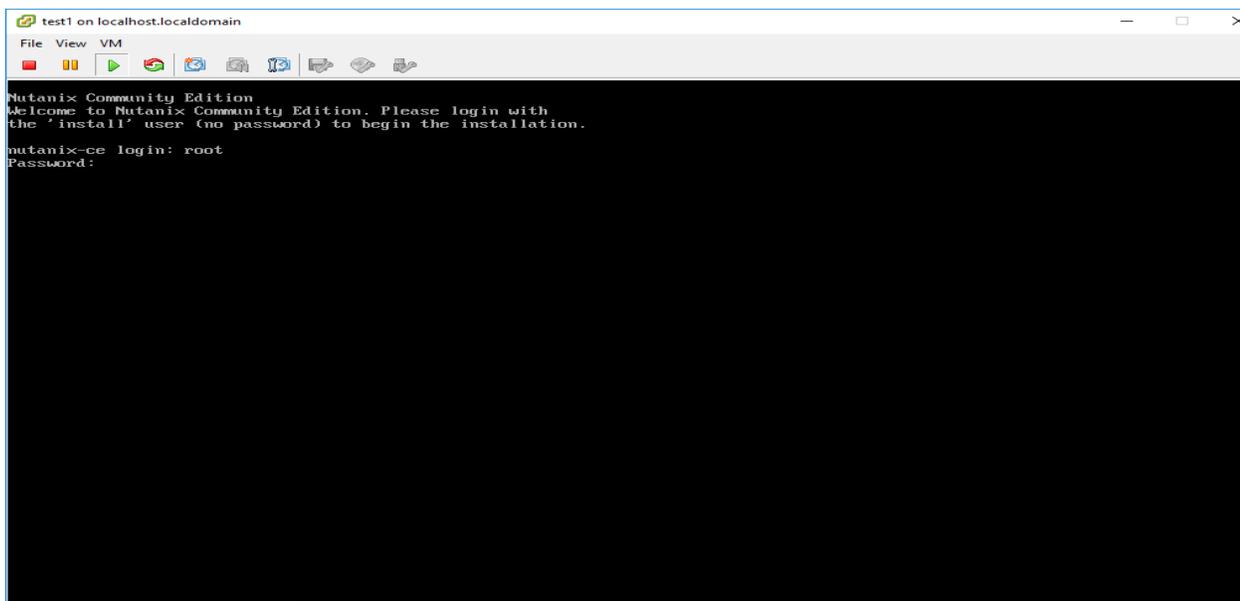


- Set the number of cores per socket to 4 as shown in the figure below and also select the SCSI controller 0 and click on change type to change its default type to paravirtual type.



## STEP 8: POWER ON THE VM AND INSTALL NUTANIX CE

Now the virtual machine is ready and go ahead and power on the virtual machine and launch the console window in vsphere. Then you will see the below screen.



- Login as a root user and the password as nutanix/4u

- Browse to the cd /home/install/phx\_iso/phenix as the below image.

```

test1 on localhost.localdomain
File View VM
Nutanix Community Edition
Welcome to Nutanix Community Edition. Please login with
the 'install' user (no password) to begin the installation.

nutanix-ce login: root
Password:
Last login: Wed May 13 21:08:36 on tty1
[root@nutanix-ce ~]# cd /home/install/phx_iso/
[root@nutanix-ce phx_iso]# ls
bin                dev                do_factory_installer.sh  etc                lib                mnt                reboot_quick        sys
ce_installer       dhcp.sh            do_installer.sh          init               livecd.sh          phoenix            rescue_shell        tmp
CMakeLists.txt    do_ce_installer.sh do_rescue_shell.sh       installer          Makefile           proc              sbin
[root@nutanix-ce phx_iso]# cd phoenix/
[root@nutanix-ce phoenix]# ls
arizona.py         consts.py          gui.py              kvm.py              log.py              patch_phoenix.py   sysUtil.py
ce_eula.txt        dell_factory.py   hel.json            layout              Makefile            phoenix            test
CMakeLists.txt    esx.py            hel.py              layout_tools.py    megadisk_inventory.sh  svm.py             util
common             firstboot         hyperv.py           legacy              minimum_reqs.py      svm_template
config             get_cvm_ip.py     imagingUtil.py     locale_select.py   partitions.py        system_info.py
[root@nutanix-ce phoenix]# _

