

Veritas Cluster Server training.

I shall not be liable for any damages that you may incur as a result of attempting anything in this training program.

I am responsible for any errors in this document and I apologize in advance for this. I am still in the learning mode and will try and make corrections as things come up.

When in DOUBT please google and search for articles on Symantec's site, for instance I did not know that one does not have to assign IP addresses to the private interconnects, I stood corrected after reading:

<http://www.symantec.com/connect/forums/configure-llt-and-gab-veritas-cluster>

Please take all careful precautions while handling electrical equipment. Some equipment may be physically heavy and/or cause you distress in some ways.

Please look up the specifications(physical/electrical/environmental/chemical etc.) of any hardware you use and take requisite precautions.

The training notes that follow are a kind of guide for you to refer back to. I intend on making this an interactive session, so please try to follow the video, you can always go back to the notes and rewind the video as well.

Veritas is a trademark of Symantec corp and/or of Veritas. RHEL is a trademark of Red Hat. Oracle is a trademark of Oracle. There are other various trademarks no attempt is made to infringe. Apology as necessary is made in advance.

This guide may prepare you to work with the daily routine tasks of veritas cluster server on RHEL. This guide may be useful for the following tasks:

- 1)Understand the meaning of clustering and High Availability.
- 2)Understand the components of Veritas Cluster server.
- 3)Installing Veritas Cluster Server
- 4)Installing Oracle Database as a highly available service (this means it will run on one node if another node fails).
- 5)Basic maintenance of a Veritas Cluster.
- 6)Troubleshooting
- 7)Interview questions.

This guide is best followed by actually doing all the steps presented herein.

I am using a HP DL580 G5 with 4 quad core processors and 32GB of RAM. It takes SFF(small form factor) SAS drives.

I have in the past successfully used an i7 PC with windows operating system(you may use RHEL), 8GB of Ram(this is the very least recommended) and about 350GB of disk space. A network interface is required. An internet connection is required. Free accounts at Symantec(for downloading the trial version of Veritas Cluster Server – SFHA – Storage Foundation High Availability) and at Oracle(trial edition of Oracle Database Enterprise version) are essential.

In Virtualbox(free virtualization technology that works on both windows and on Linux) I am giving 2048MB to both VM's and 1024MB to the virtualised storage(openfiler). I am designating one cpu each for all three. So, for the i7 case, on an 8GB physical host this leaves  $8-2-1=3$ GB for the underlying operating system. And, leaves,  $4-1(\text{openfiler})-1(\text{node1})-1(\text{node2})=1$  core(cpu) for the base physical system hosting everything of the 4core cpu i7 processor.

Where to buy this hardware: I suggest googling: HP DL 580 G5 and i7 16GB and checking out Ebay.


Down to it.


What is a cluster:

A computer cluster consists of a set of loosely or tightly connected computers(here called nodes) that work together so that, in many respects, they can be viewed as a single system. Computer clusters have each node set to perform the same task, controlled and scheduled by software(in our case by VCS).

The components of a cluster are usually connected to each other through fast local area networks ("LAN"), with each node (computer used as a server) running its own instance of an operating system. All of the nodes use IDENTICAL hardware and IDENTICAL operating system.

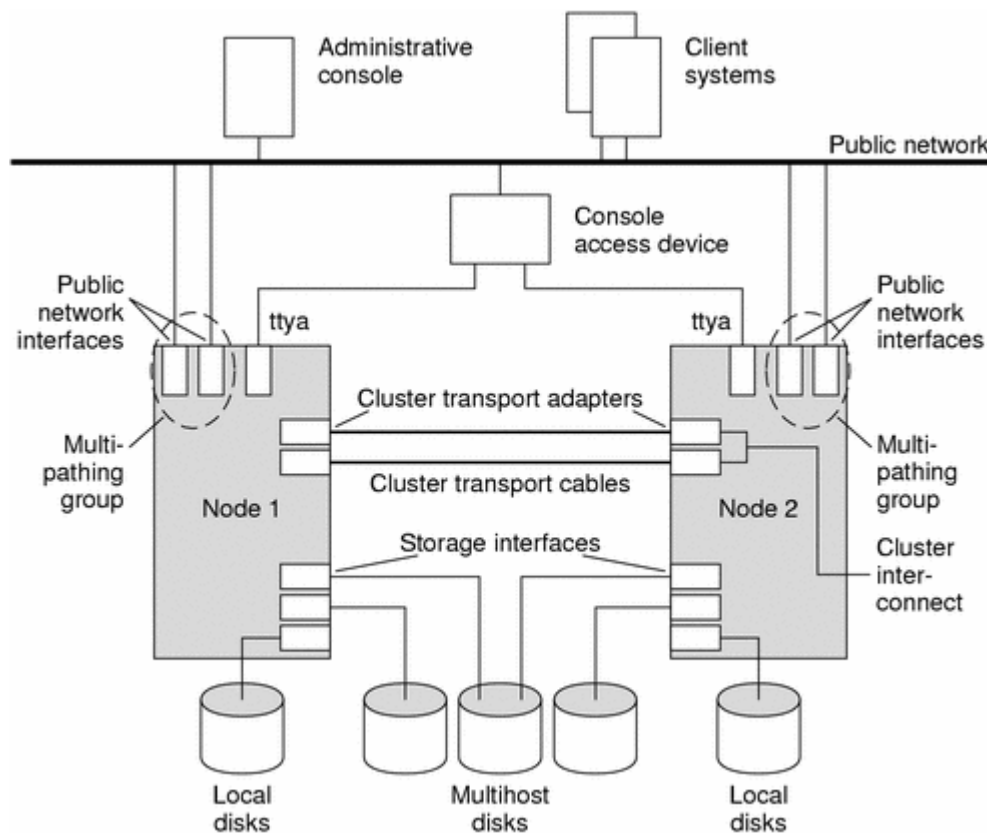
They are usually deployed to improve performance and availability over that of a single computer, while typically being much more cost-effective than single computers of comparable speed or availability.





#68709

This image below shows how more resilience is built into a cluster, this type of cluster is what you are likely to encounter at work. Please NOTE the presence on dual NIC's for the public interface, they are multipathed or bonded together, so that if one NIC on the public network goes down the other takes over.



Also please note the presence of three storage interfaces per node , one is going to the local disk, two on each node are going to the shared storage, again this is multipathing of the storage adapters, commonly called HBA's(host based adapter), again this can be accomplished in a variety of ways, you might have heard of powerpath/vxdmp/linux native multipathing etc.

A cluster typically has two or more nodes(computers) connected to each other by TWO networks. A high speed private network which usually consists of two links (two network cards dedicated for the private interconnect) and a second network which is the public network via which clients(you and me and application users) can access the cluster(for them it is a single computer).

The other facet to a cluster is the requirement of shared storage. Shared storage are disk(s) which are connected to all the nodes of the cluster.

The premise is not that complicated. The premise of clustering being that if one node of the cluster goes down/crashes the other node picks up. To this end both nodes of the cluster need to be aware if another one crashes, this is accomplished usually by the presence of the private network of high speed interconnect. When one node crashes the application hosted on the cluster “fails over” to the other node of the cluster. This is HIGH AVAILABILITY.



You might have noticed how ebay/google/amazon etc do not crash, this is again because behind the ebay.com/google.com/amazon.com is a cluster of two or more computers which provides redundancy to the entire setup.

A cluster hence consists of hardware and the cluster software(which actually provides the clustering between the nodes of the cluster).

Types of clustering software:

There are several commercially available clustering softwares, the important ones are:

- a)Veritas Cluster Server
- b)Sun Clustering
- c)RHEL/CentOS native clustering

In this guide we will focus on Veritas Cluster Server.

The components of VCS:

Veritas cluster concepts

VCS Basics

### **VCS Basics**

A single VCS cluster consists of multiple systems connected in various combinations to shared storage devices. VCS monitors and controls applications running in the cluster, and restarts applications in response to a variety of hardware or software faults. Client applications continue operation with little or no downtime. Client workstations receive service over the public network from applications running on the VCS systems. VCS monitors the systems and their services. VCS systems in the cluster communicate over a private network.

### **Switchover and Failover**

A switchover is an orderly shutdown of an application and its supporting resources on one server and a controlled startup on another server.

A failover is similar to a switchover, except the ordered shutdown of applications on the original node may not be possible, so the services are started on another node. The process of starting the application on the node is identical in a failover or switchover.

### **CLUSTER COMPONENTS:**

#### **Resources**

Resources are hardware or software entities, such as disk groups and file systems, network interface cards (NIC), IP addresses, and applications. Controlling a resource means bringing it online (starting), taking it offline (stopping), and monitoring the resource.

#### **Resource Dependencies:**

Resource dependencies determine the order in which resources are brought online or taken offline when their associated service group is brought online or taken offline. In VCS terminology, resources

are categorized as *parents* or *children*. Child resources must be online before parent resources can be brought online, and parent resources must be taken offline before child resources can be taken offline.

### **Resource Categories:**

**On-Off:** VCS starts and stops On-Off resources as required. For example, VCS imports a disk group when required, and exports it when it is no longer needed.

**On-Only:** VCS starts On-Only resources, but does not stop them. For example, VCS requires NFS daemons to be running to export a file system. VCS starts the daemons if required, but does not stop them if the associated service group is taken offline.

**Persistent:** These resources cannot be brought online or taken offline. For example, a network interface card cannot be started or stopped, but it is required to configure an IP address. VCS monitors Persistent resources to ensure their status and operation. Failure of a Persistent resource triggers a service group failover.

### **Service Groups**

A *service group* is a logical grouping of resources and resource dependencies. It is a management unit that controls resource sets. A single node may host any number of service groups, each providing a discrete service to networked clients. Each service group is monitored and managed independently. Independent management enables a group to be failed over automatically or manually idled for administration or maintenance without necessarily affecting other service groups. VCS monitors each resource in a service group and, when a failure is detected, restarts that service group. This could mean restarting it locally or moving it to another node and then restarting it.

### **Types of Service Groups:**

#### **Fail-over Service Groups**

A failover service group runs on one system in the cluster at a time.

#### **Parallel Service Groups**

A parallel service group runs simultaneously on more than one system in the cluster.

#### **Hybrid Service Groups**

A hybrid service group is for replicated data clusters and is a combination of the two groups cited above. It behaves as a failover group *within* a system zone and a parallel group *across* system zones. It cannot fail over across system zones, and a switch operation on a hybrid group is allowed only if both systems are within the same system zone.

### **The ClusterService Group**

The Cluster Service group is a special purpose service group, which contains resources required by VCS components. The group contains resources for Cluster Manager (Web Console), Notification, and the wide-area connector (WAC) process used in global clusters.

The ClusterService group can fail over to any node despite restrictions such as “frozen.” It is the first service group to come online and cannot be auto disabled. The group comes online on the first node that goes in the running state.

### **Agents**

Agents are VCS processes that manage resources of predefined resource types according to commands received from the VCS engine, HAD. A system has one agent per resource type that monitors all resources of that type; for example, a single IP agent manages all IP resources.

When the agent is started, it obtains the necessary configuration information from VCS. It then periodically monitors the resources, and updates VCS with the resource status. VCS agents are multithreaded, meaning a single VCS agent monitors multiple resources of the same resource type on one host. VCS monitors resources when they are online *and* offline to ensure they are not started on systems on which they are not supposed to run. For this reason, VCS starts the agent for any resource configured to run on a system when the cluster is started. If no resources of a particular type are configured, the agent is not started.

### **Agent Operation:**

Online—Brings a specific resource ONLINE from an OFFLINE state.

Offline—Takes a resource from an ONLINE state to an OFFLINE state.

Monitor—Tests the status of a resource to determine if the resource is online or offline.

Clean—Cleans up after a resource fails to come online, fails to go offline, or fails while in an ONLINE state.

Action—Performs actions that can be completed in a short time (typically, a few seconds), and which are outside the scope of traditional activities such as online and offline.

Info—Retrieves specific information for an online resource.

### **Multiple Systems**

VCS runs in a replicated state on each system in the cluster. A private network enables the systems to share identical state information about all resources and to recognize which systems are active, which are joining or leaving the cluster, and which have failed. The private network requires two communication channels to guard against network partitions.

For the VCS private network, two types of channels are available for heartbeating: network connections and heartbeat regions on shared disks. The shared disk region heartbeat channel is used for heartbeating only, not for transmitting information as are network channels. Each cluster configuration requires at least two channels between systems, one of which *must* be a network connection. The remaining channels may be a combination of network connections and heartbeat regions on shared disks. This requirement for two channels protects your cluster against network partitioning. Also it's recommended to have at least one heart beat disk region on each I/O shared between systems. E.g. two-system VCS cluster in which sysA and sysB have two private network connections and another connection via the heartbeat disk region on one of the shared disks. If one of the network connections fails, two channels remain. If both network connections fail, the condition is in jeopardy, but connectivity remains via the heartbeat disk.

### **Shared Storage**

A VCS hardware configuration typically consists of multiple systems connected to share storage via I/O channels. Shared storage provides multiple systems an access path to the same data, and enables VCS to restart applications on alternate systems when a system fails.

### **Cluster Control, Communications, and Membership**

Cluster communications ensure VCS is continuously aware of the status of each system's service groups and resources.

### **High-Availability Daemon (HAD)**

The high-availability daemon, or HAD, is the main VCS daemon running on each system. It is responsible for building the running cluster configuration from the configuration files, distributing the information when new nodes join the cluster, responding to operator input, and taking corrective action when something fails. It is typically known as the VCS engine. The engine uses agents to monitor and manage resources. Information about resource states is collected from the agents on the local system and forwarded to all cluster members. HAD operates as a *replicated state machine* (RSM). This means HAD running on each node has a completely synchronized view of the resource status on each node. The RSM is maintained through the use of a purpose-built communications package consisting of the protocols *Low Latency Transport* (LLT) and *Group Membership Services/Atomic Broadcast* (GAB).

### **Low Latency Transport (LLT)**

VCS uses private network communications between cluster nodes for cluster maintenance. The Low Latency Transport functions as a high-performance, low-latency replacement for the IP stack, and is used for all cluster communications.

### **Traffic Distribution**

LLT distributes (load balances) internode communication across all available private network links. This distribution means that all cluster communications are evenly distributed across all private network links (maximum eight) for performance and fault resilience. If a link fails, traffic is redirected to the remaining links.

### **Heartbeat**

LLT is responsible for sending and receiving heartbeat traffic over network links. This heartbeat is used by the Group Membership Services function of GAB to determine cluster membership. The system administrator configures LLT by creating the configuration files `/etc/llthosts`, which lists all the systems in the cluster, and `/etc/llttab`, which describes the local system's private network links to the other systems in the cluster.

### **Group Membership Services/Atomic Broadcast (GAB)**

The Group Membership Services/Atomic Broadcast protocol (GAB) is responsible for cluster membership and cluster communications.

### **Cluster Membership**

GAB maintains cluster membership by receiving input on the status of the heartbeat from each node via LLT. When a system no longer receives heartbeats from a peer, it marks the peer as DOWN and excludes the peer from the cluster.

### **Cluster Communications**

GAB's second function is reliable cluster communications. GAB provides guaranteed delivery of point-to-point and broadcast messages to all nodes.

The system administrator configures GAB driver by creating a configuration file (`/etc/gabtab`).

Now we will describe how we will use freely available and not so freely available software for learning Veritas Cluster. We will essentially virtualise everything.

At the workplace, you will have two or more physical servers as nodes of the cluster. Each physical machine will be a node. For our training we will use two virtual machines running on Oracle's Virtualbox as the “two physical machines”.

At the workplace you will most likely have two NICs, bonded together, dedicated to one private interconnect. I apologize that this is not clear in any of the diagrams. So you will have 4 NIC's dedicated to 2 private interconnects per node. Again this is called NIC bonding. In our training we will not use bonding, rather we will assign two bridged interfaces for the two NIC's that serve the purpose of private interconnect.

At the workplace you will have two HBA's(SAN adapters) bonded together(to appear one), per node of the cluster connected to the SAN. This connection will be emulated here over the public NIC via the use of iSCSI technology. Hence we will for our training, virtualize the SAN(storage area network). We will use openfiler for this purpose.

For our two nodes we will use two RHEL 6.6 Virtual Machines. A virtual machine is a computer that is virtual. To host virtual computers we need a virtualization software and the hardware that supports it.

Again you may install Linux on your hardware and install Oracle VirtualBox on top of Linux or if you are using a PC with windows operating system, install virtualbox on it.

The first step is to enable hardware virtualization technology on your computer. Please try following these links, they may help:

<http://www.sysprobs.com/disable-enable-virtualization-technology-bios>

<http://www.tomshardware.com/answers/id-1715638/enable-hardware-virtualization-bios.html>

If you have difficulty with this step. Please consider joining a forum and asking for help. You will not be able to proceed in configuring anything till you go past this step.

The next step is to install the virtualization software over which our virtual computers henceforth known as virtual machines(VM's) will run on, there are many options, we will use the freely available Oracle Virtualbox, you may download it from:

<https://www.virtualbox.org/>

Please also obtain the RHEL 6.x iso image

Please also obtain the openfiler iso image. Why do we need openfiler: We will virtualize shared storage using it. Instead of buying expensive hardware to use for the shared storage we will instead go with the freely available virtualized storage application called openfiler. You may download openfiler from:  
<https://www.openfiler.com/community/download>

Please go ahead and install Oracle Virtual Box. After this is done, please follow the following steps:



Select New and choose the path of the RHEL 6.x iso image and add a total of three bridged network interfaces as shown in the below steps.

Create Virtual Machine✕



## Name and operating system

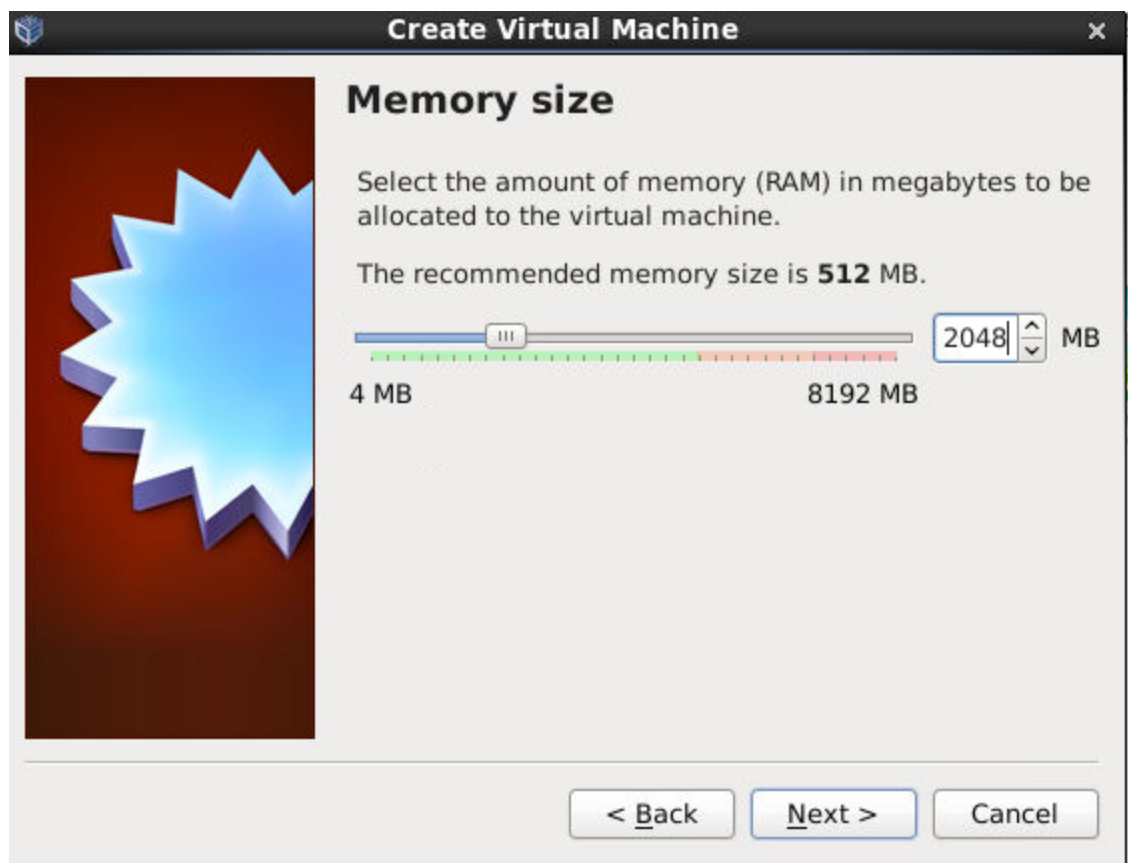
Please choose a descriptive name for the new virtual machine and select the type of operating system you intend to install on it. The name you choose will be used throughout VirtualBox to identify this machine.

Name:

Type:  

Version:

Give it a name, say, node1



Giving it 2048MB of RAM





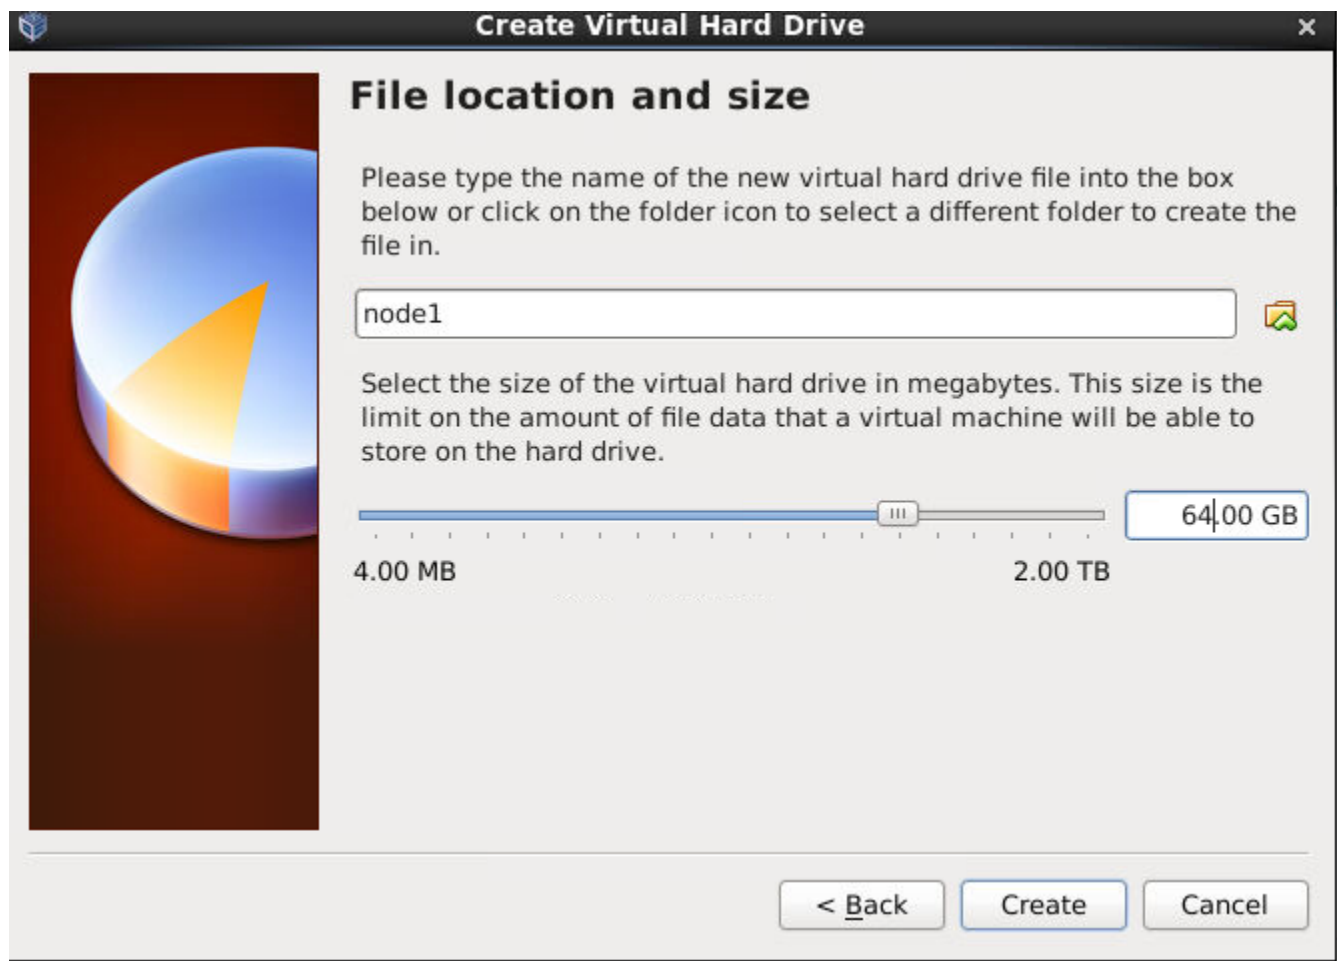
click "Create"



Click "Next"