```

- Now that you have accessed the directory now all you need to do some changes within the directory for your Nutanix CE to install successfully without any error to do that first type in the command **/home/install/phx\_iso/phenix/minimum\_reqs.py** this will allow you to enter inside the code and do some hardcode editing to fit in our current requirement.
- Ones you have entered inside the code drag down the panel and down you will find the below line of code which you need to comment by using # symbol in front it as shown in the image below.

```

try:
    # If the node has a bad disk the dmesg buffer can get flooded with sense
    # errors. Before disk interaction livecd.sh will save the output in a
    # temporary file which will be available to read from during the UT-d check.
    with open("/tmp/dmesg_out") as dm:
        dmesg = dm.read()
except IOError:
    ret, dmesg, err = sysUtil.shell_cmd(["dmesg"])

checkMemory(meminfo)
checkCores(cpuinfo)
## checkUtx(cpuinfo)
## checkIsIntel(cpuinfo)

if os.environ.has_key('COMMUNITY_EDITION'):
    boot_disk = sysUtil.find_boot_disk(param_list)
    if not boot_disk:
        raise StandardError('Unable to detect boot device from /proc/mounts.')
    if not param_list or param_list.svm_install_type == 'clean':
        CE_checkDisks(boot_dev=boot_disk.dev)
else:
    checkUtd(dmesg)
    checkSSD()
    checkDiskModels(param_list)
    checkInstallationDiskSize(param_list)

```

- 
-

- Remember the minimum SSD requirement for installing Nutanix is greater than or equal to 200, so the Nutanix CE installation process will throw error if you do not have minimum 200gb SSD.
- You need to lower the IOPS thresholds (SSD\_rdIOPS\_thresh and SSD\_wrIOPS\_thresh from 5000 to 1000 as shown below) to do that type in the below command **/home/install/phx\_iso/phoenix/sysUtil.py**.
- Set the SSD\_rdIOPS\_thresh = 1000
- Set the SSD\_wrIOPS\_thresh = 1000

```
# Contains custom SUM resource settings per model. Defaults will be used
# if a particular model does not exist in this structure.
SUM_GB_RAM = 16
SUM_NUM_UCPUS = 8
svm_resource_map = {
    'NX-1020': { 'ucpus': 4,          'ram': 12 },
    'NX-1050': { 'ucpus': 8,          'ram': SUM_GB_RAM },
    'NX-4170': { 'ucpus': 16,         'ram': 64 },
    'NX-8150': { 'ucpus': SUM_NUM_UCPUS, 'ram': 32 },
    'NX-9040': { 'ucpus': SUM_NUM_UCPUS, 'ram': 32 },
}

RE_WWN_SEARCH = re.compile(r"(\Logical Unit id: +0x(\w+)$)!"
                             "(\LU WWN Device Id: +(\w+)?$)", re.MULTILINE)
RE_SERIAL_SEARCH = re.compile(r"(\Serial InNumber: +(\w+)$)", re.MULTILINE)
RE_MODEL_SEARCH = re.compile(r"(\Device Model: +(\w+)$)", re.MULTILINE)
RE_FIRMWARE_SEARCH = re.compile(r"(\Revision: +(\w+)$)", re.MULTILINE)

perf_cache = {}
SSD_rdIOPS_thresh = 1000
SSD_wrIOPS_thresh = 1000

def shell_cmd(cmd, wait=True, fatal=True, ttyout=False):
    """
    This function will execute a command on a spawned
    shell process and will check return status if set
    and will dump output to screen (ttyout) if set.
    """
```

- Exit from the root user and login as an install user to start installation process

```
Nutanix Community Edition
Welcome to Nutanix Community Edition. Please login with
the 'install' user (no password) to begin the installation.

Hint: Num Lock on

nutanix-ce login: install
```

- Leave it to default and use tab to proceed to next step.

```
<< Nutanix Community Edition Installer >>

Please select your keyboard layout from the following list.

*****Keyboard Layout*****
| ua-utf-us          |
| ua-us             |
| uk                |
| unicode           |
|* us *             |
| us-acentos        |
| us-alt-intl       |
| us-altgr-intl     |
*****
Cancel   Proceed
```

- Leave it to default and use tab to proceed to next step.

```
<< Nutanix Community Edition Installer >>

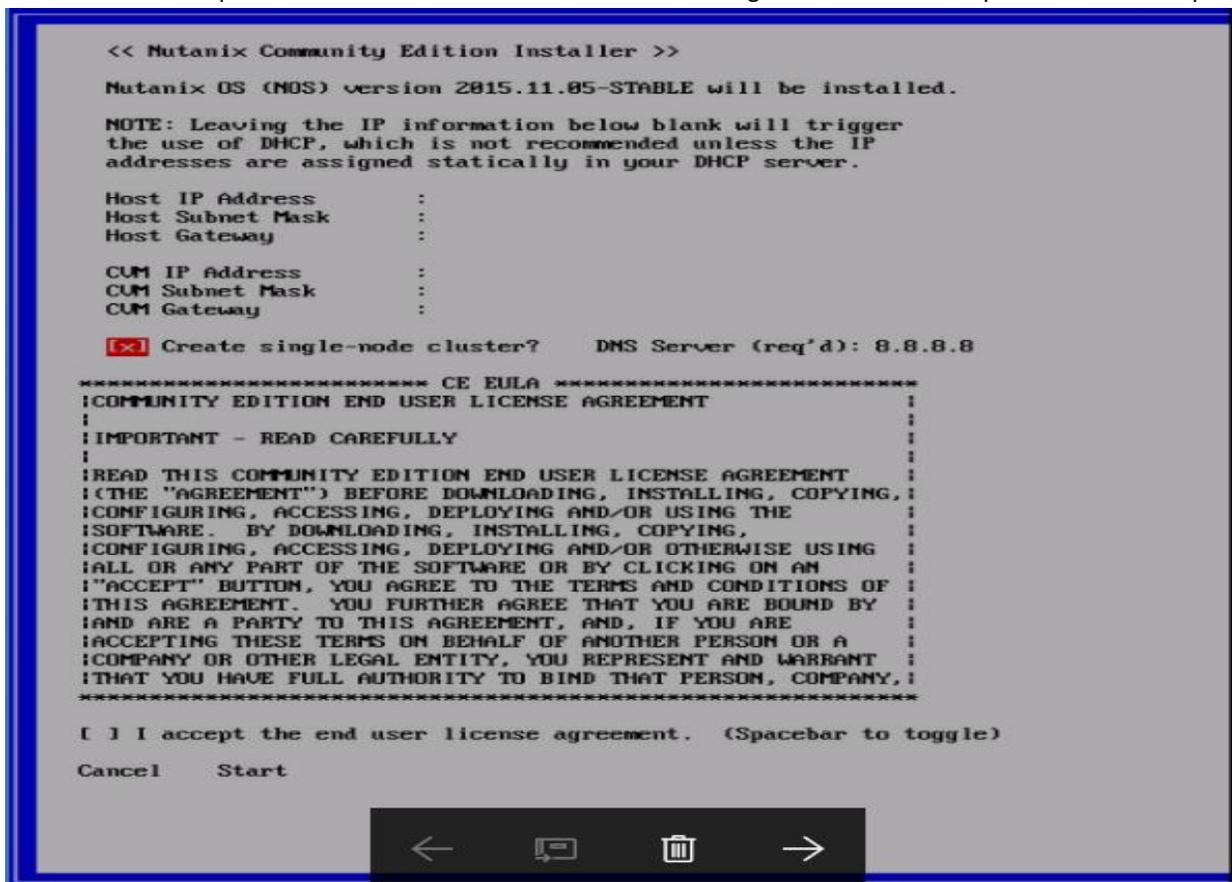
WARNING: Destructive IO tests will be run on the following disks in order to confirm acceptable performance.
If the disks listed below still have any data on them, please cancel and backup your data first.

***** Disks *****
| sdb: Model [Virtual disk], Size [536.87] GB, Serial [None] |
| sdc: Model [Virtual disk], Size [161.06] GB, Serial [None] |
*****
Cancel   Proceed
```

- Please wait for 10-30 minutes to install Nutanix

INFO: Executing performance tests on disks. Please expect a 10-30 second delay per device.

- Enter the ip address and subnet mask and check “create single node cluster” and proceed to next step.