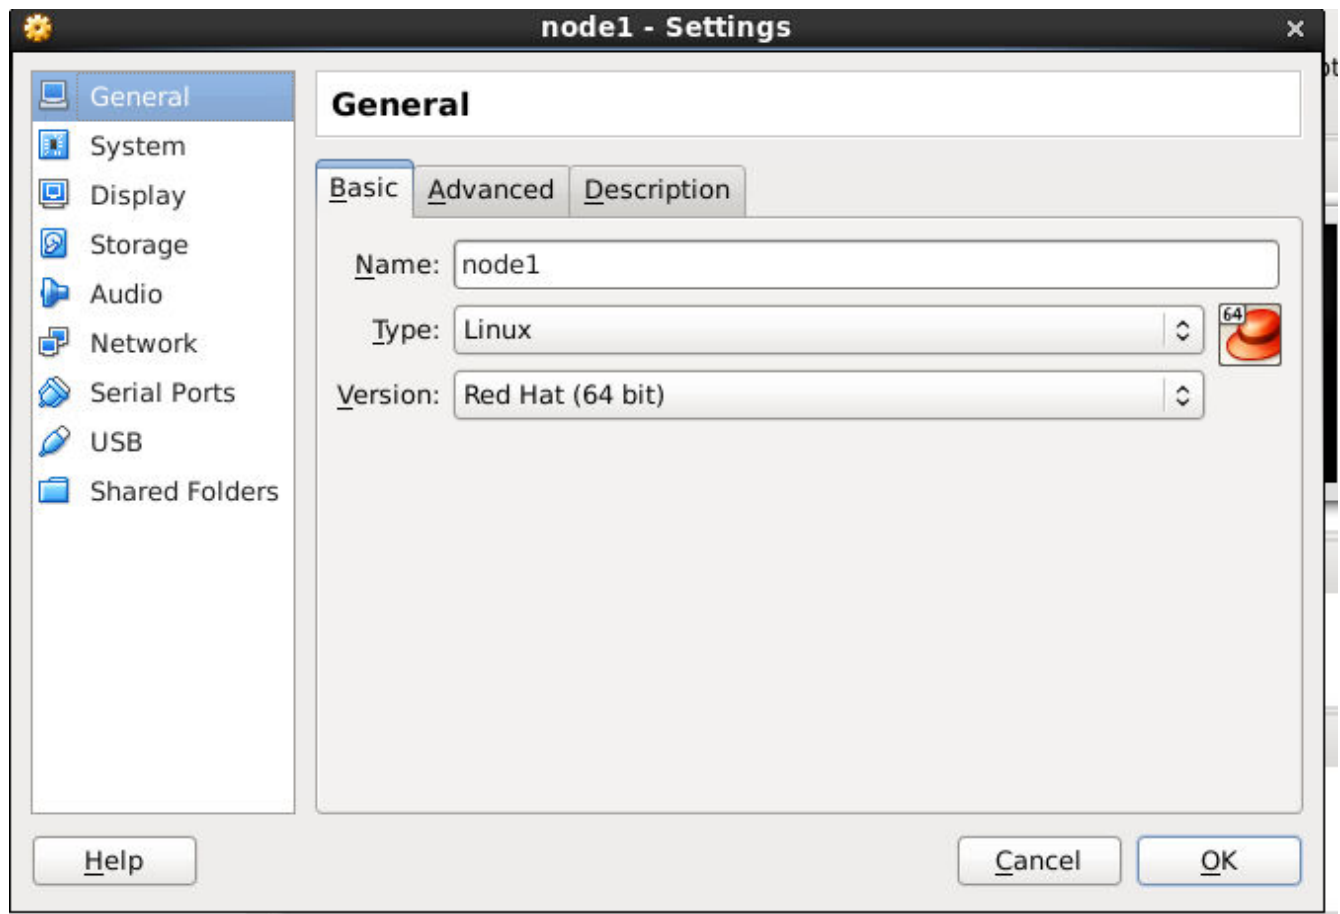


Click “Next”

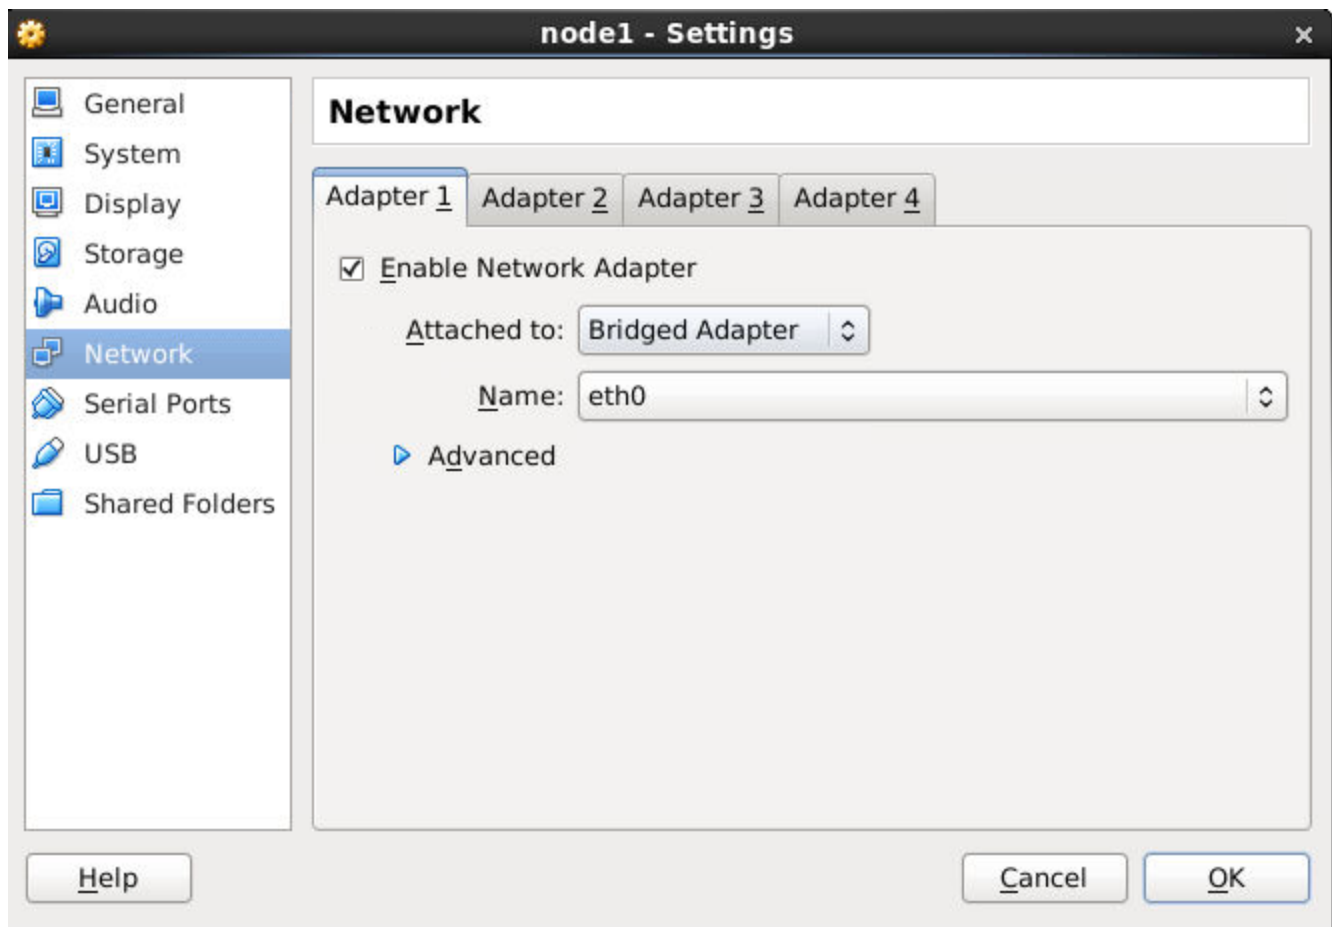


Make it a 64GB disk

Click settings

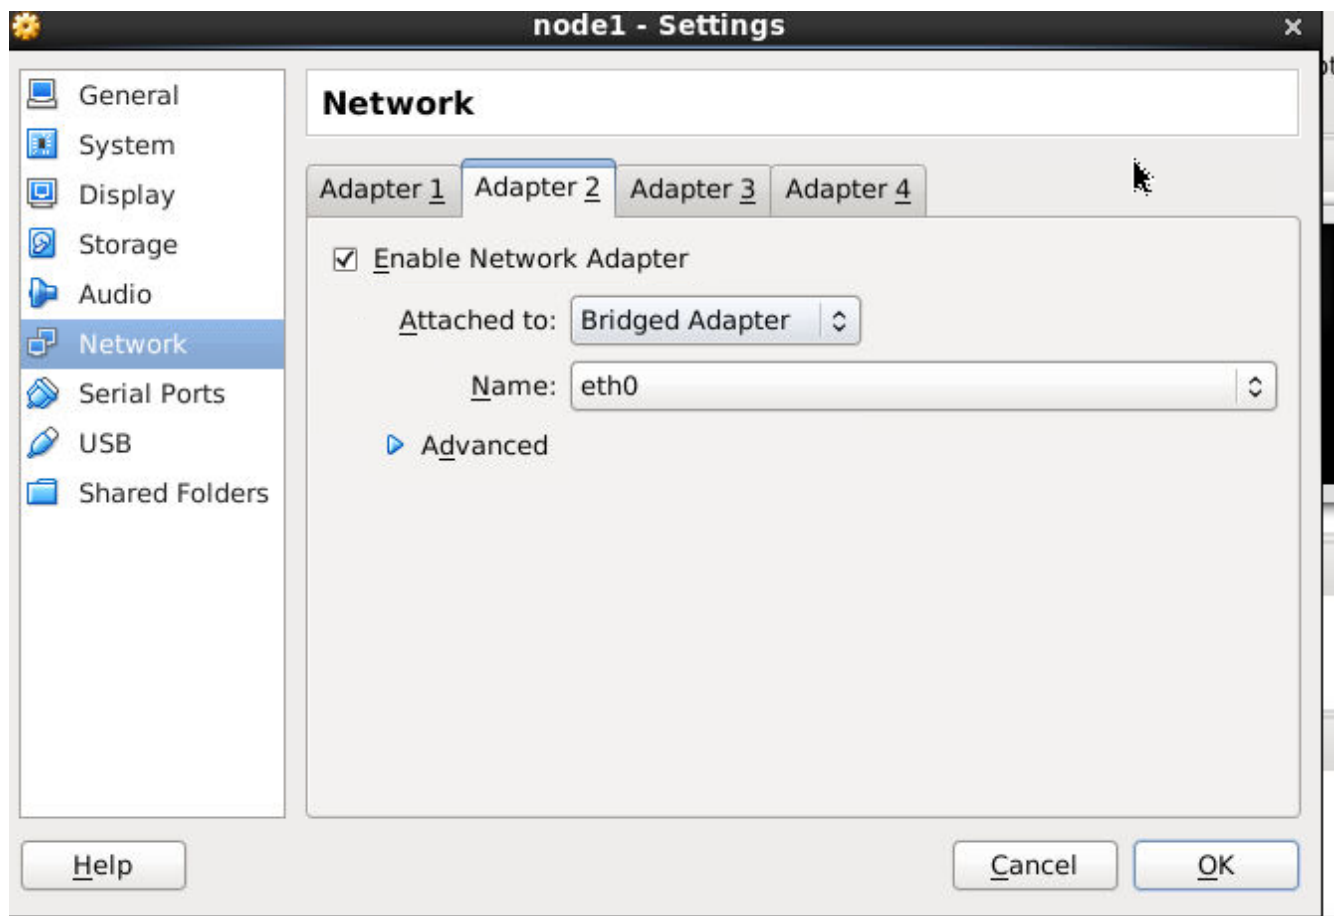


Click Network

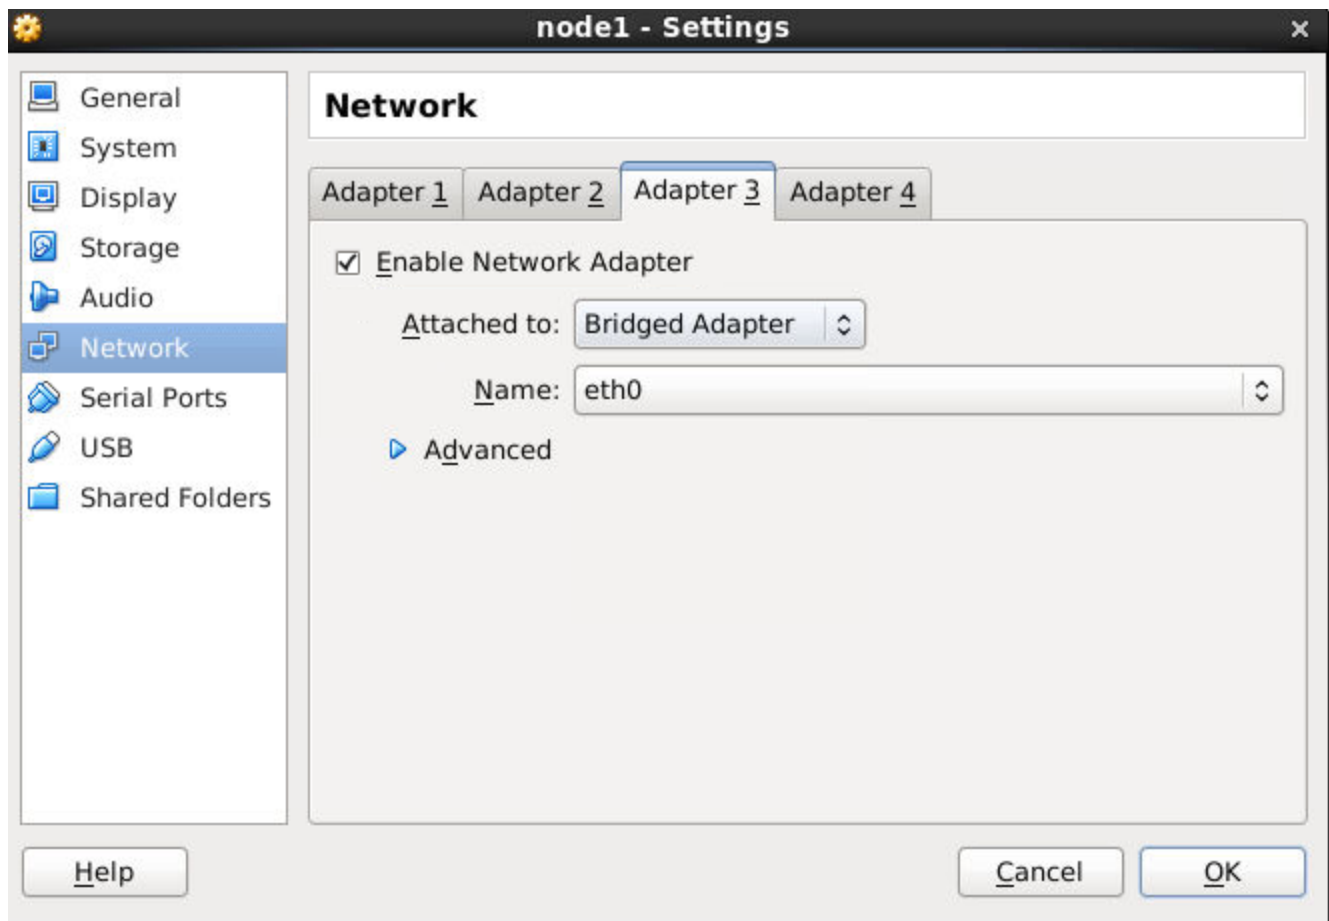


Choose "Bridged Adapter"

Select "Adapter2"



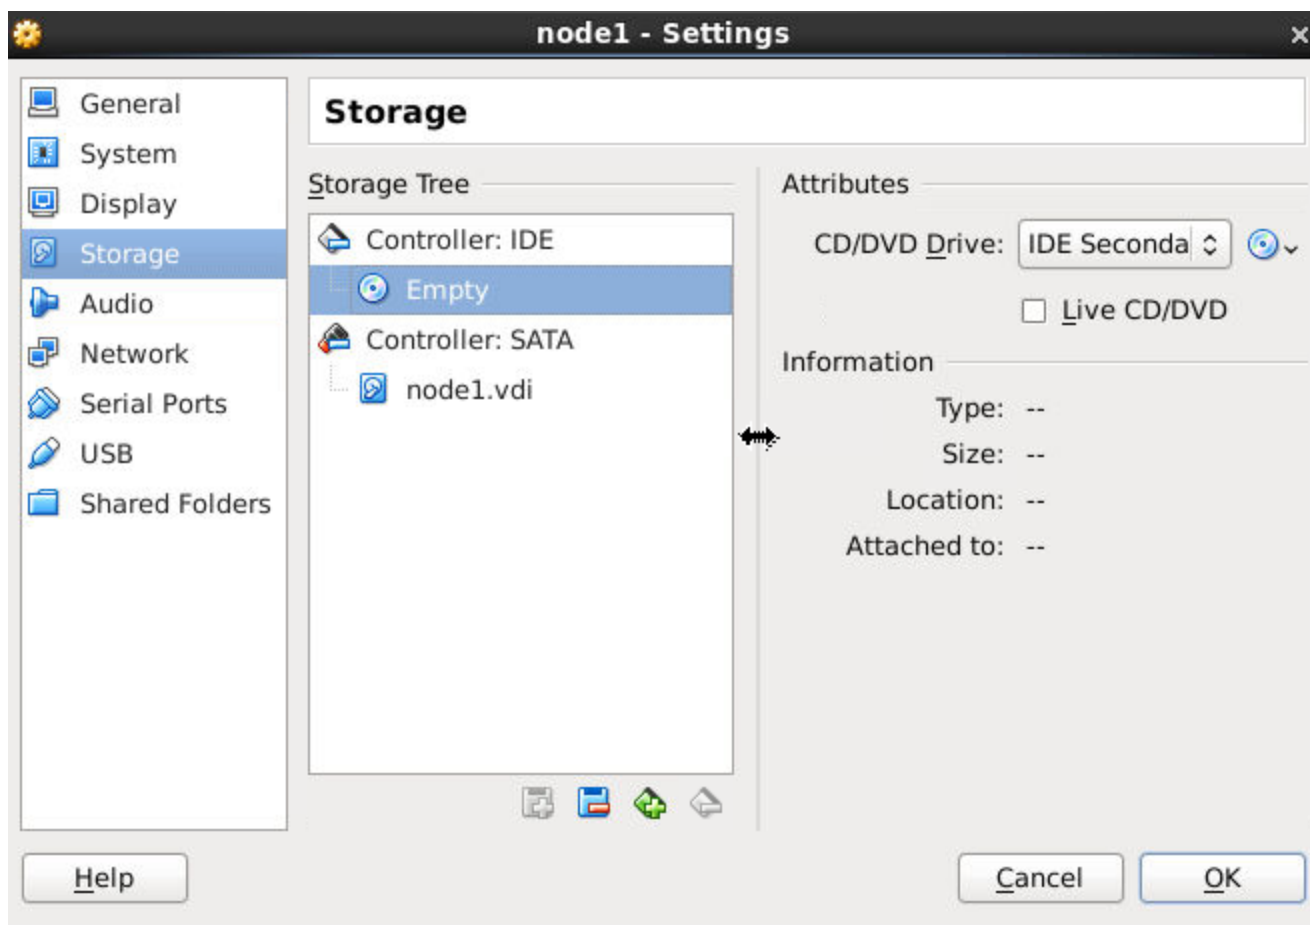
Same for "Adapter 3"



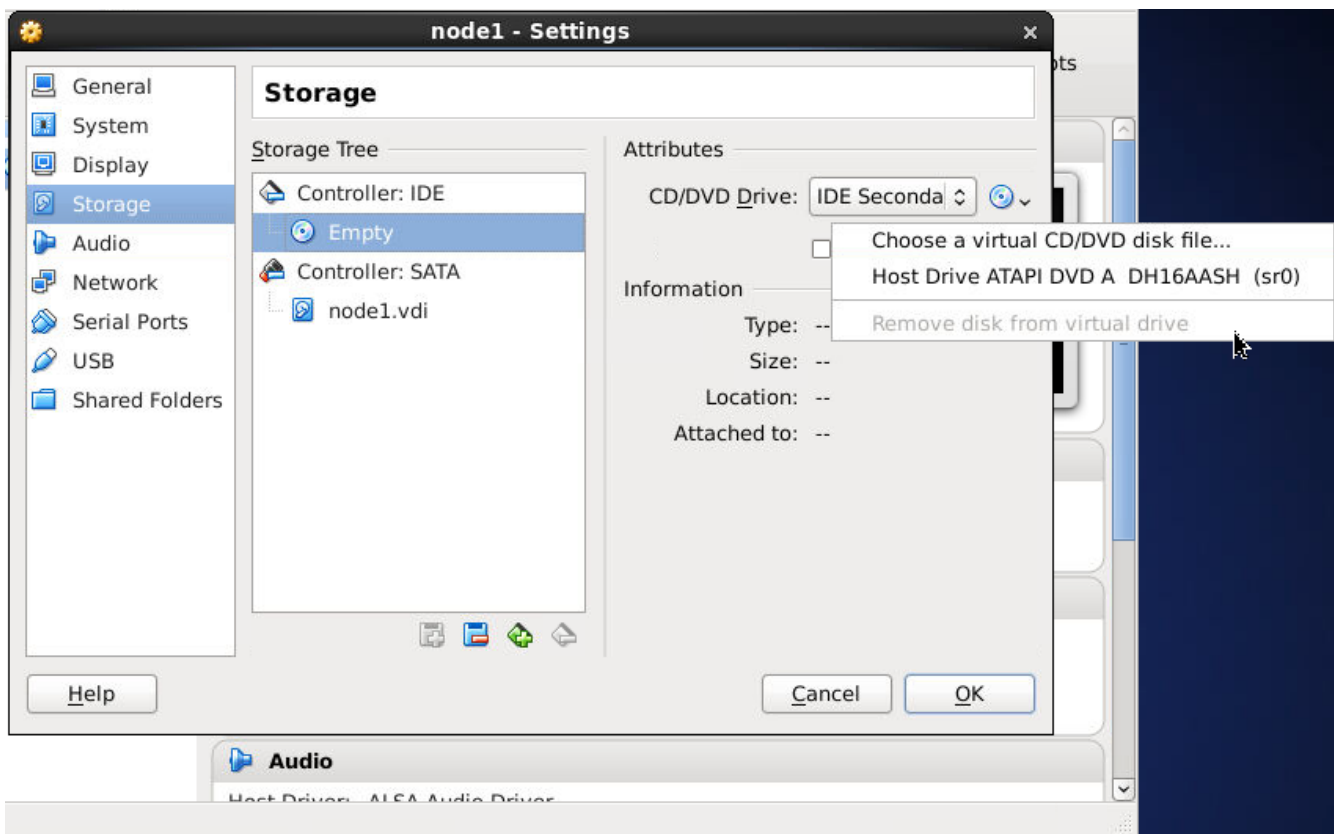
Click “Storage Next”

Click the “empty” so it is highlighted






Click on the CD icon on the extreme right.



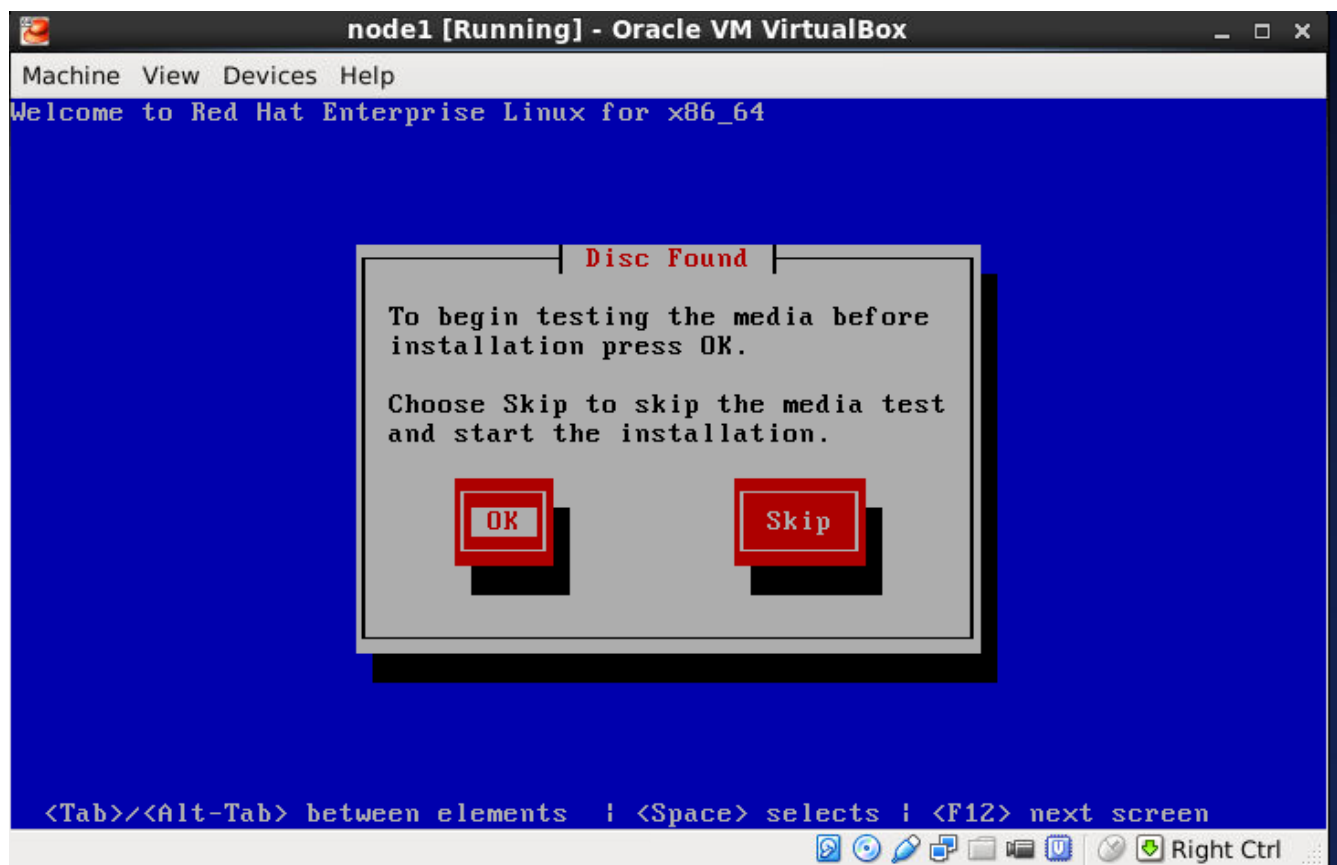
Choose the virtual CD/DVD disk file



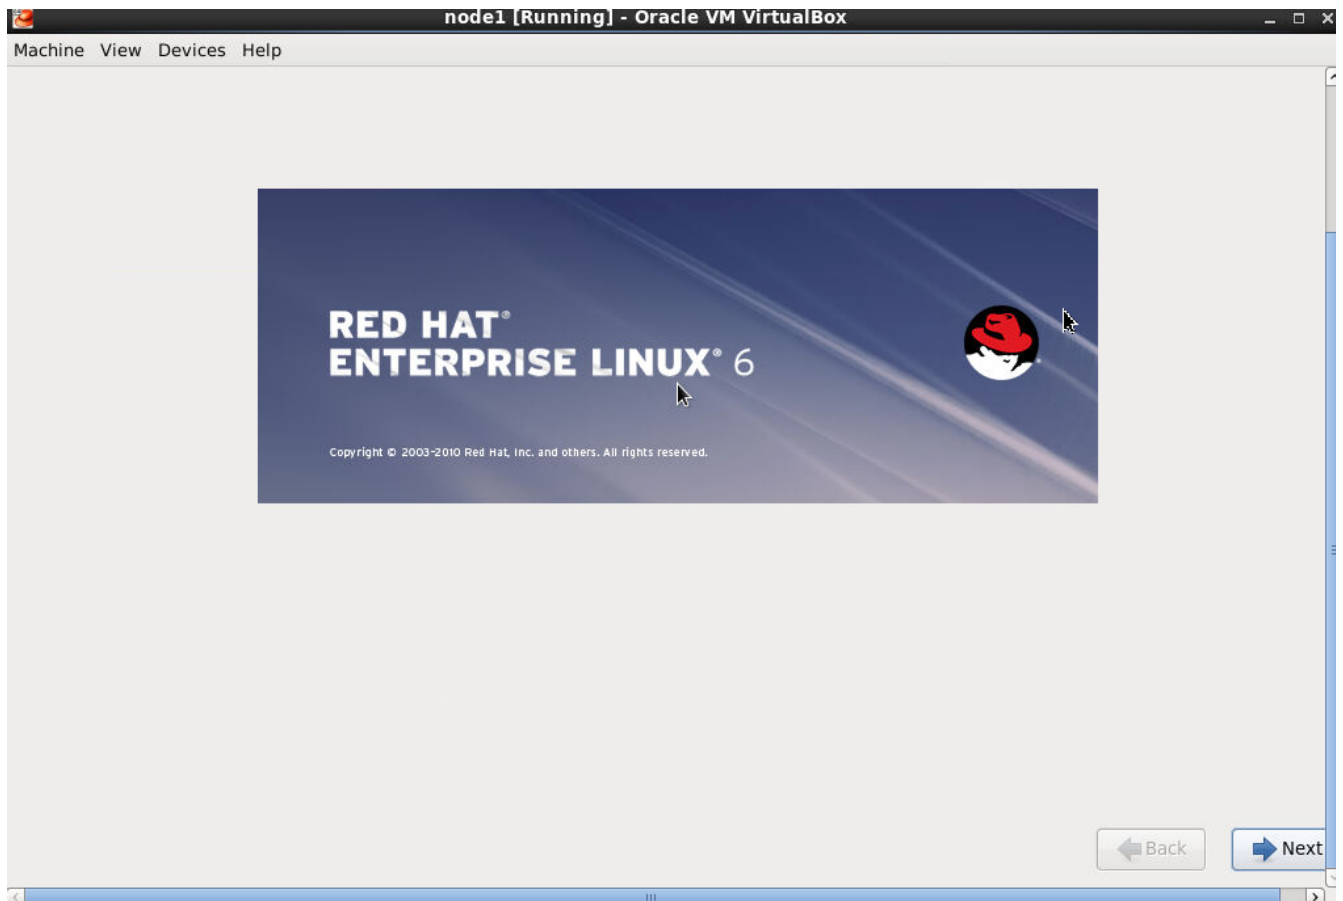
Click "open"

Click "OK"

Now we are ready to start this VM. Please click on "Start".