- After successful installation of Nutanix you will see the screen below and press enter to continue

## Configuring a Nutanix CE cluster

Although it is possible to configure a cluster from the main cluster\_init page:

[http://CVM-IP-ADDRESS:2100/cluster\\_init.html](http://CVM-IP-ADDRESS:2100/cluster_init.html)

Some operations cannot be completed through this cluster\_init interface, or from the Create single-node cluster option during the install, namely creating a cluster where the host IP Address and CVM IP Address are on the same subnet (such as 192.168.x.x):

Although we will use the command-line method to configure the cluster in this guide, an example of the cluster\_init interface is shown below:

The screenshot shows the Nutanix cluster configuration interface. It includes the following fields and tables:

Cluster Name:

Cluster External IP:

Cluster Max Redundancy Factor:

CVM DNS Servers:

CVM NTP Servers:

Hypervisor DNS Servers:

Hypervisor NTP Servers:

Cluster-wide Network settings:

	Controller	Hypervisor	IPMI
Subnet Mask	255.255.255.0	255.255.255.0	
Default Gateway	192.168.0.254	192.168.0.254	

Discovered nodes:

Select	Node	Controller IP	Hypervisor IP	IPMI IP
<input checked="" type="checkbox"/>	1c864377/A	192.168.0.111	192.168.0.11	
<input checked="" type="checkbox"/>	4e4834a6/A	192.168.0.112	192.168.0.12	
<input checked="" type="checkbox"/>	39da2ce6/A	192.168.0.113	192.168.0.13	

### INSTALL AN SSH CLIENT:

To access the CVM host over SSH you will need an SSH client:

- SSH – Under Linux
- SSH – Under MacOS
- Putty – Windows

### SSH under Linux

ssh CVM-IP-ADDRESS (Where CVM-IP-ADDRESS is the IP Address of the CVM)

## Single-node and Multi-node clusters

Due to the way Nutanix CE protects and distributes data across hosts, only the following cluster sizes are available:

- Single-node cluster 1 Host no data redundancy
- Multi-node cluster 2 hosts not available
- Multi-node cluster 3 hosts with data redundancy
- Multi-node cluster 4 hosts with data redundancy

## CREATE A SINGLE-NODE CLUSTER

To create a single-node cluster, SSH into the host and provide the following user information:

User: nutanix

Password: nutanix/4u

Then run the following commands:

```
cluster -s CVM-IP -f create
```

```
ncli cluster add-to-name-servers servers="DNS-SERVER"
```

ncli cluster get-name-servers (Where CVM-IP is the IP Address of the CVM, and DNS-SERVER is a single IP Address or comma separated list of IP Addresses of DNS Servers)

NOTE: If DNS servers are not configured correctly at this stage registering the installation against your NEXT account in the next section will fail.

## CREATE A MULTI-NODE CLUSTER

To create a multi-node cluster, SSH into the host and provide the following user information:

User: nutanix

Password: nutanix/4u

Then run the following commands:

```
cluster -s CVM-IP -f create
```

```
ncli cluster add-to-name-servers servers="DNS-SERVER"
```

ncli cluster get-name-servers (Where CVM-IP is comma separated list of IP Addresses of the CVM's on each of the host you wish to include in the cluster, and DNS-SERVER is a single IP Address or comma separated list of IP Addresses of DNS Servers)

NOTE: If DNS servers are not configured correctly at this stage registering the installation against your NEXT account in the next section will fail.

```
INFO: Copying firstboot scripts into /var/stage/firstboot
INFO: Copying SSH keys
INFO: Installing firstboot marker file
INFO: Imaging thread 'hypervisor' has completed successfully
INFO: Imaging thread 'sdci: WRITE SAME failed. Manually zeroing.
INFO: Creating layout file for CommunityEdition in position A
INFO: Injecting post-cluster create settings into CUM
INFO: Copying diagnostic UM into SUM
INFO: Copying diagnostic UM into SUM
INFO: Imaging of SUM has completed successfully!
INFO: Imaging thread 'svm' has completed successfully
INFO: Cleaning up
INFO: Imaging process completed successfully!

Updating the iniramfs... done.

Waiting for the Nutanix Controller VM to start up...
[ 706.992029] kvm [98121]: vcpu1 unhandled wrmsr: 0x38d data 0
[ 707.022068] kvm [98121]: vcpu0 unhandled wrmsr: 0x38d data 0
[ 707.023652] kvm [98121]: vcpu1 unhandled wrmsr: 0x38d data 0
[ 707.113974] kvm [98121]: vcpu0 unhandled rdmsr: 0x38d
[ 707.117984] kvm [98121]: vcpu1 unhandled rdmsr: 0x38d
[ 711.464133] kvm [98121]: vcpu0 unhandled rdmsr: 0x38d
[ 710.662716] kvm [98121]: vcpu0 unhandled rdmsr: 0x38d
[ 718.664874] kvm [98121]: vcpu0 unhandled rdmsr: 0x38d
[ 718.703040] kvm [98121]: vcpu0 unhandled rdmsr: 0xc9
[ 718.704655] kvm [98121]: vcpu0 unhandled rdmsr: 0x1a6
[ 718.705076] kvm [98121]: vcpu0 unhandled rdmsr: 0x3f6
[ 718.857626] kvm [98121]: vcpu1 unhandled rdmsr: 0x38d
[ 718.859579] kvm [98121]: vcpu1 unhandled rdmsr: 0x38d
.....[ 993.091079] sdb: WRITE SAME failed. Manually zeroing.
.....

Success! Press <Enter> to return the login prompt.
Please refer to the documentation for next steps.
In case of a clean install, the login banner will
show the IP address of the CUM.
In case of a repair, the CUM will take a couple more
minutes to come up.
```