Tab to "Skip" and hit enter






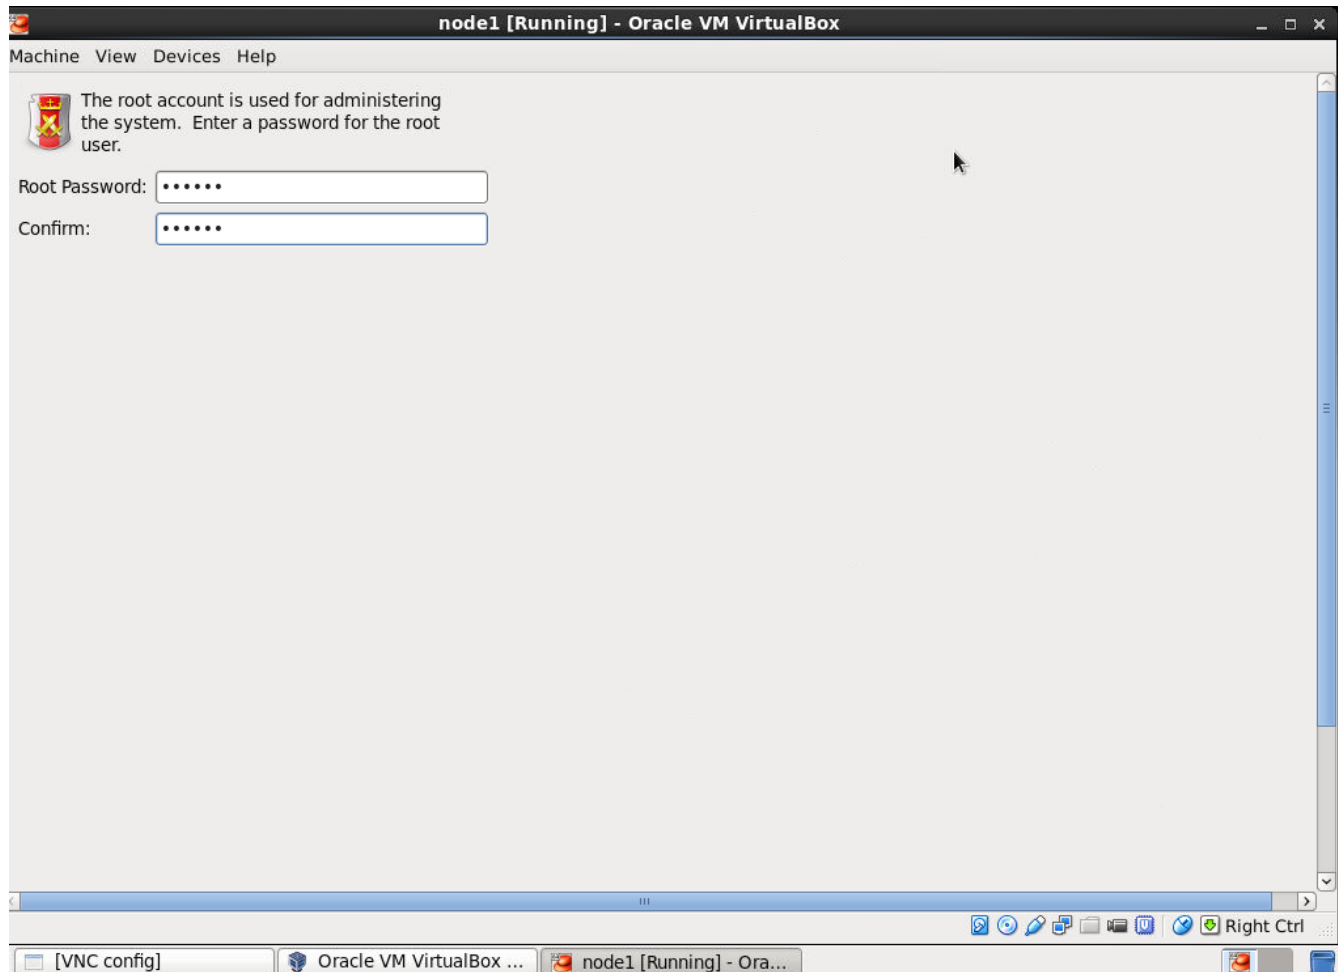
Next Next Next



“Yes, discard any data”

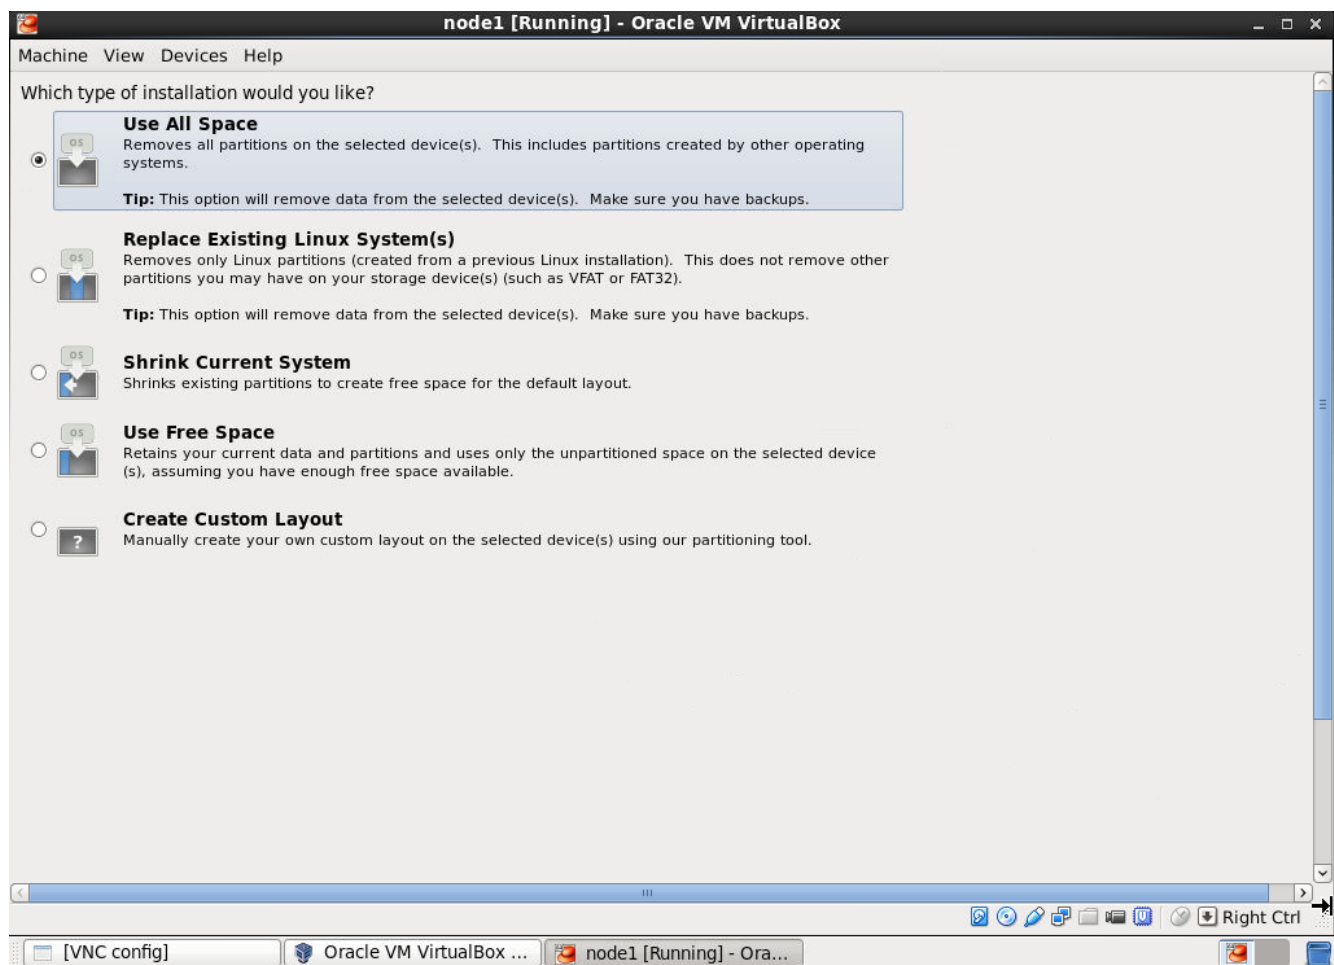


Select the default New York time zone



Select a password.

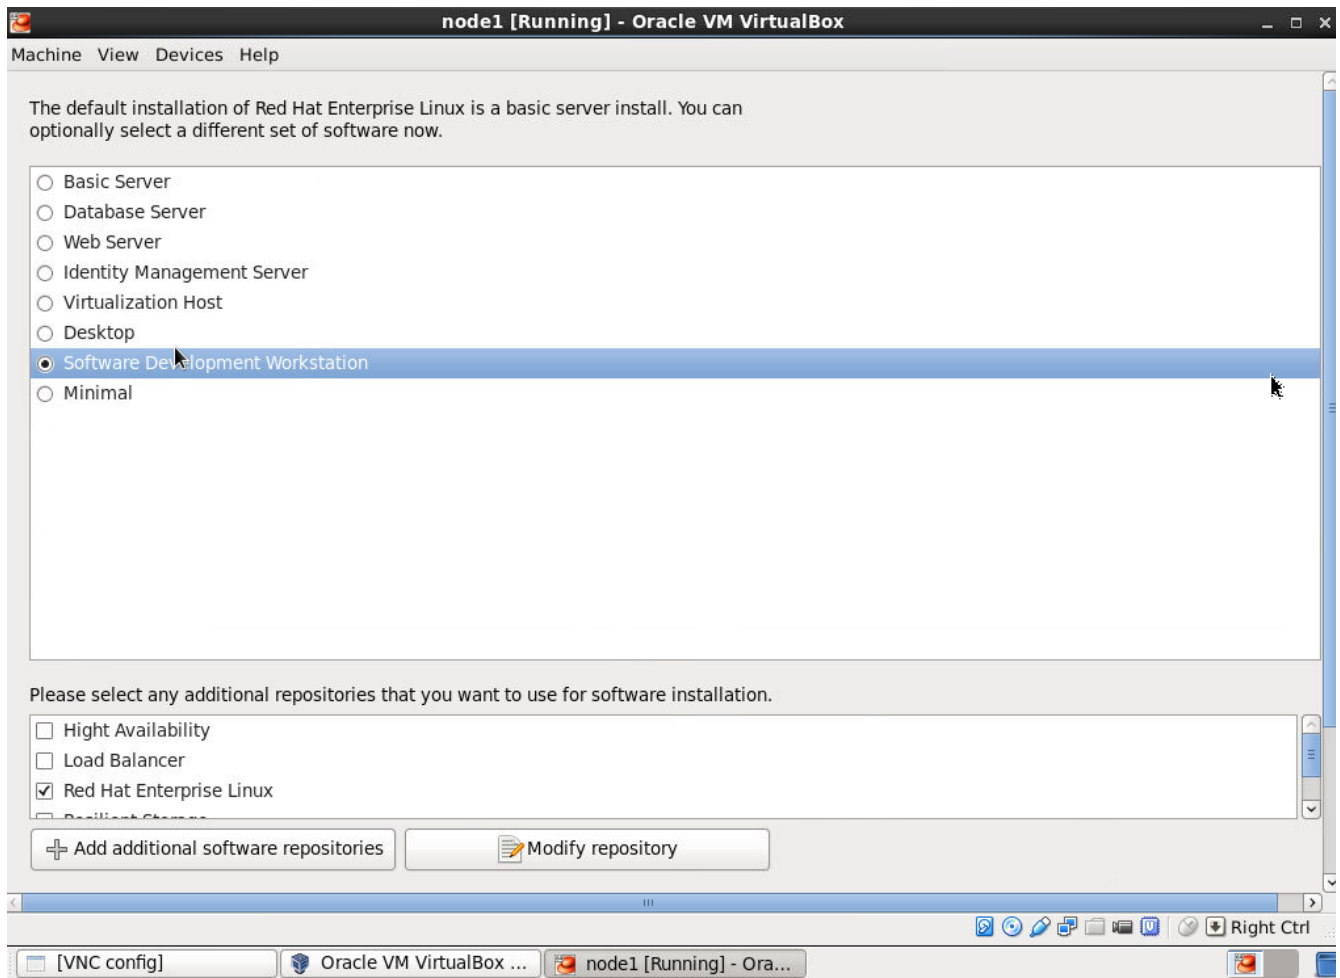


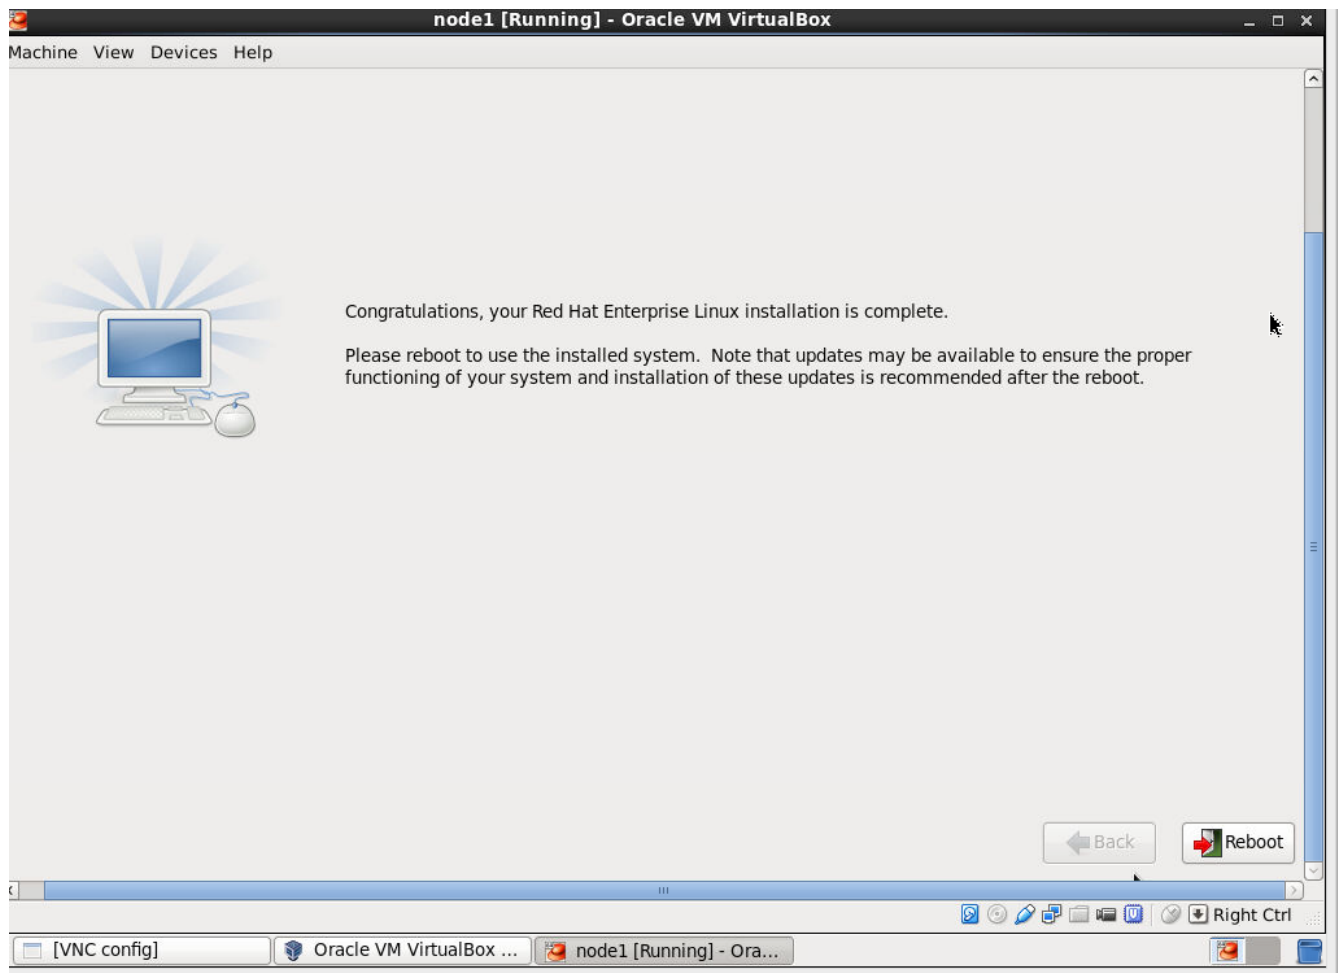


Use All Space

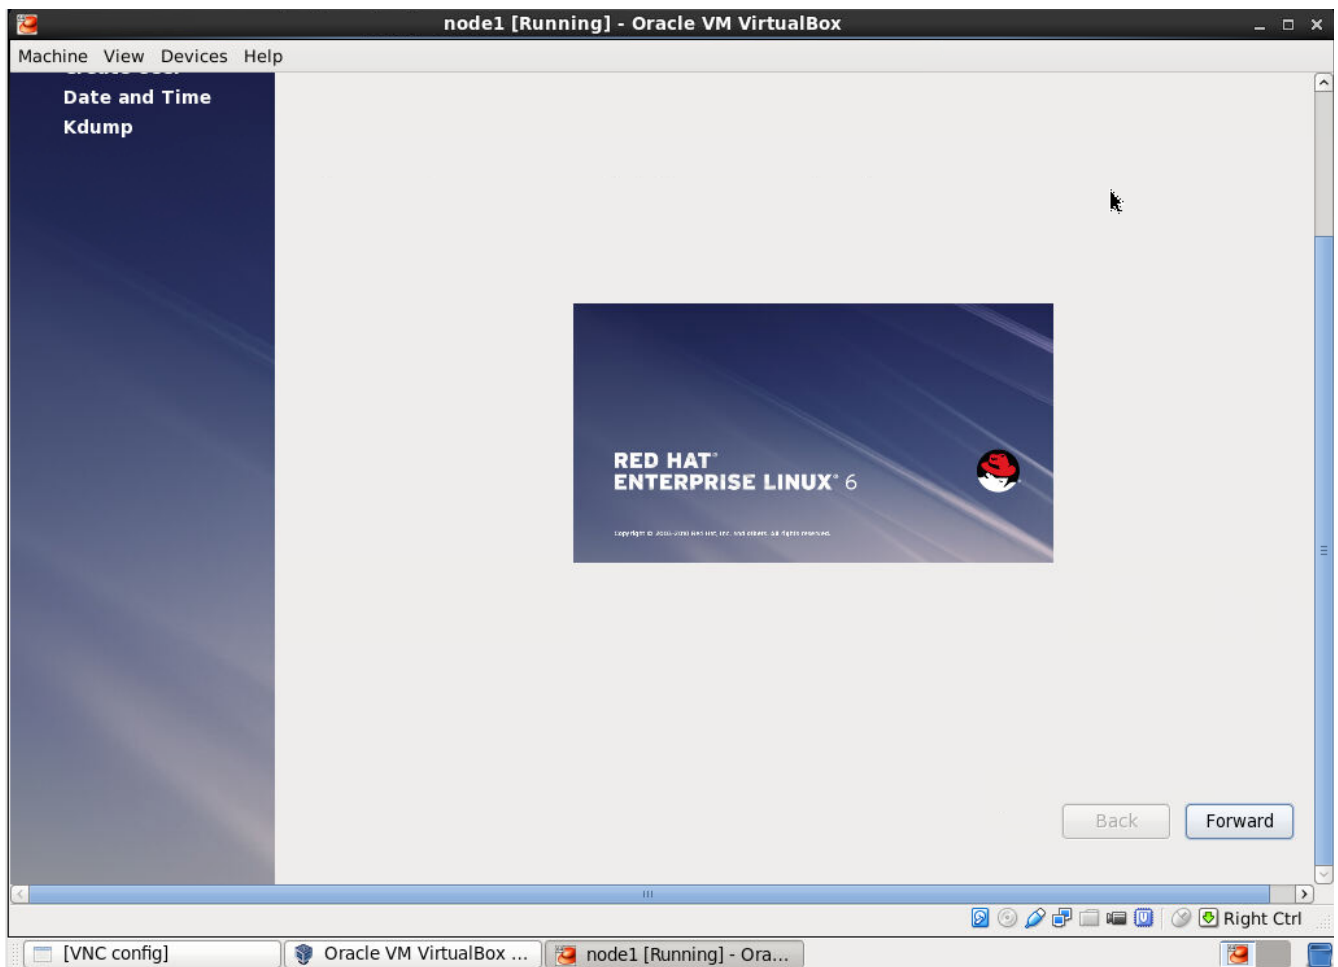


“Write Changes to Disk”

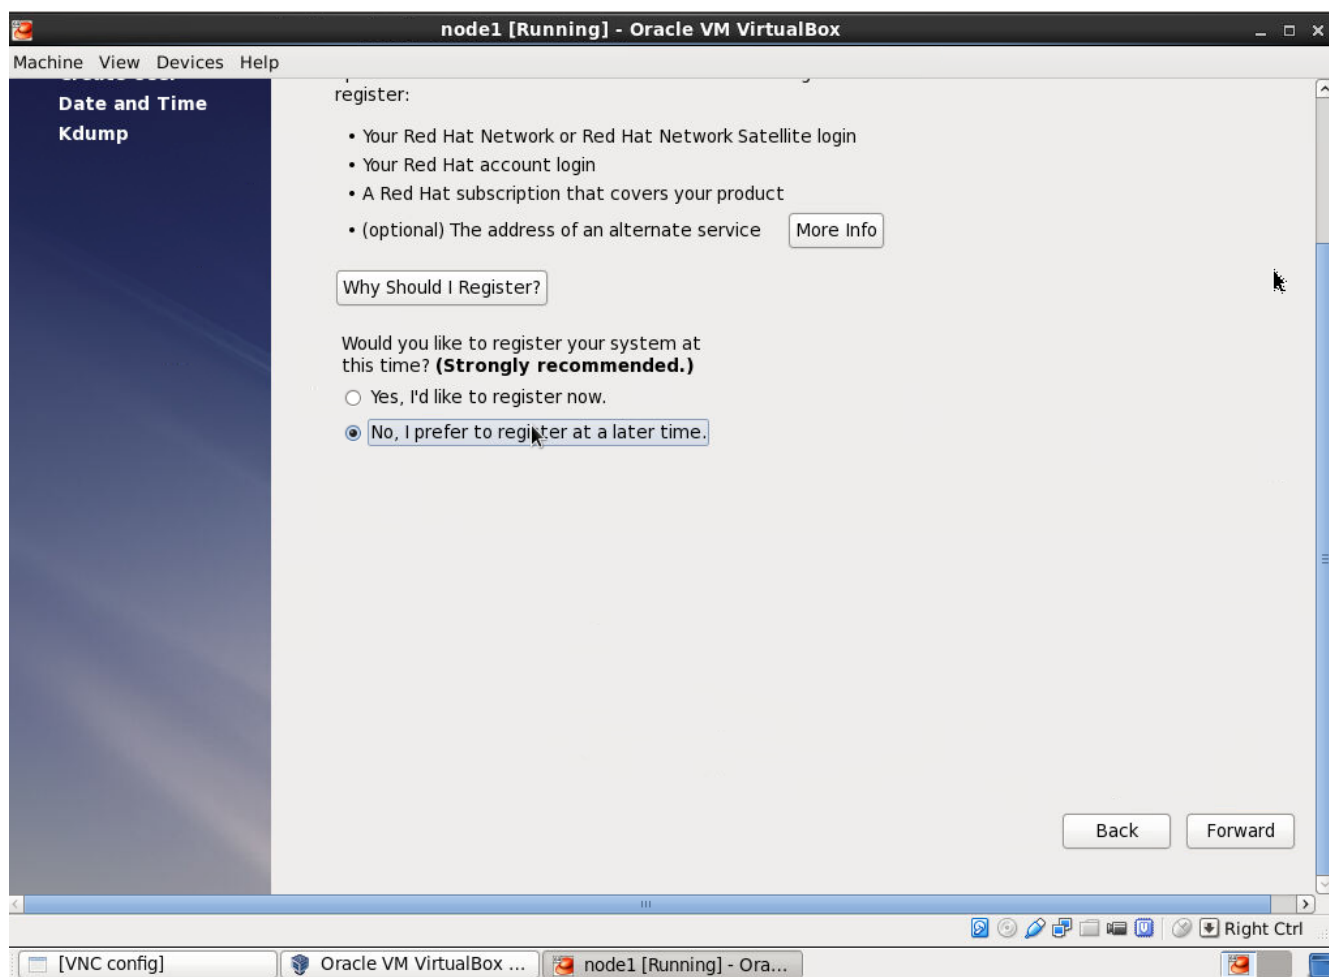




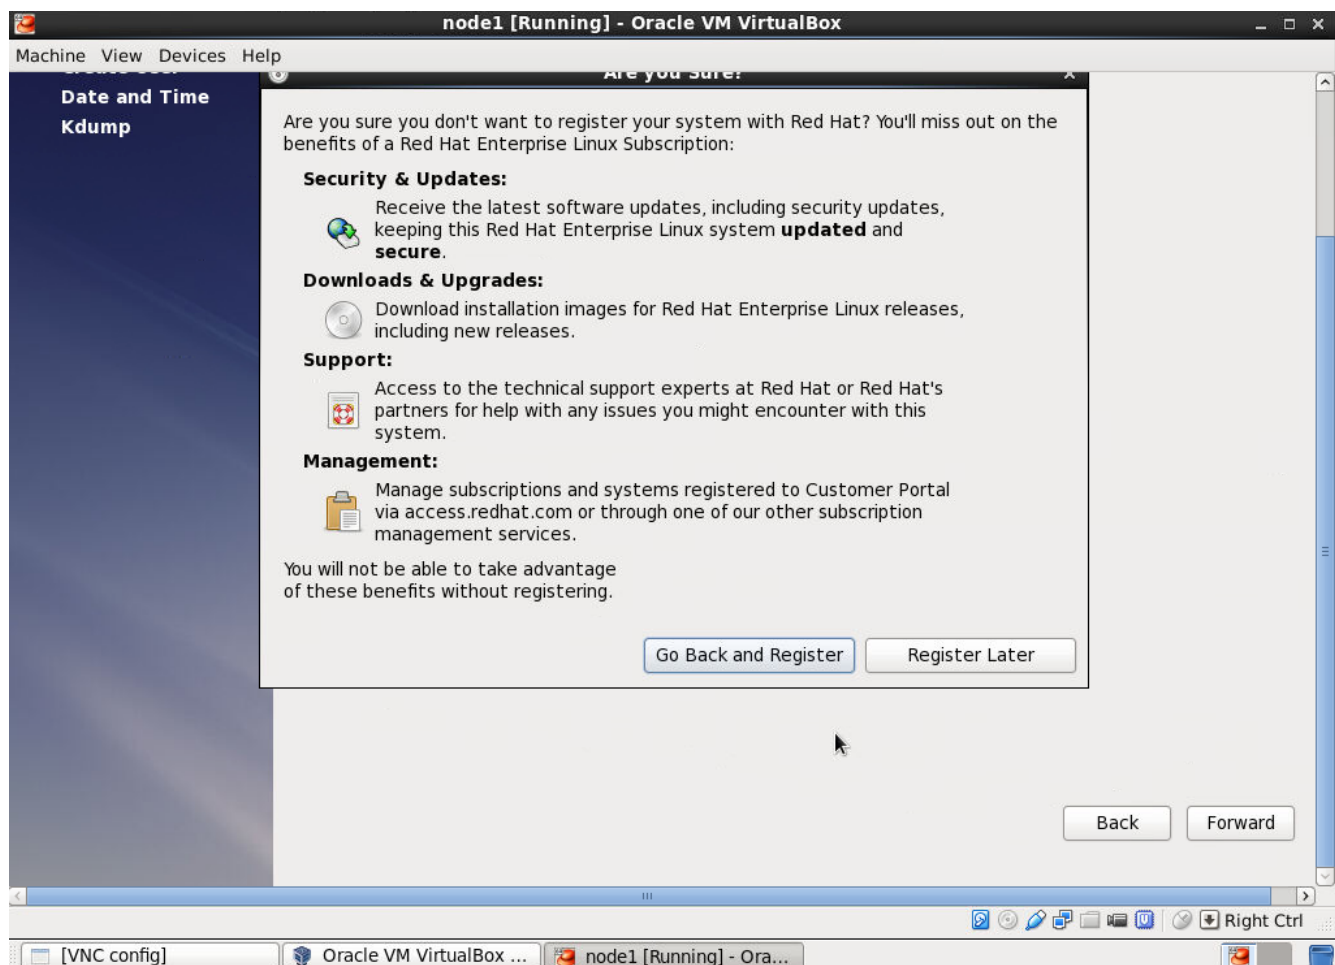
Reboot



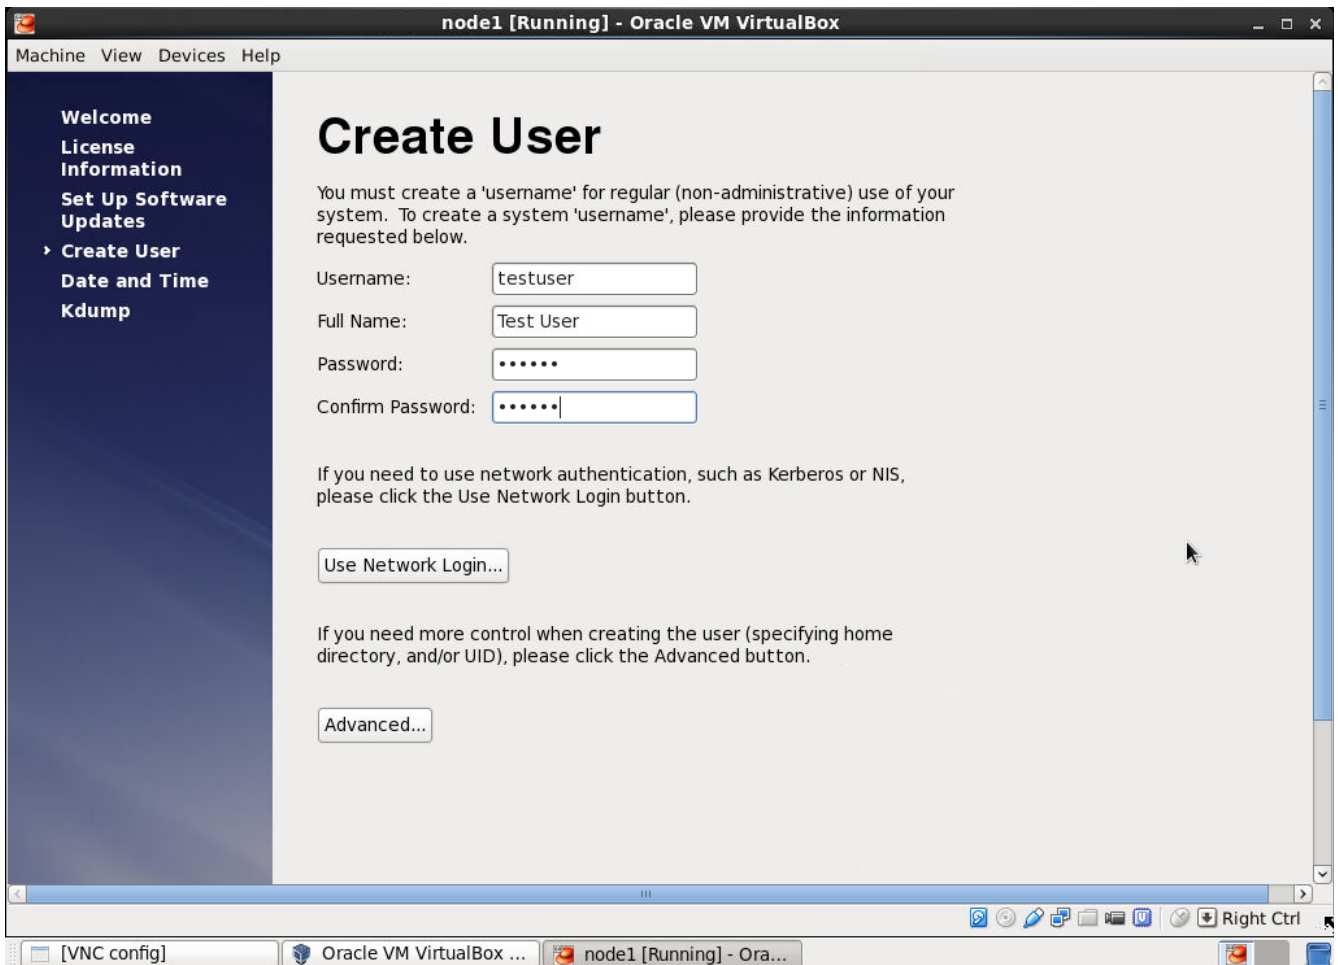
Forward

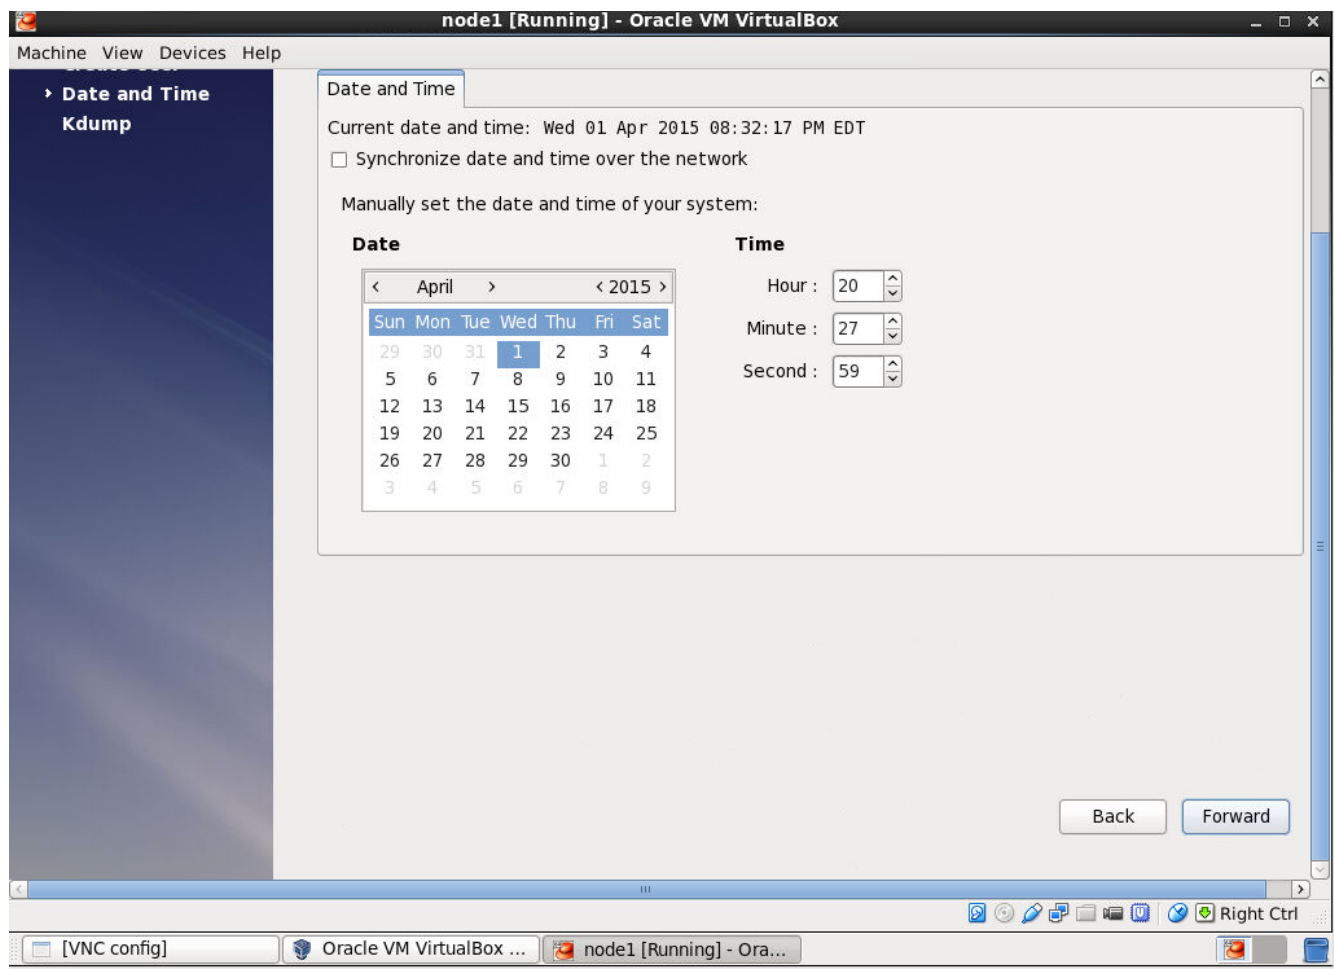


Forward



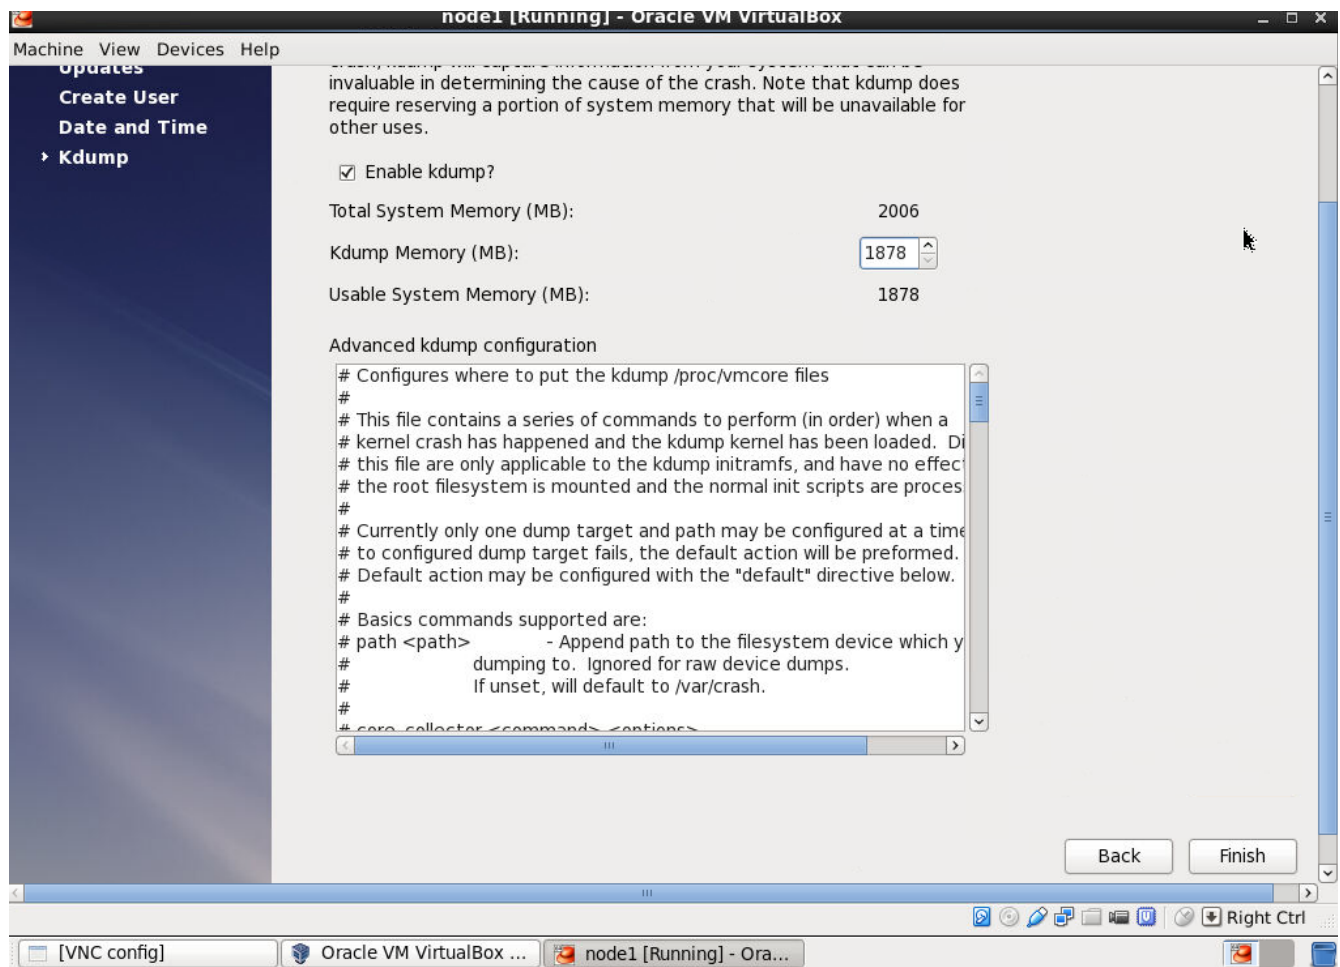
Register Later  
then  
Forward






Forward





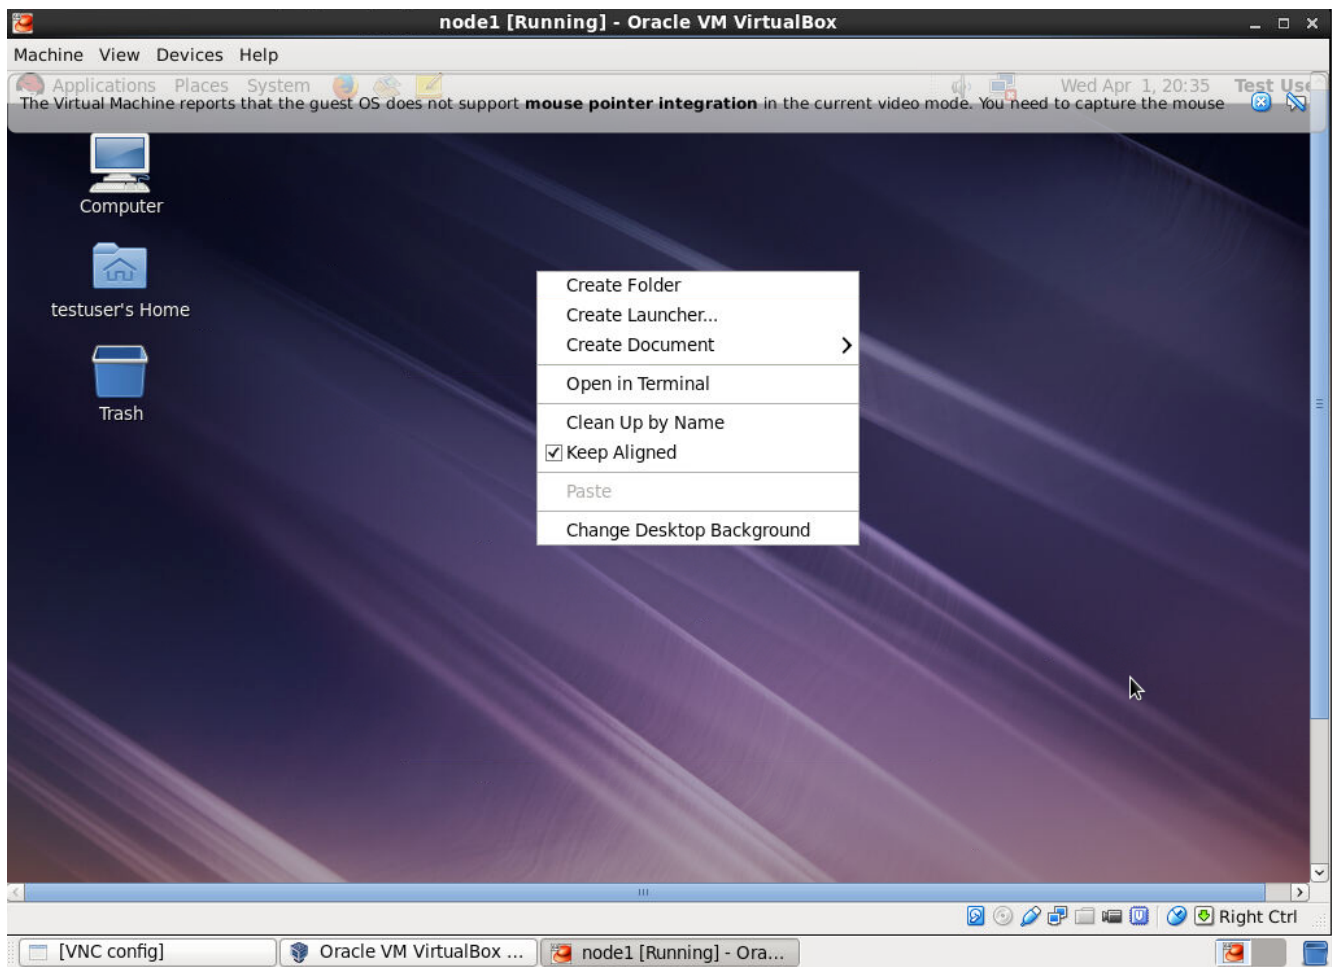
Finish

Reboot

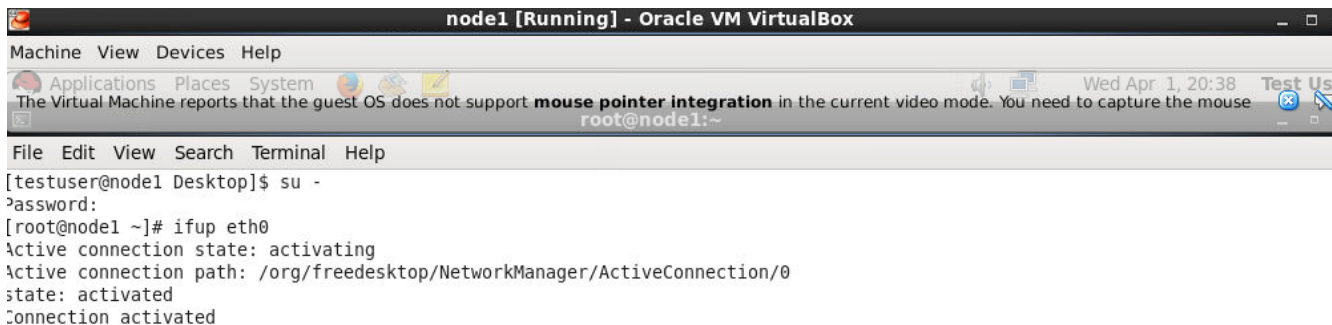


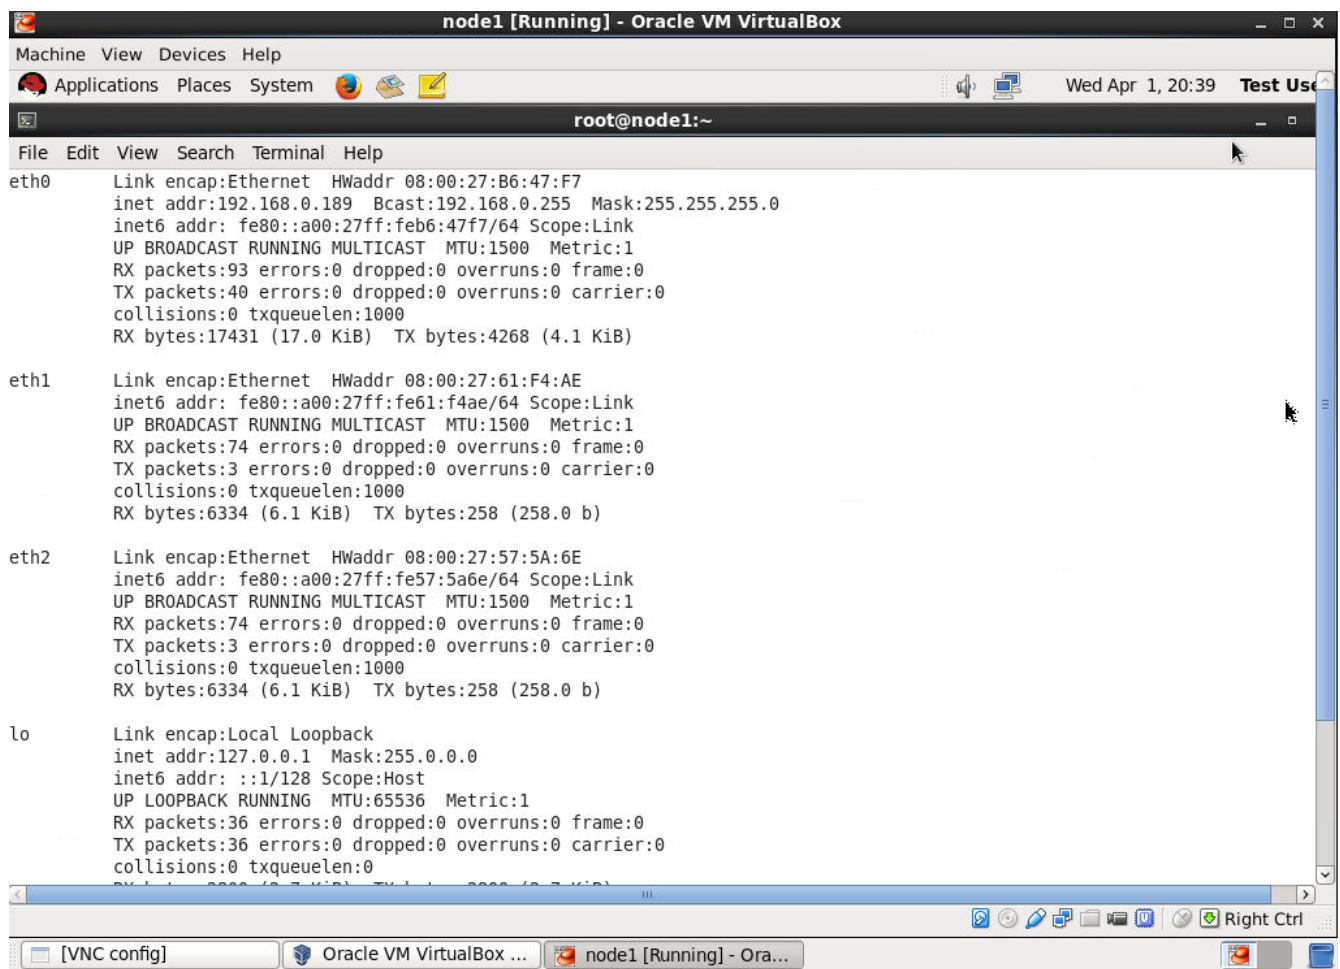
Login

Then Right Click



“Open in Terminal”





```
node1 [Running] - Oracle VM VirtualBox
Machine View Devices Help
Applications Places System
root@node1:~
File Edit View Search Terminal Help

eth0  Link encap:Ethernet HWaddr 08:00:27:B6:47:F7
      inet addr:192.168.0.189 Bcast:192.168.0.255 Mask:255.255.255.0
      inet6 addr: fe80::a00:27ff:feb6:47f7/64 Scope:Link
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:93 errors:0 dropped:0 overruns:0 frame:0
      TX packets:40 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:17431 (17.0 KiB) TX bytes:4268 (4.1 KiB)

eth1  Link encap:Ethernet HWaddr 08:00:27:61:F4:AE
      inet6 addr: fe80::a00:27ff:fe61:f4ae/64 Scope:Link
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:74 errors:0 dropped:0 overruns:0 frame:0
      TX packets:3 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:6334 (6.1 KiB) TX bytes:258 (258.0 b)

eth2  Link encap:Ethernet HWaddr 08:00:27:57:5A:6E
      inet6 addr: fe80::a00:27ff:fe57:5a6e/64 Scope:Link
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:74 errors:0 dropped:0 overruns:0 frame:0
      TX packets:3 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:6334 (6.1 KiB) TX bytes:258 (258.0 b)


lo    Link encap:Local Loopback
      inet addr:127.0.0.1 Mask:255.0.0.0
      inet6 addr: ::1/128 Scope:Host
      UP LOOPBACK RUNNING MTU:65536 Metric:1
      RX packets:36 errors:0 dropped:0 overruns:0 frame:0
      TX packets:36 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
```

So we observe that the eth0 dhcp assigned address is 192.168.0.189, we will make this a static address.