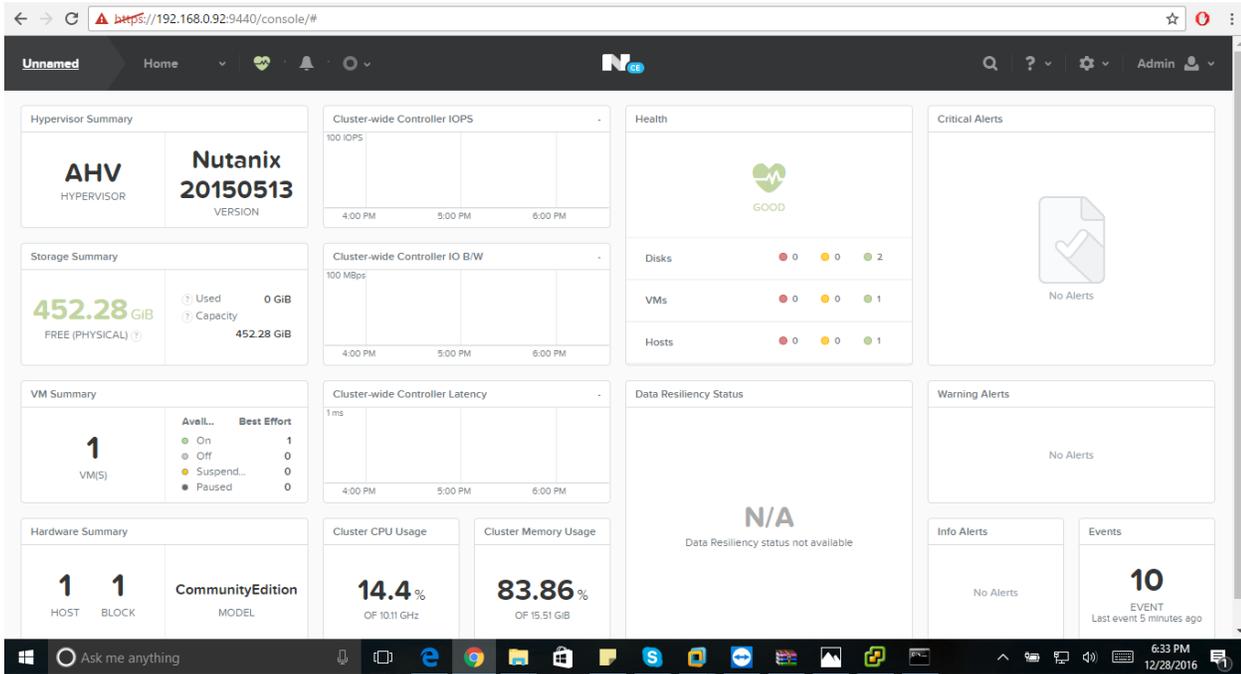
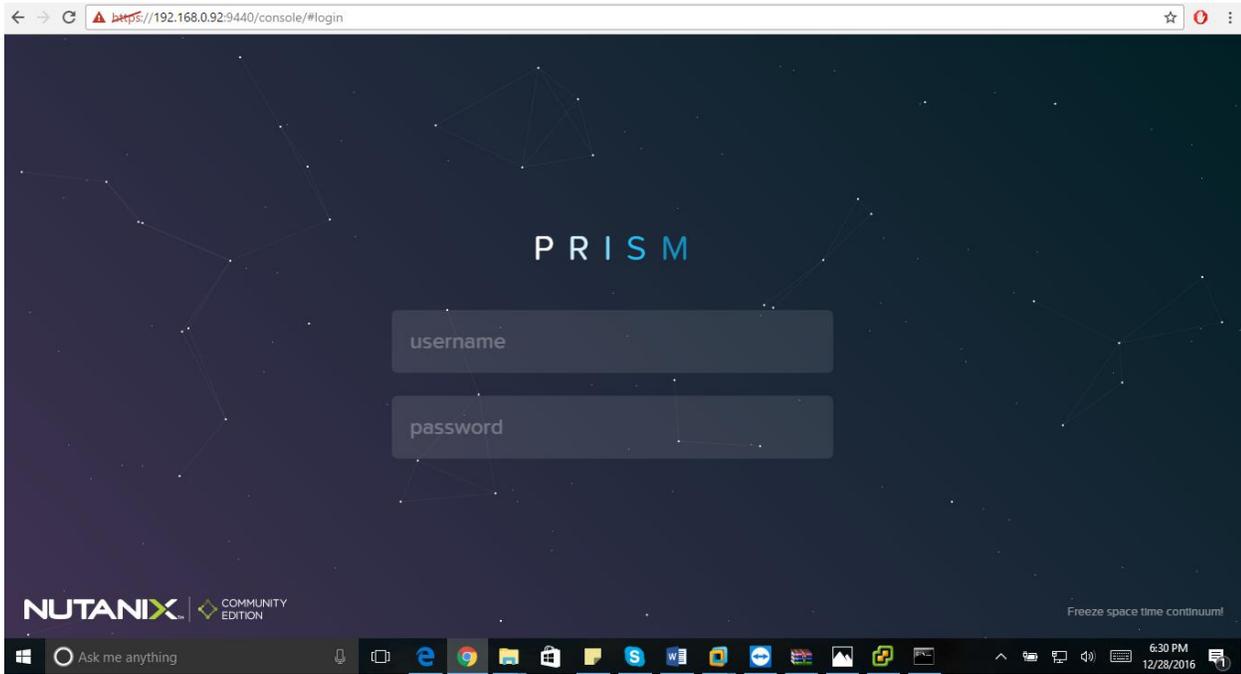
➤ Use CVM IP to login to the prism

```
Nutanix Community Edition
Nutanix CUM IP: 192.168.0.92

Hint: Num Lock on

NTNX-b0217637-A login: _
```