At this time, please download putty from:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>

and connect to 192.168.0.189 via putty

A terminal window titled 'root@node1:~' with a black background and white text. The text shows a login sequence: 'login as: root', 'root@192.168.0.189's password:', and a green cursor on the line '[root@node1 ~]#'.

```
login as: root
root@192.168.0.189's password:
[root@node1 ~]#
```

Great we are connected via putty.

```
[root@node1 ~]# vi /etc/sysconfig/network-scripts/ifcfg-eth0
```

```
DEVICE=eth0
#HWADDR=08:00:27:B6:47:F7
TYPE=Ethernet
UUID=afb738f2-64e1-4203-8ad6-79a3dd5679d4
ONBOOT=yes
NM_CONTROLLED=no
BOOTPROTO=static
IPADDR=192.168.0.189
DNS1=192.168.0.1
GATEWAY=192.168.0.1
```

NOTE: Please replace DNS1 and GATEWAY by your routers IP. Please note NM\_CONTROLLED=no  
Also do this:

```
[root@node1 ~]# service NetworkManager stop
Stopping NetworkManager daemon:
[root@node1 ~]# chkconfig NetworkManager off
[root@node1 ~]#
```

Then:

```
[root@node1 ~]# service network restart
Shutting down interface eth0: Device state: 3 (disconnected)
```

```
[ OK ]
Shutting down interface eth1: [ OK ]
Shutting down interface eth2: [ OK ]
Shutting down loopback interface: [ OK ]
Bringing up loopback interface: [ OK ]
Bringing up interface eth0: Active connection state: activated
Active connection path: /org/freedesktop/NetworkManager/ActiveConnection/2
[ OK ]
[root@node1 ~]#
```

Now we will disable two features, one is selinux and the other is iptables:

```
[root@node1 ~]# vi /etc/selinux/config
change "enforcing" to "disabled", it should look like this:
```

```
[root@node1 ~]# cat /etc/selinux/config

# This file controls the state of SELinux on the system.
# SELINUX= can take one of these three values:
#   enforcing - SELinux security policy is enforced.
#   permissive - SELinux prints warnings instead of enforcing.
#   disabled - No SELinux policy is loaded.
SELINUX=disabled
# SELINUXTYPE= can take one of these two values:
#   targeted - Targeted processes are protected,
#   mls - Multi Level Security protection.
SELINUXTYPE=targeted
```

```
[root@node1 ~]#
```

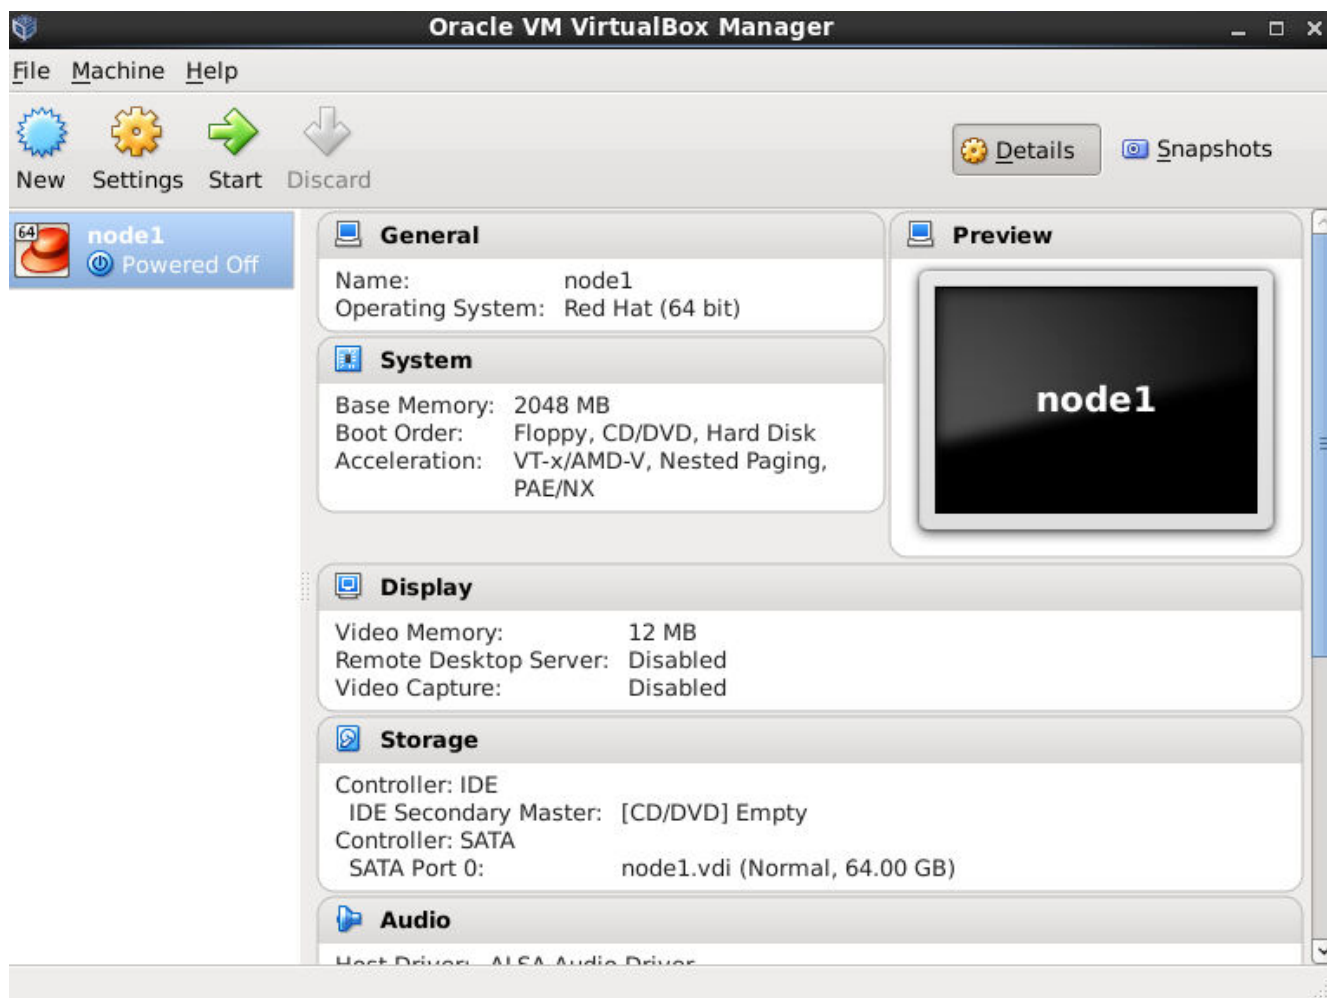
Then, do:

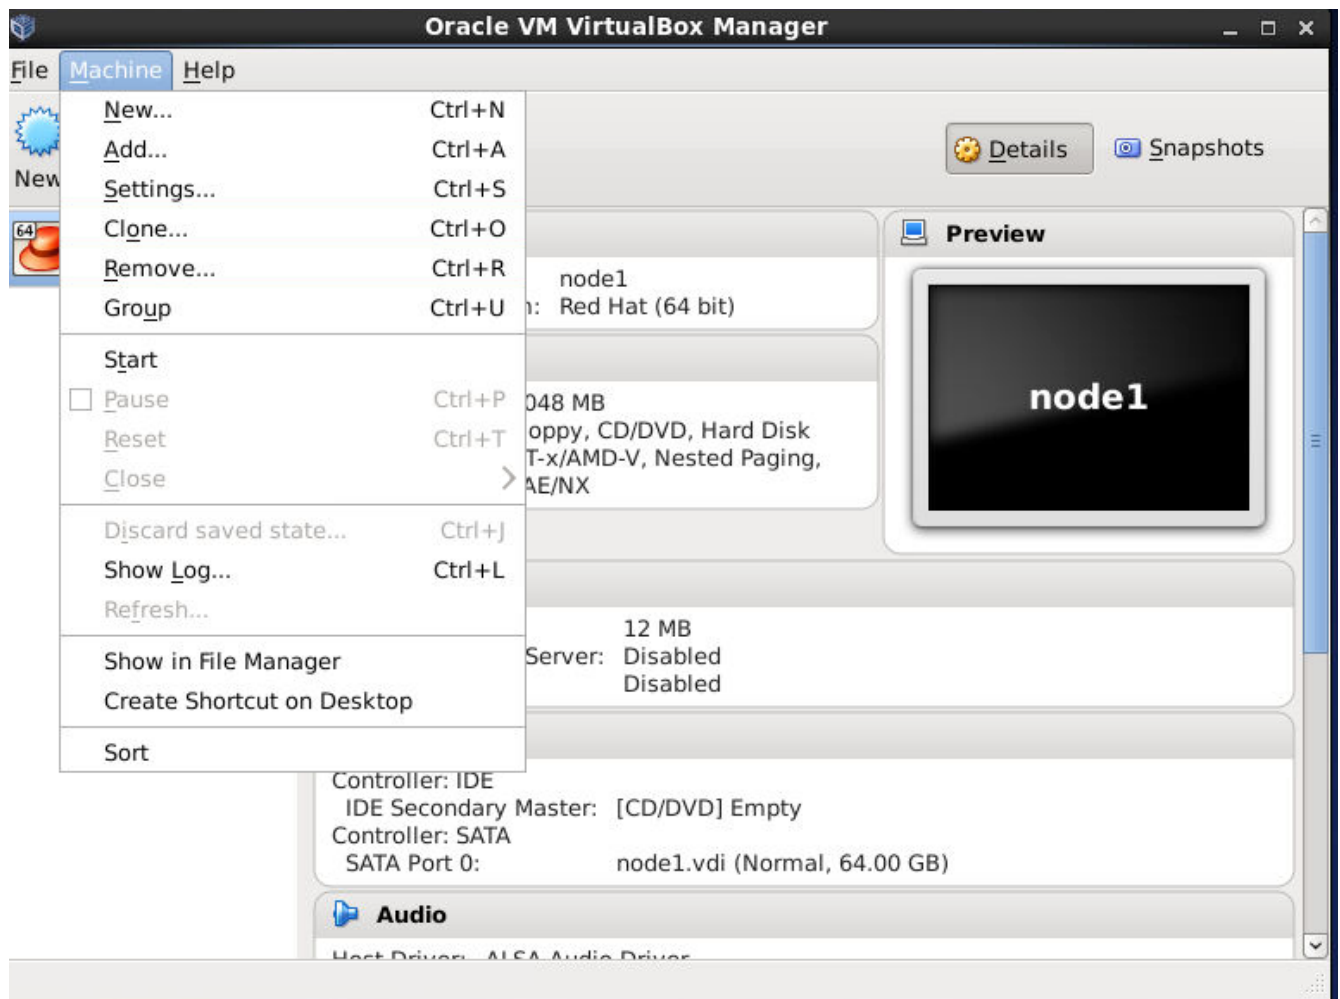
```
[root@node1 ~]# chkconfig iptables off
```

Now we will power it down and clone it into another system called node2.

```
[root@node1 ~]# init 0
```

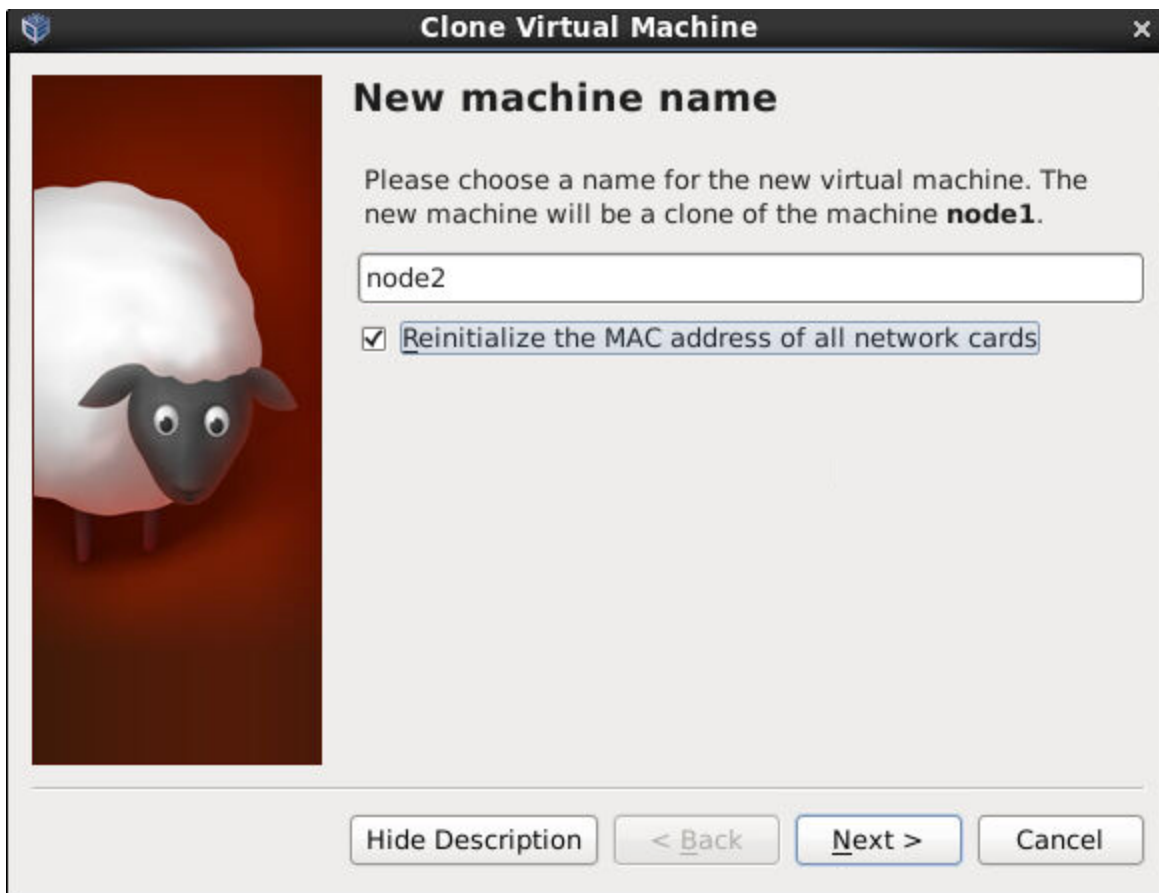
As you see below, node1 is powered off:



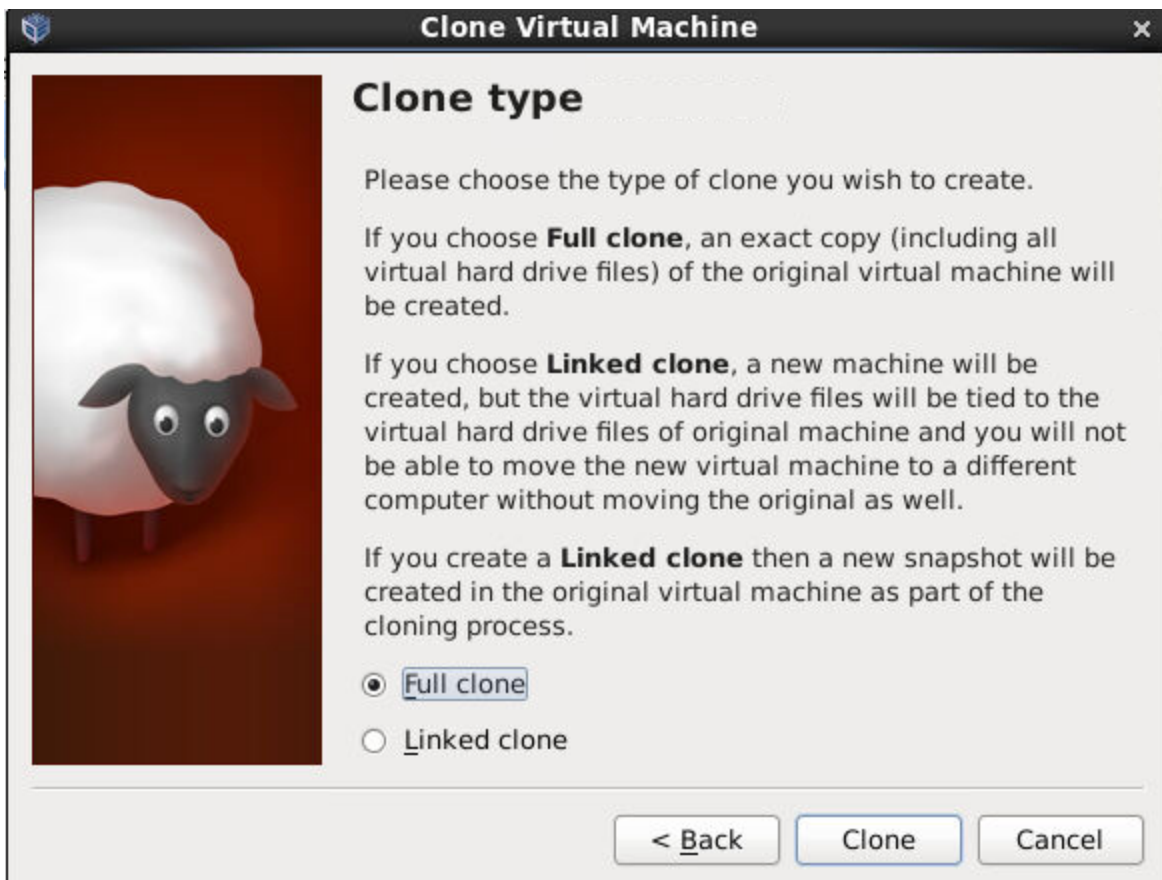


Select clone

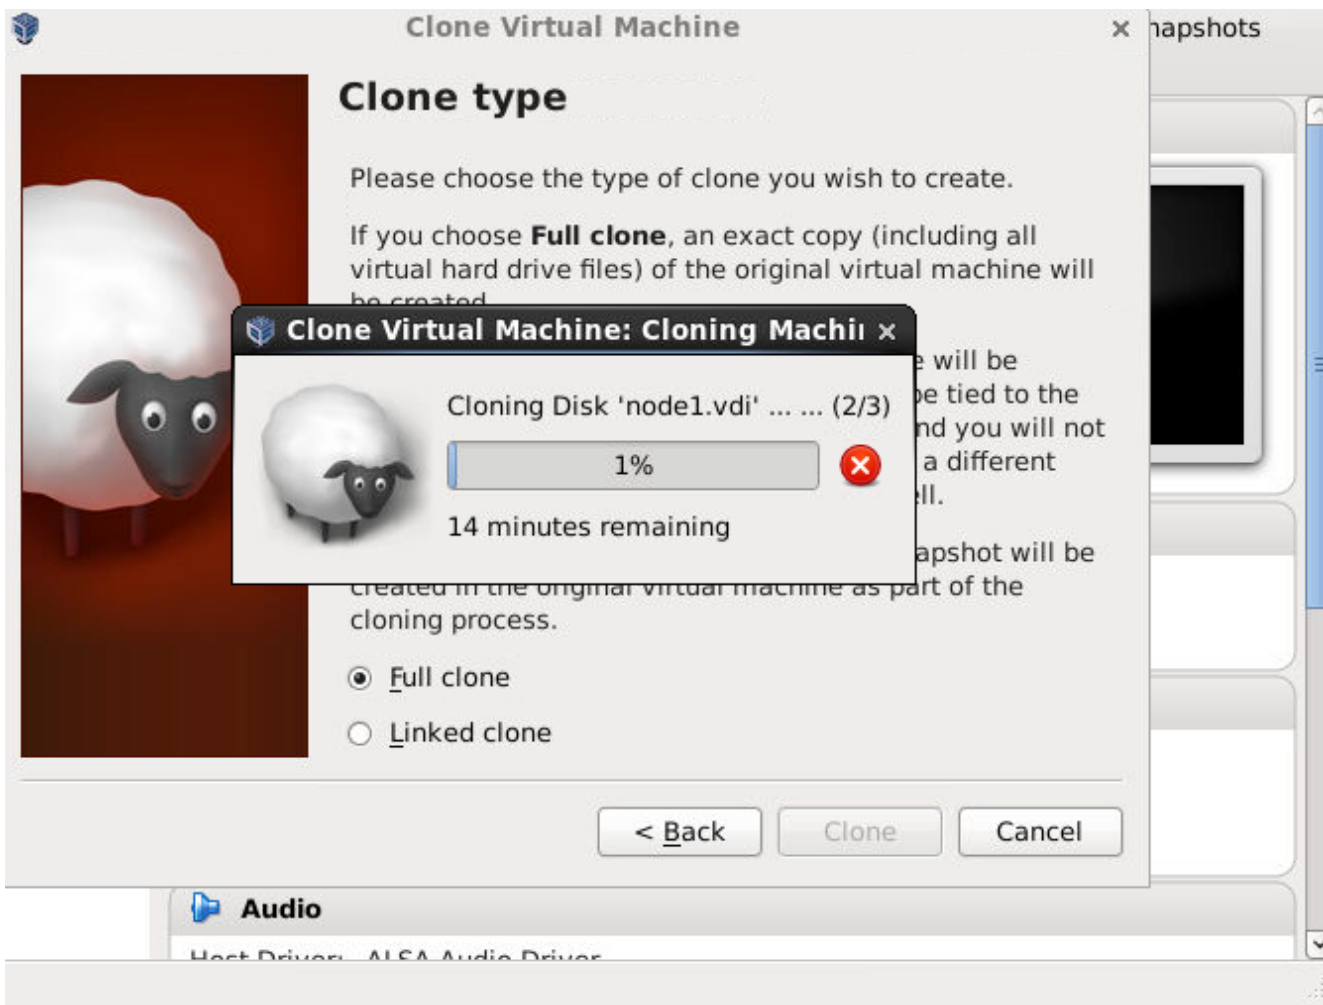




Please click "Next"




Click "Clone"





We will now power on node2 and change it's name in the operating system to node2, give it a new IP address, this will make them exactly two different systems.

Highlight node2 and start it



The screenshot shows a terminal window with a title bar that reads "root@node1:~". The menu bar includes "File", "Edit", "View", "Search", "Terminal", and "Help". The terminal content displays the following configuration for the network interface eth0:

```
DEVICE=eth0
HWADDR=08:00:27:B6:47:F7
TYPE=Ethernet
UUID=afb738f2-64e1-4203-8ad6-79a3dd5679d4
ONBOOT=yes
NM_CONTROLLED=yes
BOOTPROTO=static
IPADDR=192.168.0.190
DNS1=192.168.0.1
GATEWAY=192.168.0.1
```

Below the configuration, there are several lines of blue wavy lines, likely representing a network diagram or a series of status indicators. At the bottom of the terminal, the text "-- INSERT --" is visible.