➤ Open the web browser and enter the CVM IP and prism will load and use admin as username and admin as password.



## Starting and stopping services

### STARTING

The correct Start order is HOST > CVM > VMs

To check the status of a Nutanix CE cluster, login to a CVM as

user: nuatnix,

password: nutanix/4u

and type the following:

cluster status

To start a Nutanix CE cluster, login to a CVM as

user: nuatnix,

password:nutanix/4u

and type the following:

cluster start

To start a Nutanix CE VM, go to the VM Table under the VM menu in PRISM and select a Power option from the list.

### STOPPING

The correct Stop order is VMs > CVM > Host

To stop a Nutanix CE VM, go to the VM Table under the VM menu in PRISM and select a Power option from the list.

To check the status of a Nutanix CE cluster, login to a CVM as user: nuatnix,  
password: nutanix/4u and type the following:

cluster status

To stop a Nutanix CE cluster, login to a CVM as

user: nuatnix,

password: nutanix/4u

and type the following:

`cluster stop`

To stop a Nutanix CE CVM, login to a CVM as

`user: nutanix,`

`password:nutanix/4u`

and type the following:

`sudo shutdown -h now`

To stop a Nutanix CE Host, login to a CVM as user: root, password: nutanix/4u

and type the following:

`shutdown -h now`

## POWERING OFF A WHOLE CLUSTER

Complete shutdown

The correct Shutdown order for a whole Nutanix CE cluster is as follows:

- Shutdown and power off each individual VM through PRISM
- Log into any CVM as “nutanix” and issue a “cluster stop”, followed by a

“cluster status” (just to check that everything did shutdown cleanly)

- Log into each CVM individually as “nutanix” and issue a “sudo shutdown -h now”

- Log into each Host individually as “root” and issue a “sudo shutdown -h now”

NOTE: Each Host has its own IP Address, plus an additional IP Address for the CVM that runs on each node.

NOTE: If the commands do not work as expected be sure to check that you are logging onto the correct IP Address and as the correct user (root for the physical host and nutanix for the CVM virtual machine).