NOTE: we commented the HWADDR. Changed the IP address to 192.168.0.190

```
vi /etc/sysconfig/network
```

[illegible]

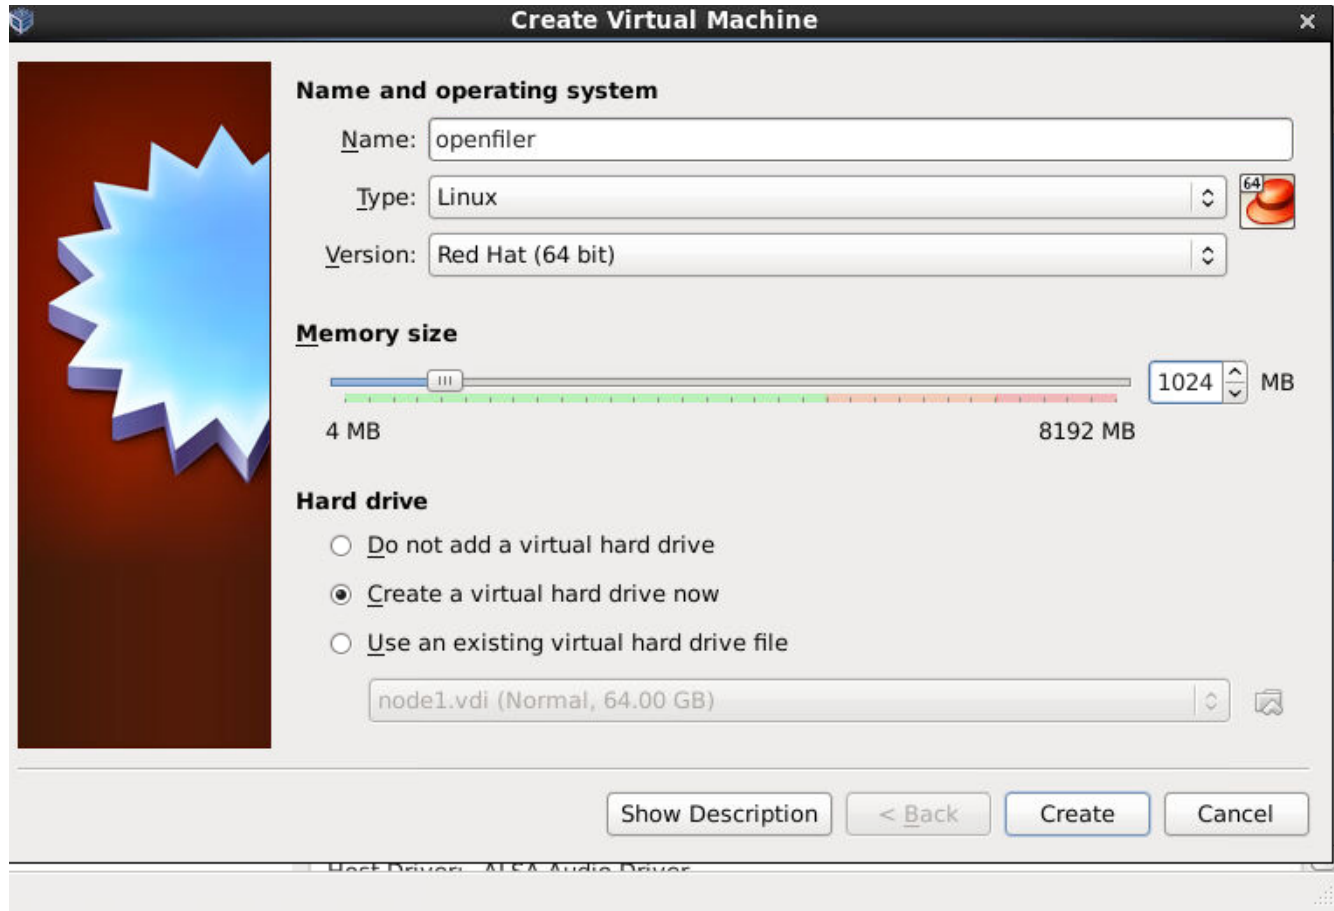
reboot

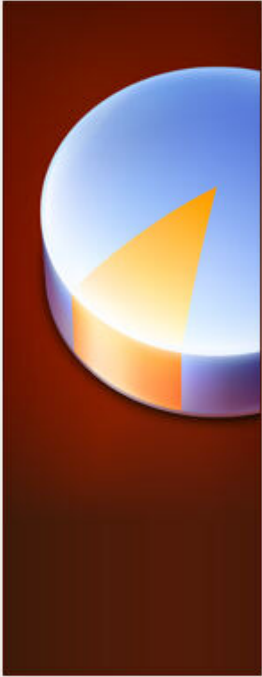
While you are at it, start node1 as well

Now we have both VM's which will be the two nodes of our cluster up and running.

Now we will install over virtual storage area network. This is openfiler. To do this we will again go to virtualbox.

Go to virtualbox and click “New”





### Create Virtual Hard Drive

**File location**

openfiler

**File size**

4.00 MB 2.00 TB 64.00 GB


**Hard drive file type**

- ☒ VDI (VirtualBox Disk Image)
- ☐ VMDK (Virtual Machine Disk)
- ☐ VHD (Virtual Hard Disk)
- ☐ HDD (Parallels Hard Disk)
- ☐ QED (QEMU enhanced disk)
- ☐ QCOW (QEMU Copy-On-Write)

**Storage on physical hard drive**

- ☒ Dynamically allocated
- ☐ Fixed size
- ☐ Split into files of less than 2GB

Show Description < Back Create Cancel




### Create Virtual Machine

## Name and operating system

Please choose a descriptive name for the new virtual machine and select the type of operating system you intend to install on it. The name you choose will be used throughout VirtualBox to identify this machine.

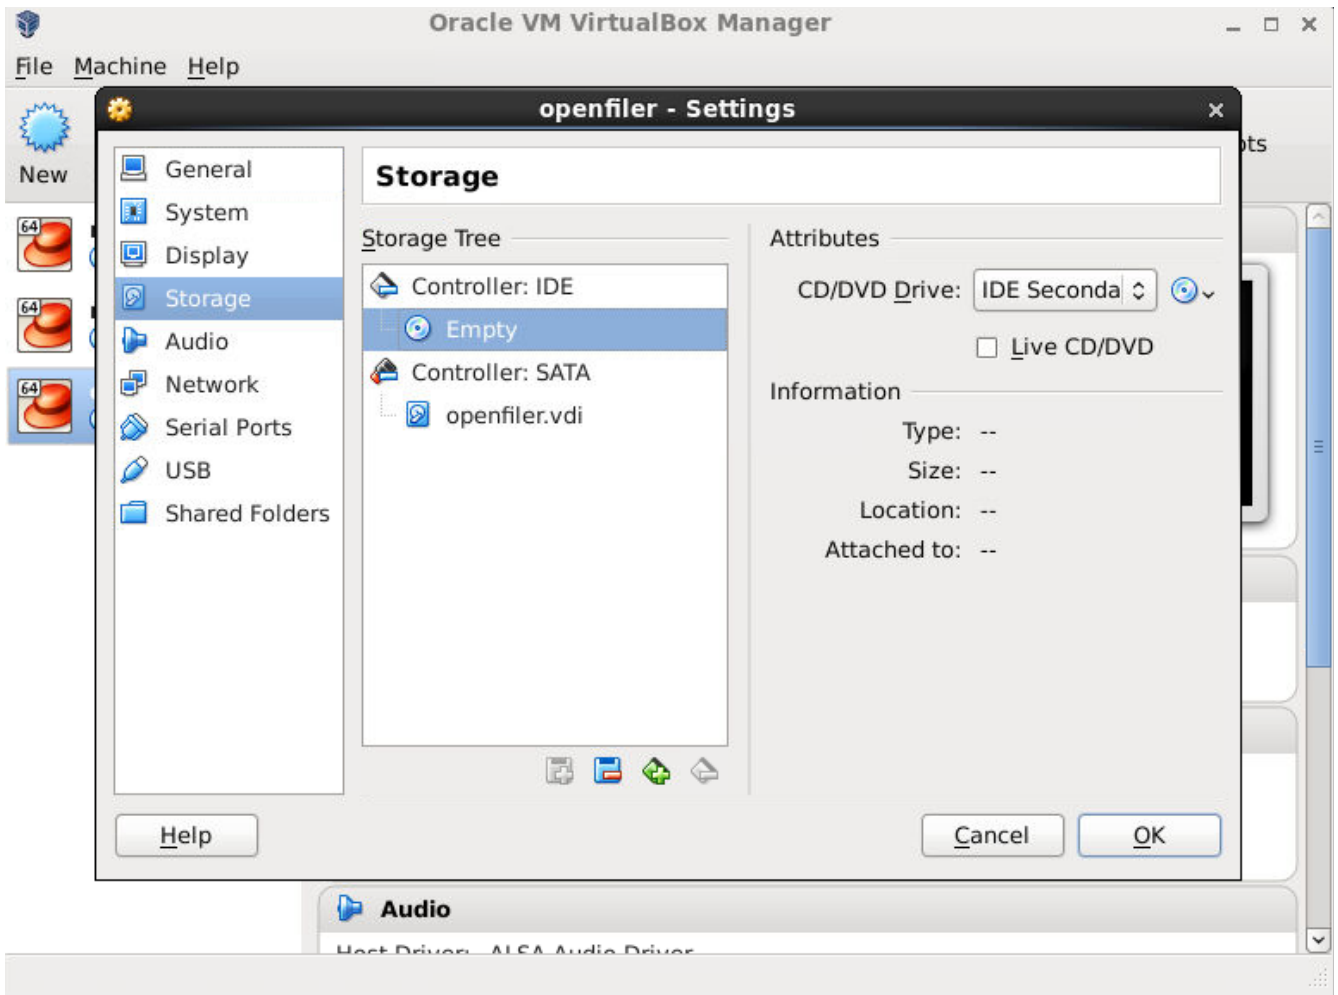
Name:

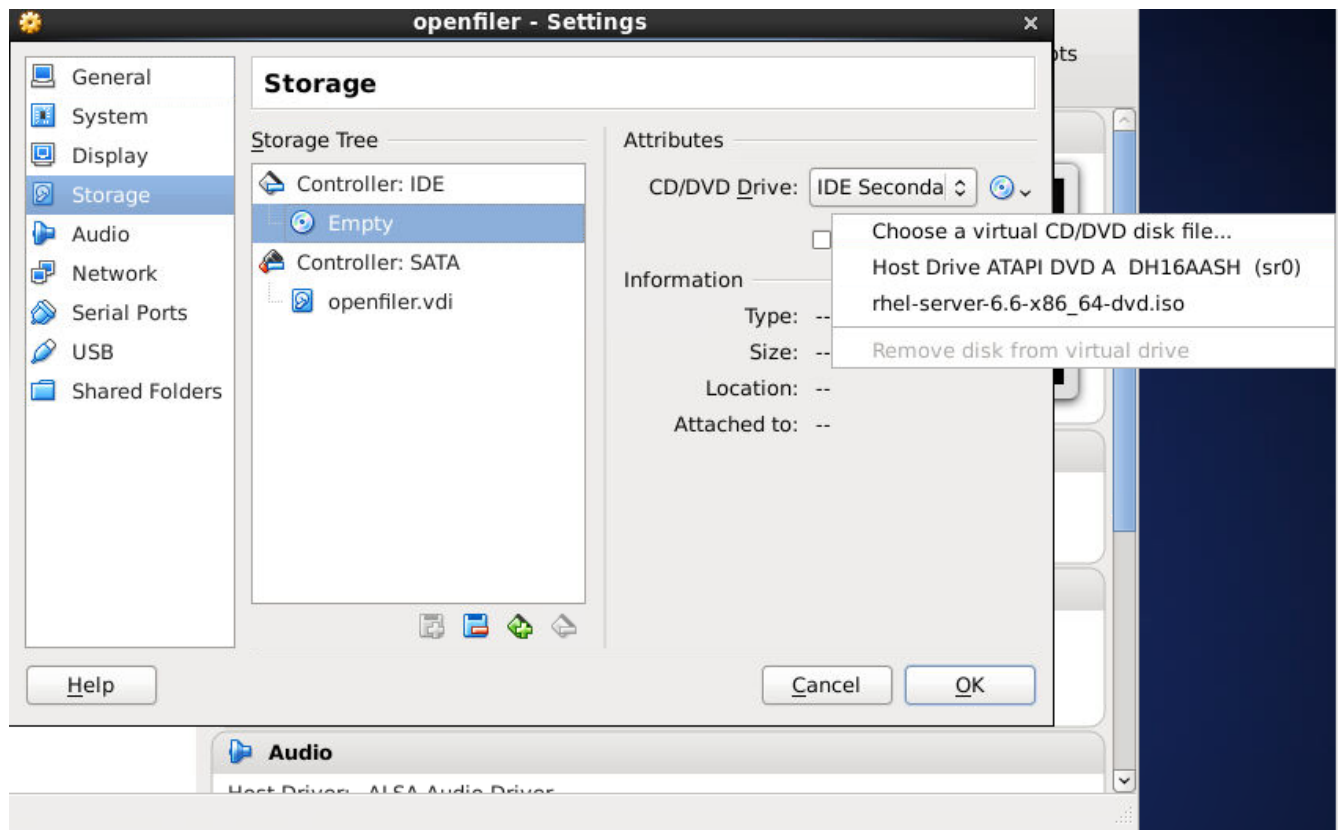
Type: Linux 


Version: Red Hat (64 bit)

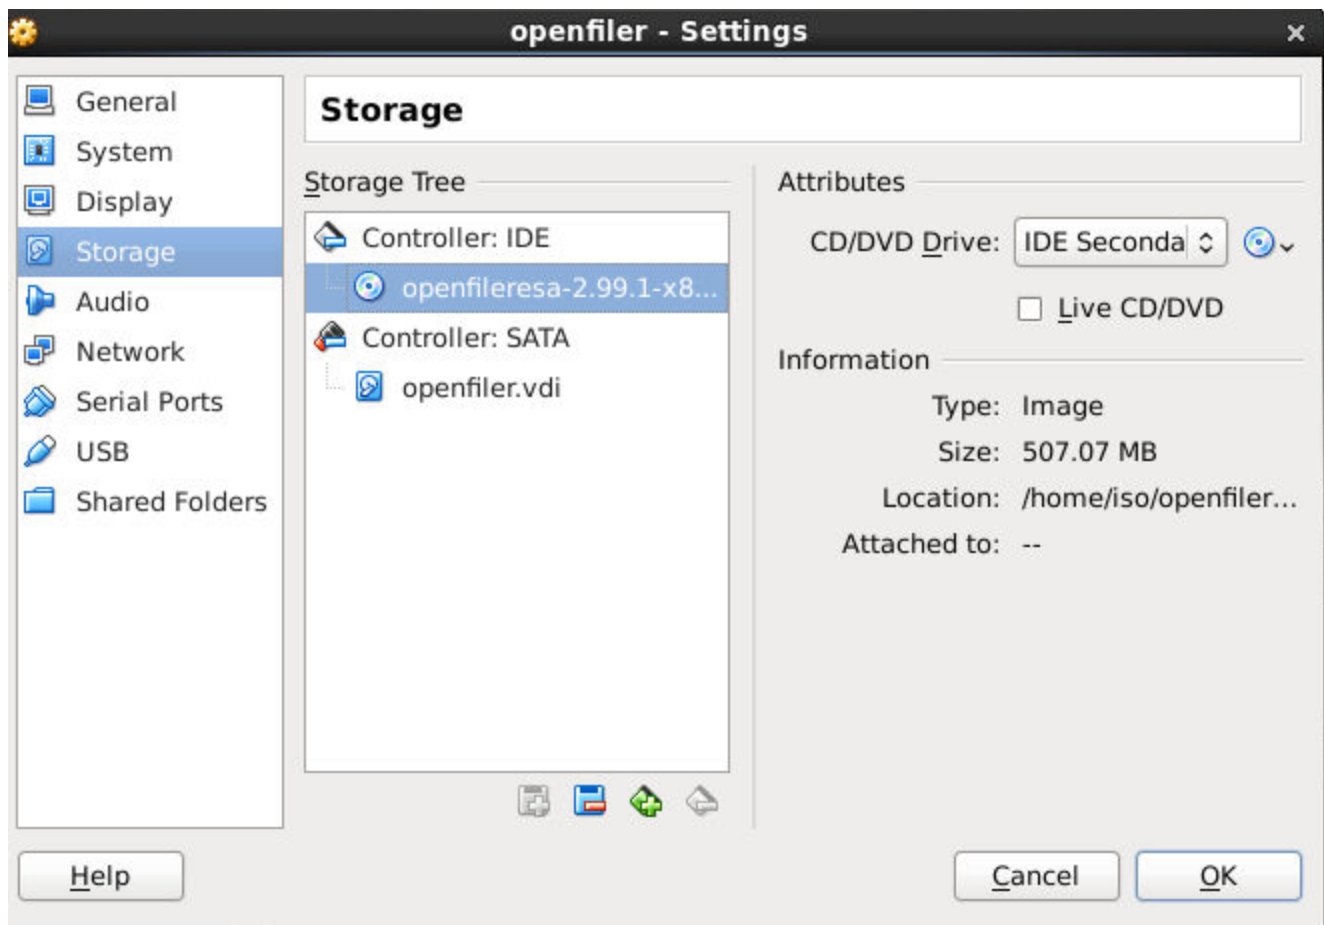
Hide Description < Back Next > Cancel



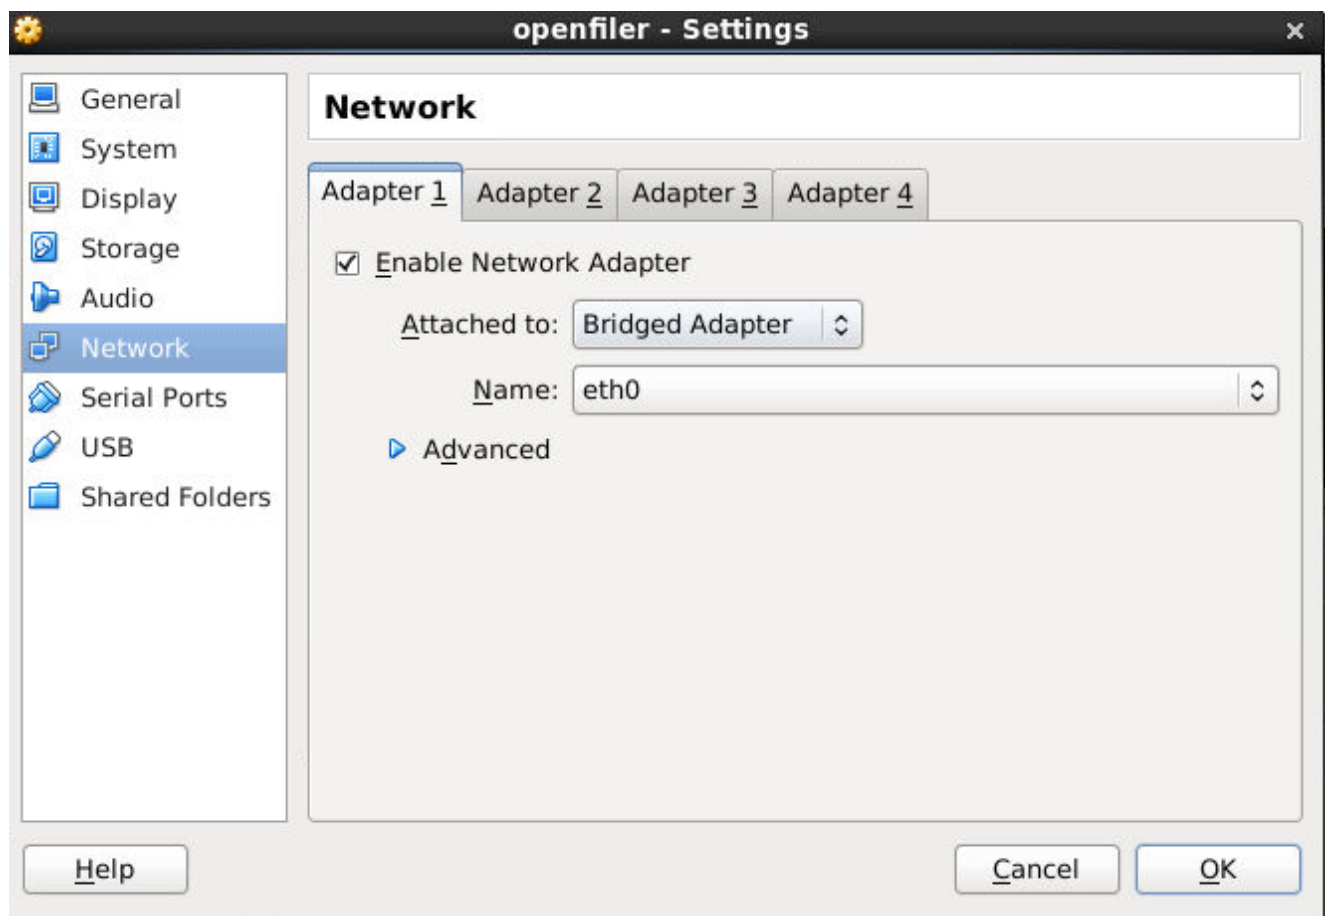




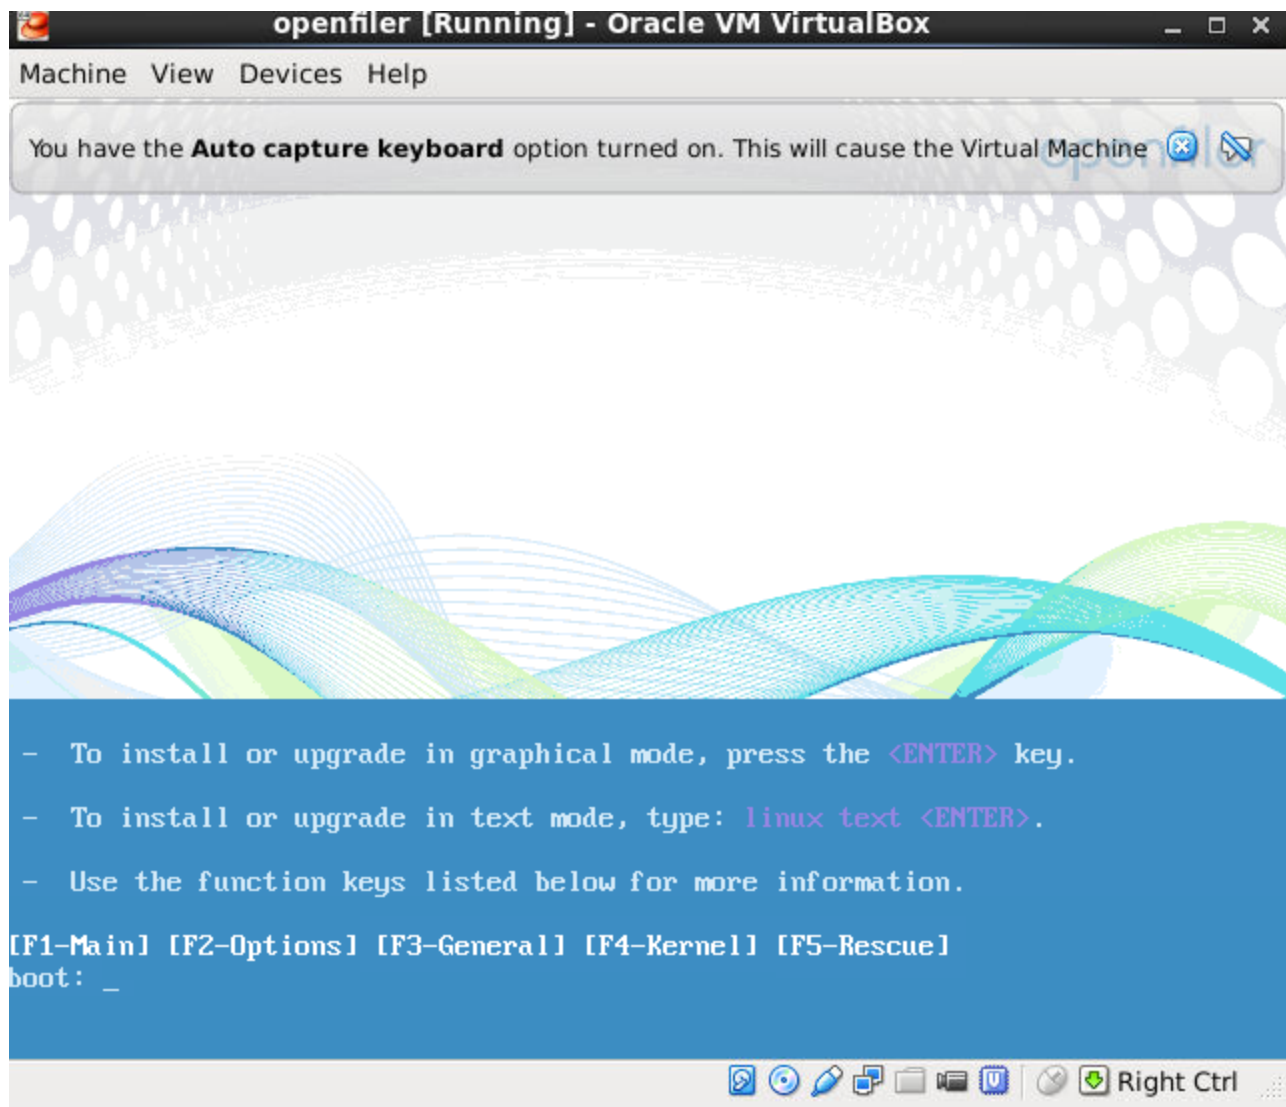




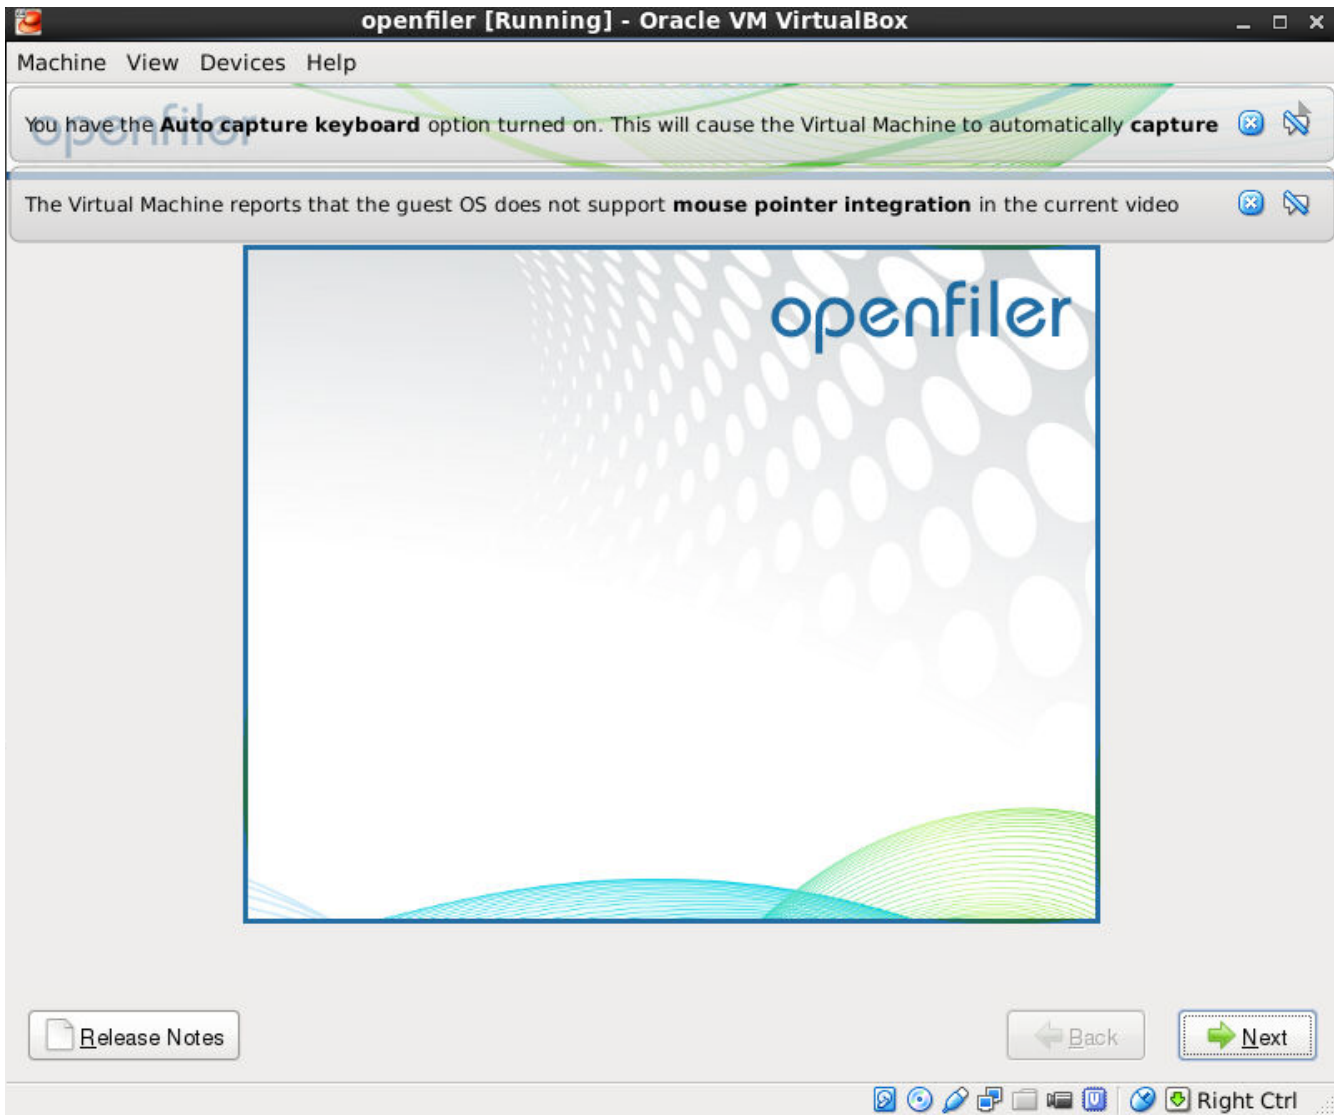
Click Network

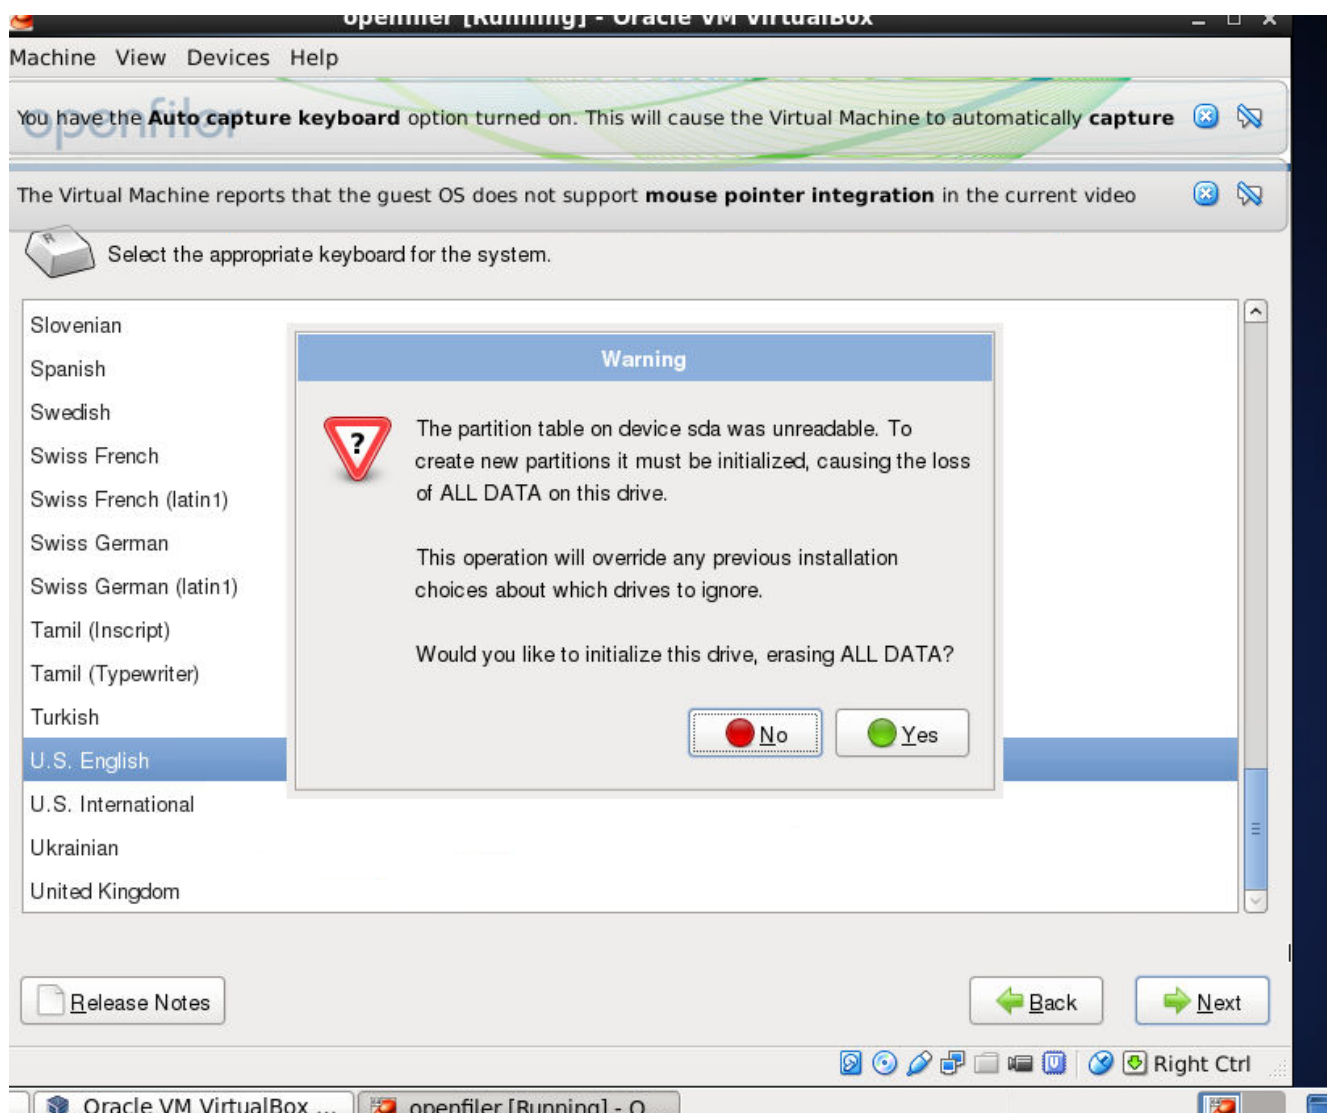


Click start

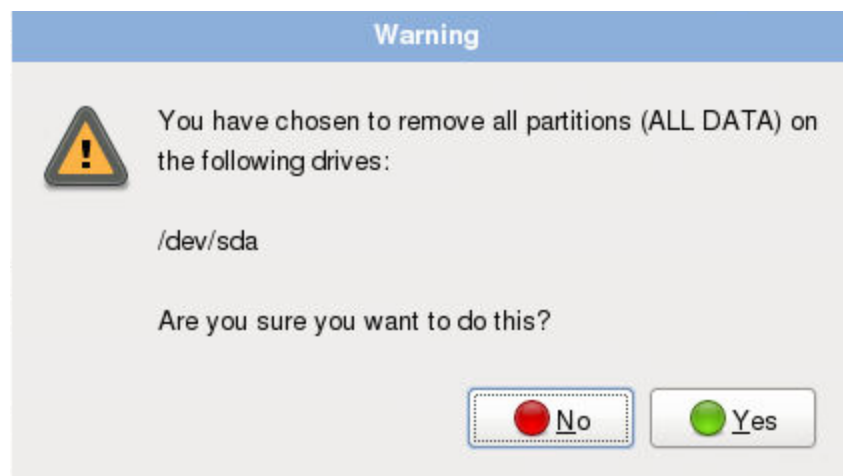
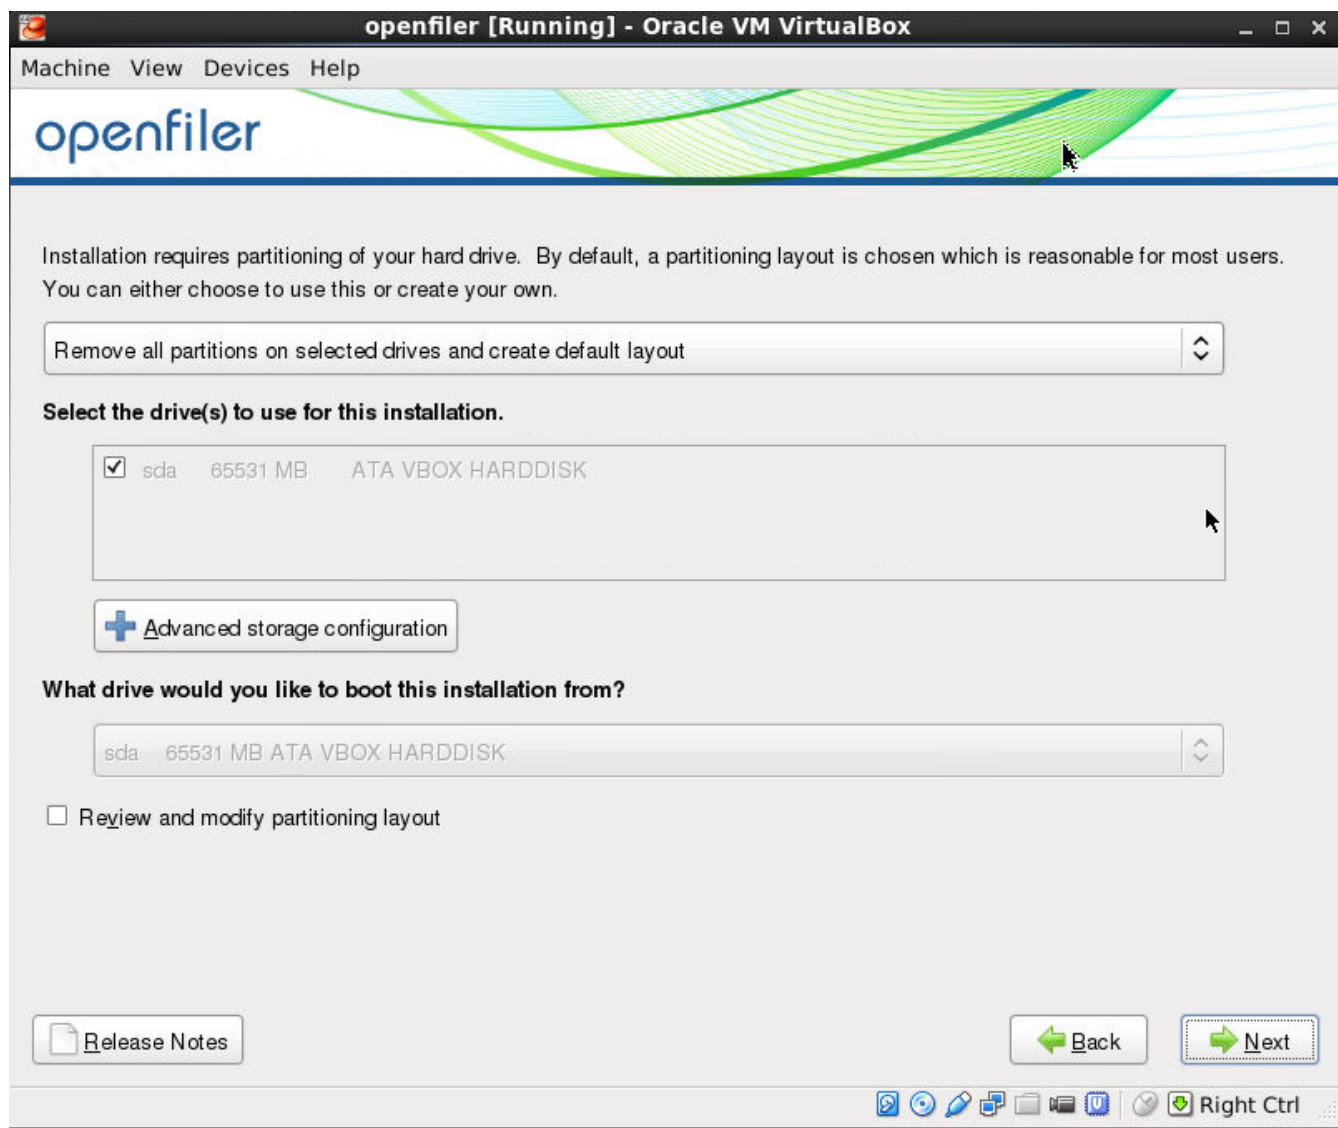


Hit Enter

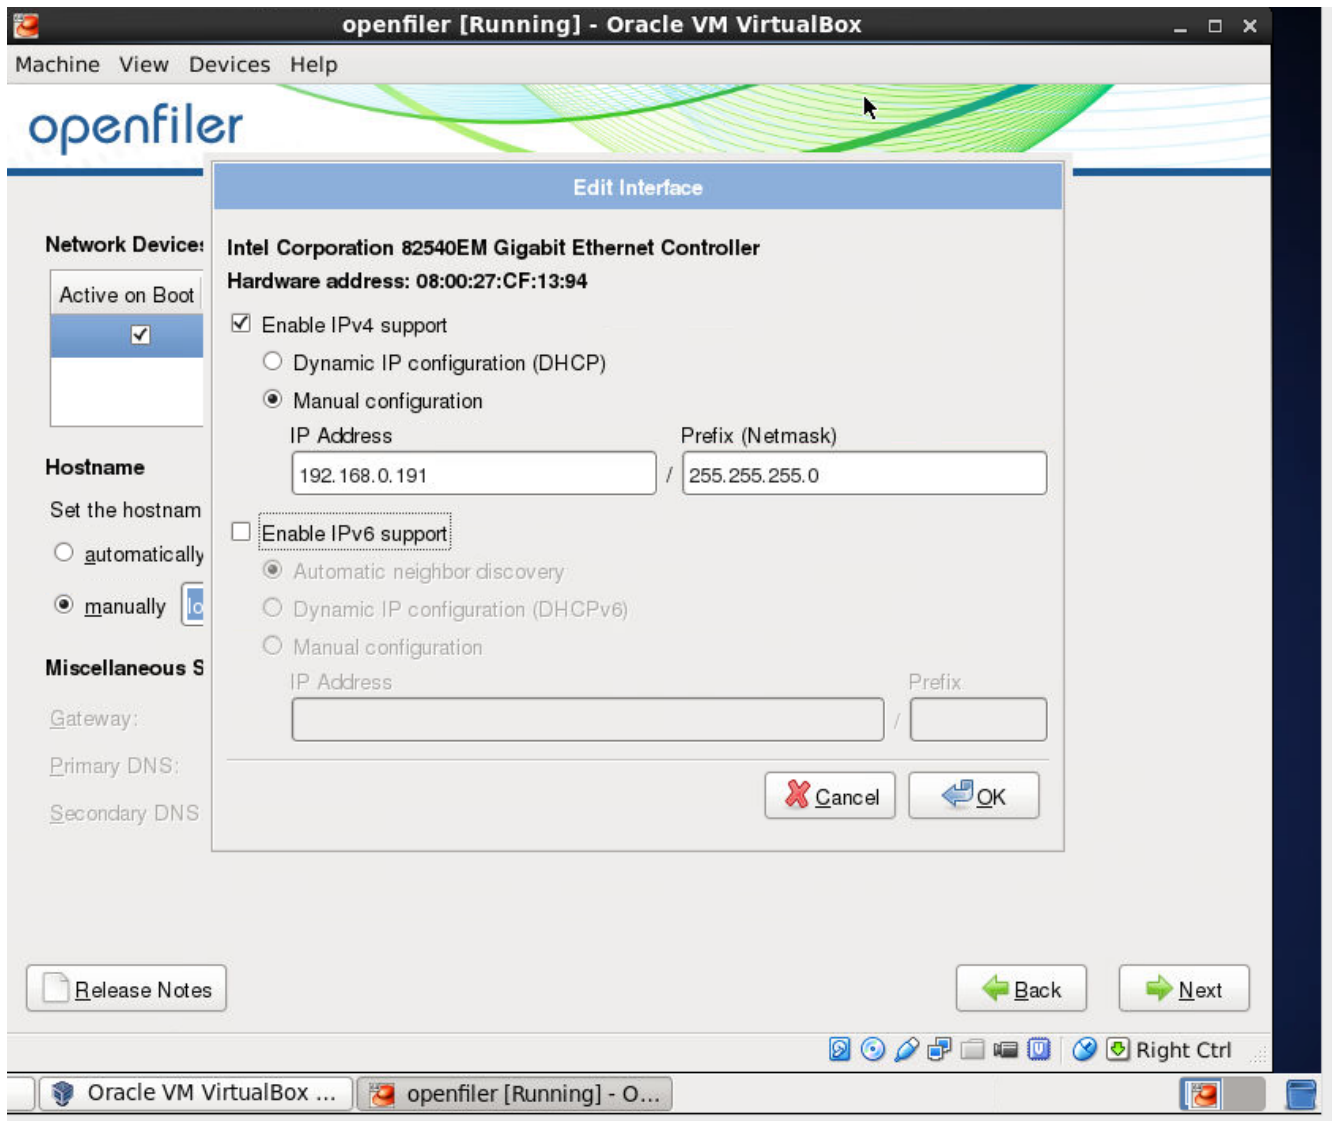


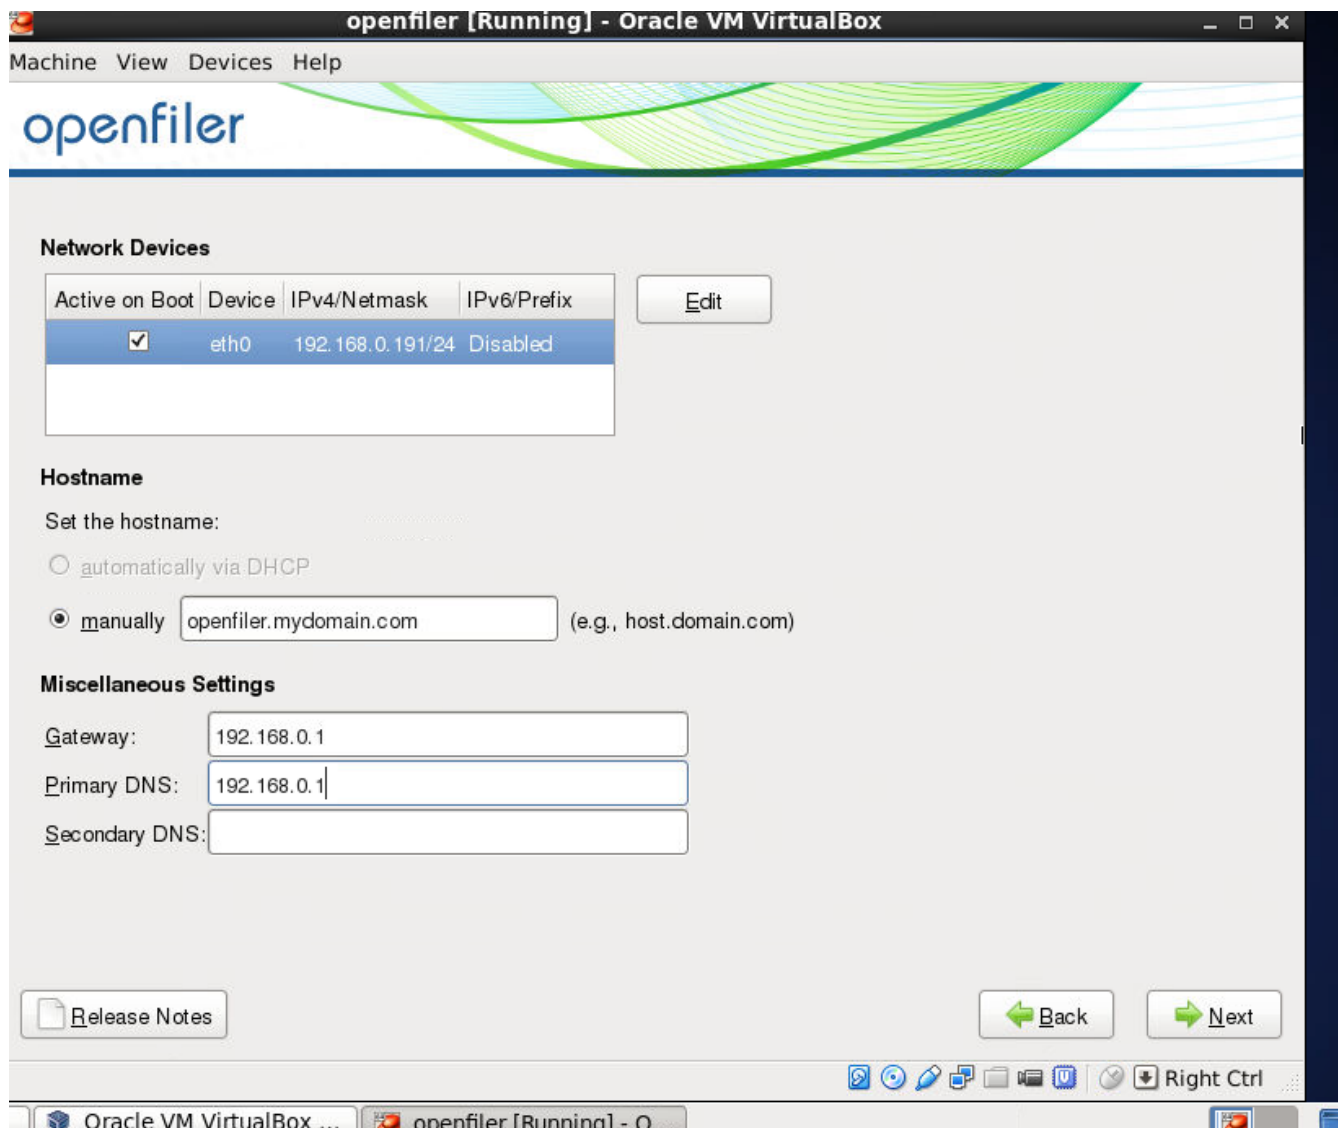


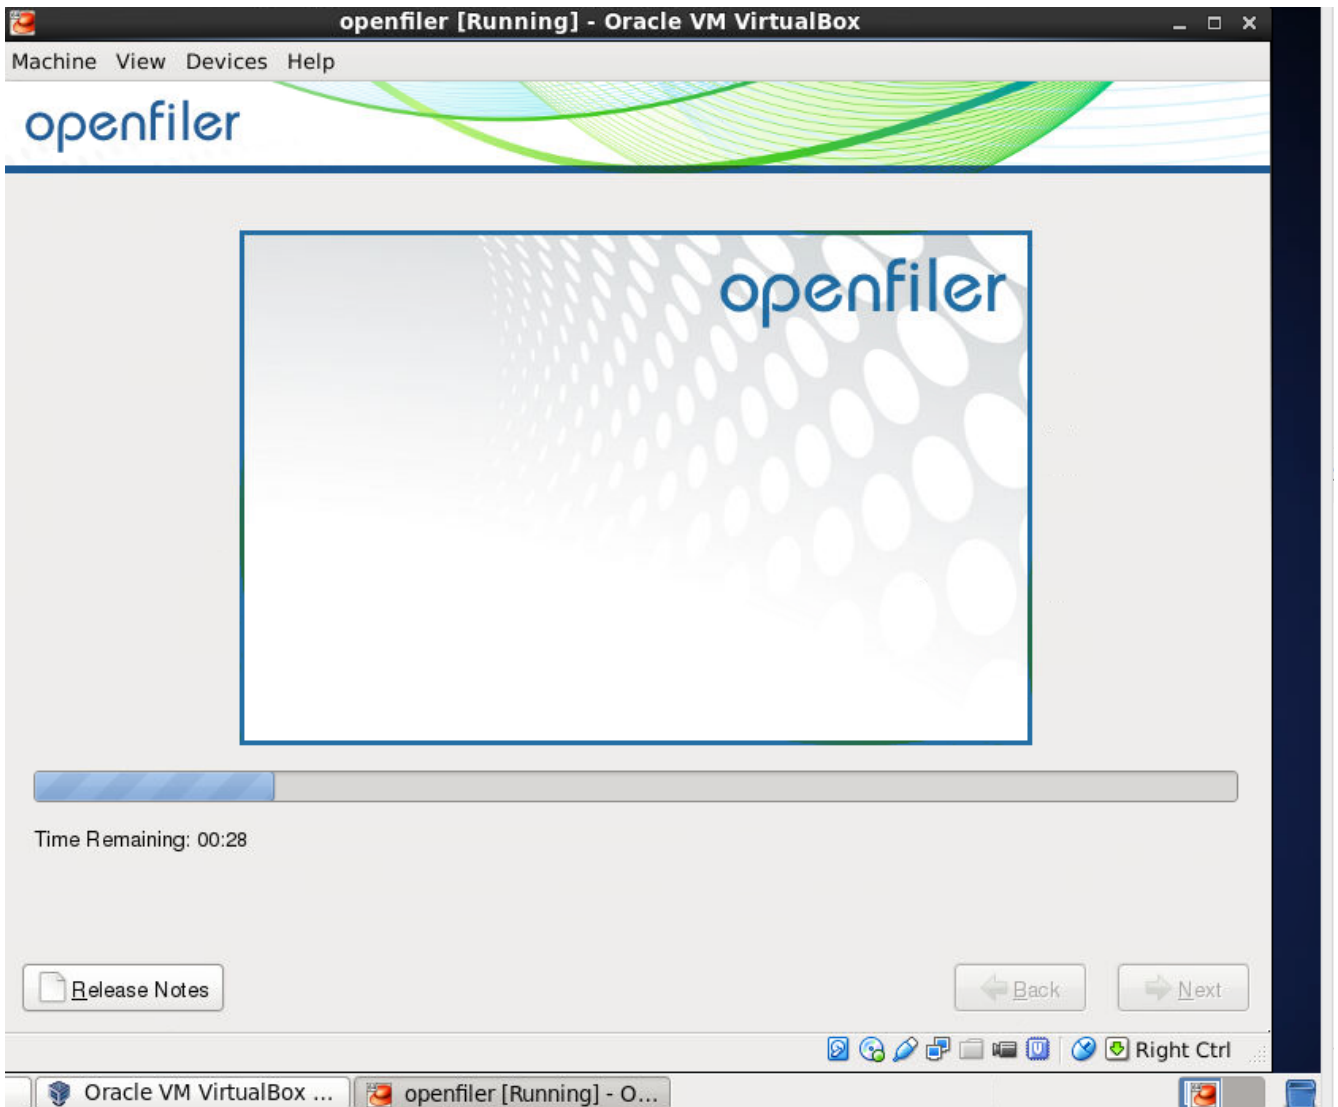


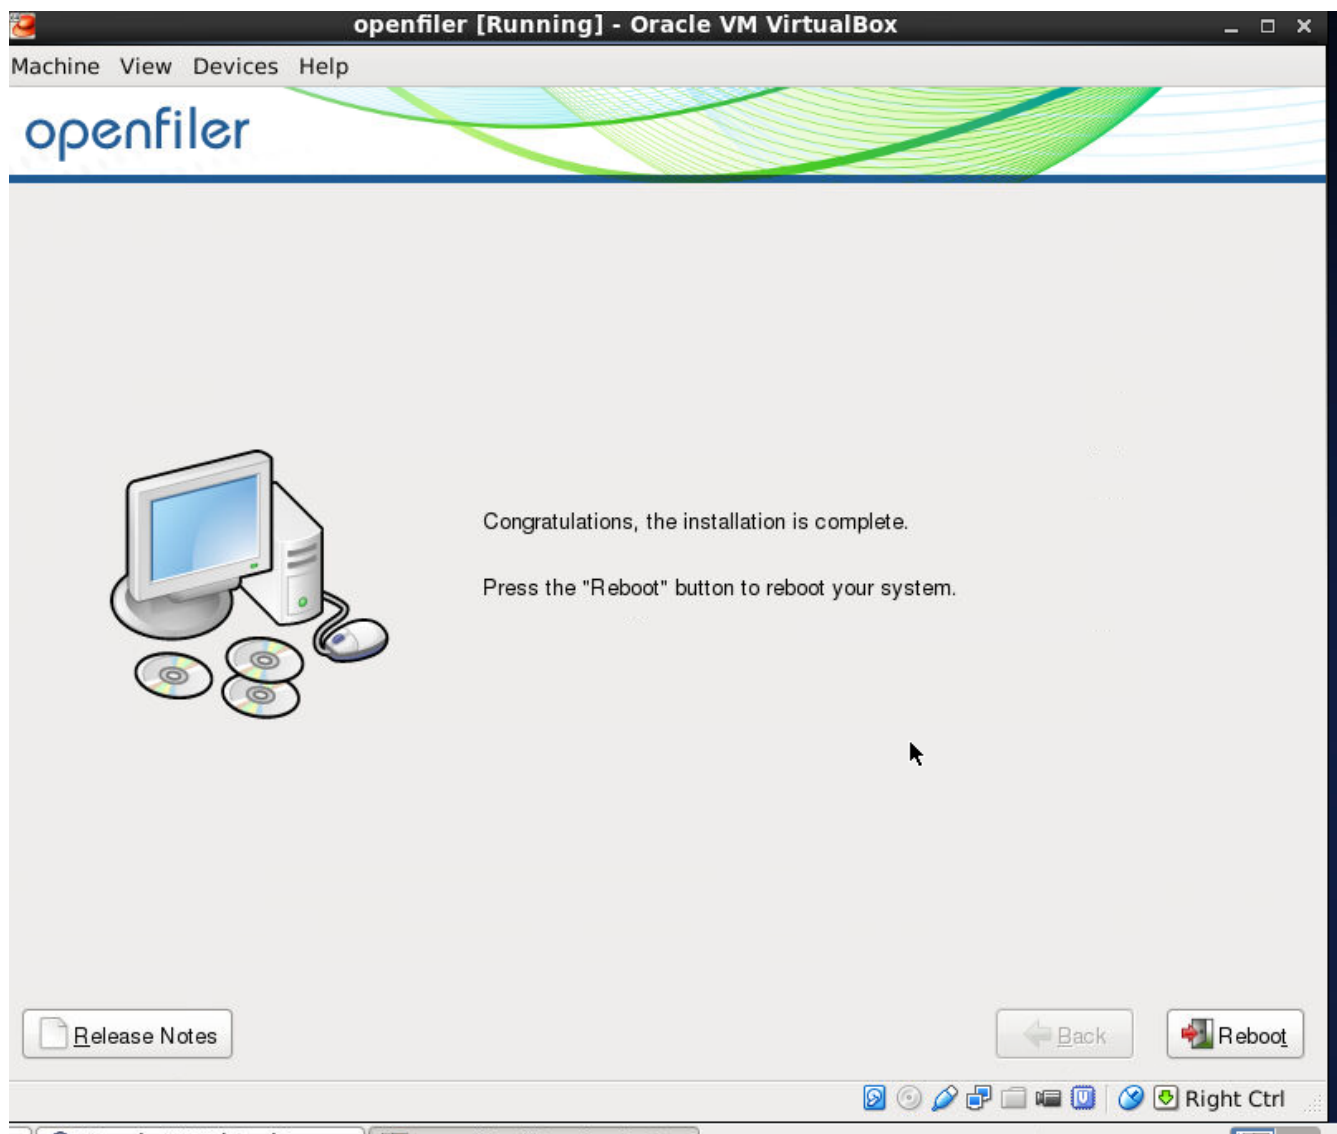


Select Yes

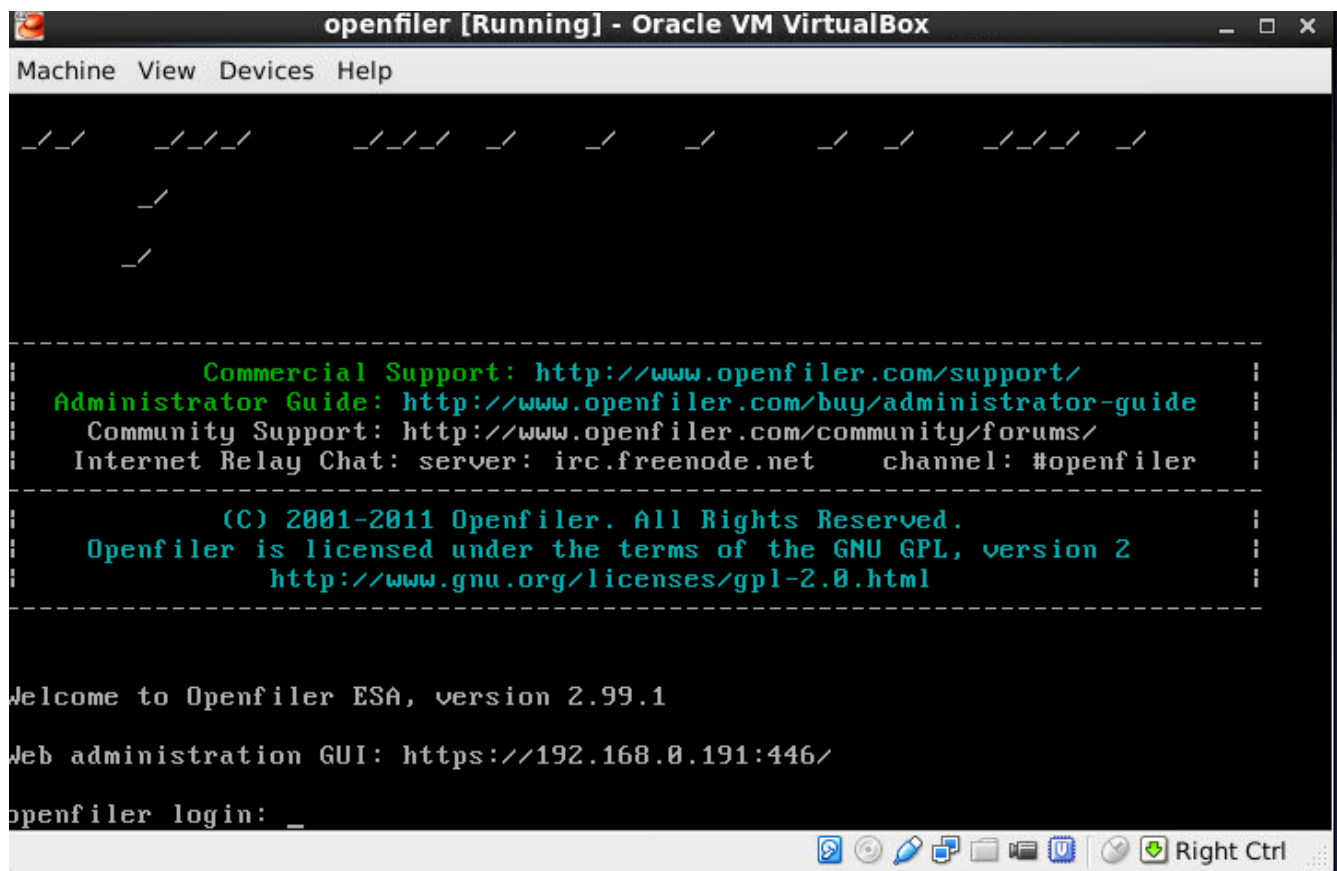




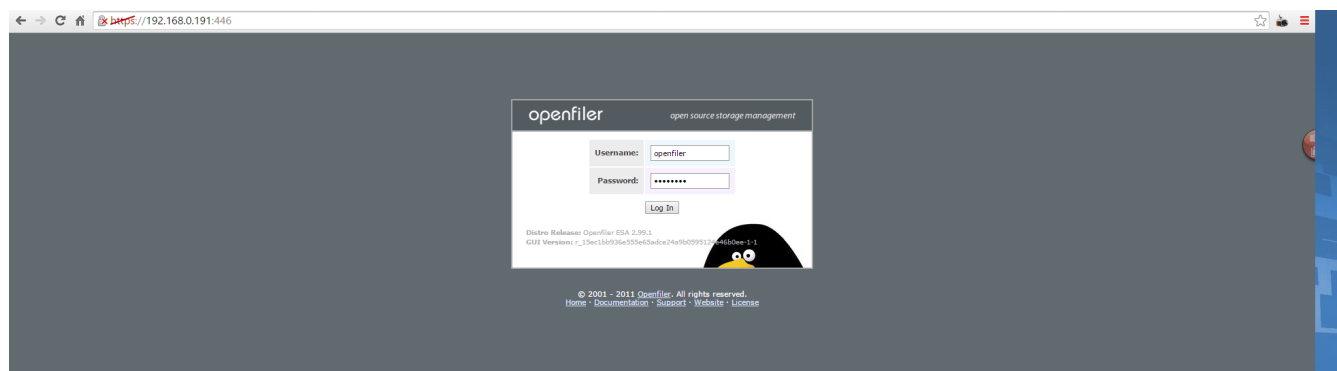




Reboot



Open a browser and go to <https://192.168.0.191:446>



username is "openfiler"  
password is "password"

Manage Services

Service	Boot Status	Modify Boot	Current Status	Start / Stop
CIFS Server	Disabled	<a href="#">Enable</a>	Stopped	<a href="#">Start</a>
NFS Server	Disabled	<a href="#">Enable</a>	Stopped	<a href="#">Start</a>
RSync Server	Disabled	<a href="#">Enable</a>	Stopped	<a href="#">Start</a>
HTTP/Dav Server	Disabled	<a href="#">Enable</a>	Running	<a href="#">Stop</a>
LDAP Container	Disabled	<a href="#">Enable</a>	Stopped	<a href="#">Start</a>
FTP Server	Disabled	<a href="#">Enable</a>	Stopped	<a href="#">Start</a>
iSCSI Target	Disabled	<a href="#">Enable</a>	Stopped	<a href="#">Start</a>
UPS Manager	Disabled	<a href="#">Enable</a>	Stopped	<a href="#">Start</a>
UPS Monitor	Disabled	<a href="#">Enable</a>	Stopped	<a href="#">Start</a>
iSCSI Initiator	Disabled	<a href="#">Enable</a>	Stopped	<a href="#">Start</a>
ACFT Daemon	Enabled	<a href="#">Disable</a>	Running	<a href="#">Stop</a>
SCSI Target	Disabled	<a href="#">Enable</a>	Stopped	<a href="#">Start</a>
FC Target	Disabled	<a href="#">Enable</a>	Stopped	<a href="#">Start</a>
Cluster Manager	Disabled	<a href="#">Enable</a>	Stopped	<a href="#">Start</a>

Services section

- Manage Services
- SMB/CIFS Setup
- LDAP Setup
- UPS Setup
- RSync Setup
- iSCSI Target Setup
- FTP Setup

Support resources

- Report bug
- Get support
- Forums
- Admin Guide

Enable the services:

iSCSI Target    Enabled    [Disable](#)    Running    [Stop](#)

Now power it down and add a drive to it. This drive will be the shared storage for our cluster. We will give it 40GB. Please observe below:

openfiler [Running] - Oracle VM VirtualBox

Machine View Devices Help

```

Commercial Support: http://www.openfiler.com/support/
Administrator Guide: http://www.openfiler.com/buy/administrator-guide
Community Support: http://www.openfiler.com/community/forums/
Internet Relay Chat: server: irc.freenode.net    channel: #openfiler

(C) 2001-2011 Openfiler. All Rights Reserved.
Openfiler is licensed under the terms of the GNU GPL, version 2
http://www.gnu.org/licenses/gpl-2.0.html

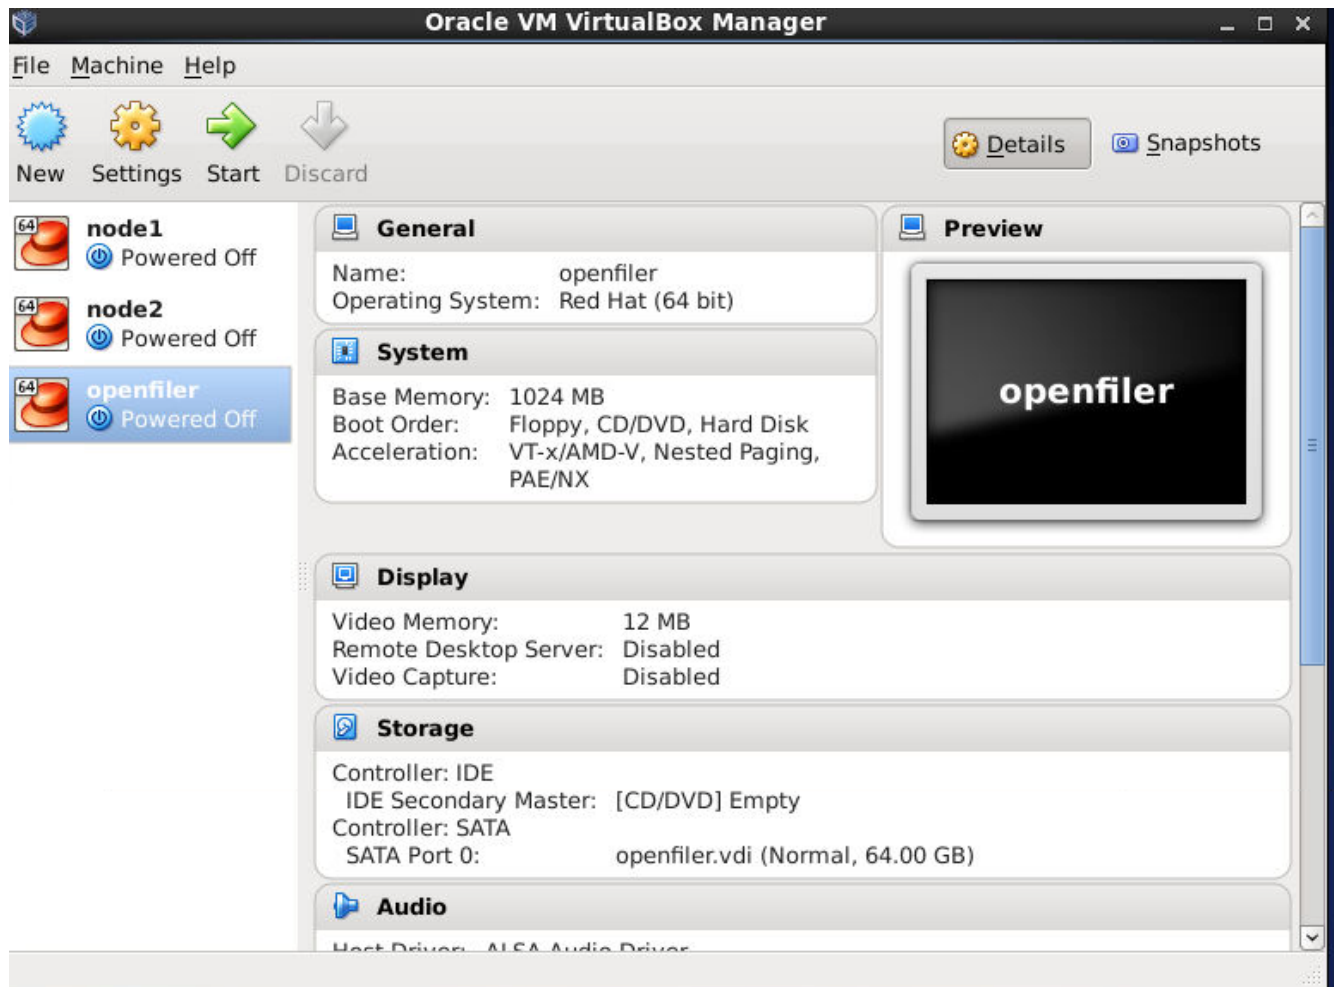
Welcome to Openfiler ESA, version 2.99.1
Web administration GUI: https://192.168.0.191:446/

openfiler login: root
Password:
[root@openfiler ~]# init 0_

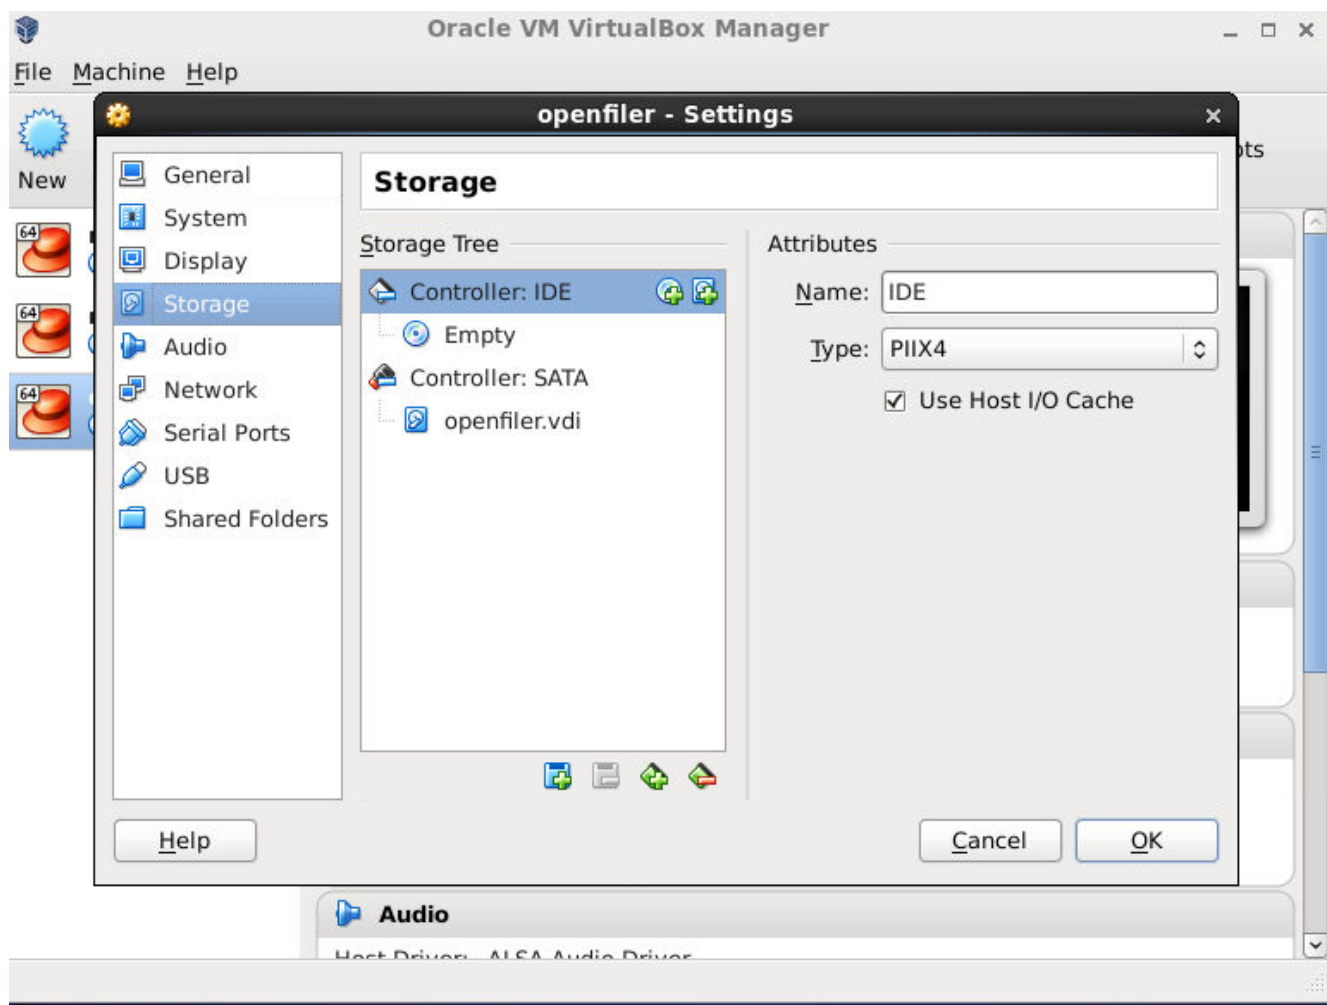
```

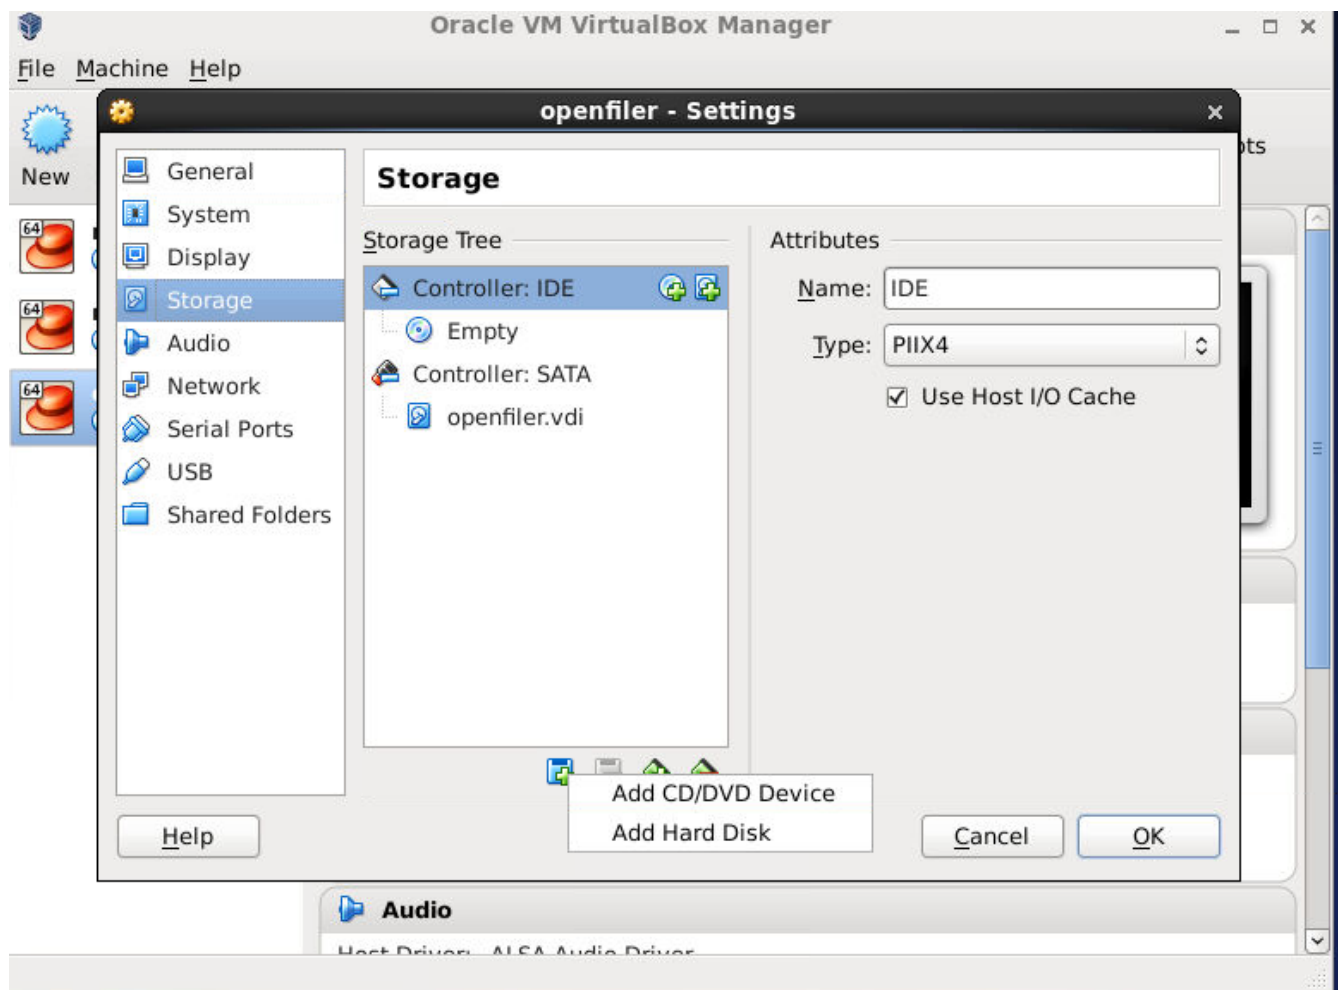
Right Ctrl

Highlight the openfiler VM:

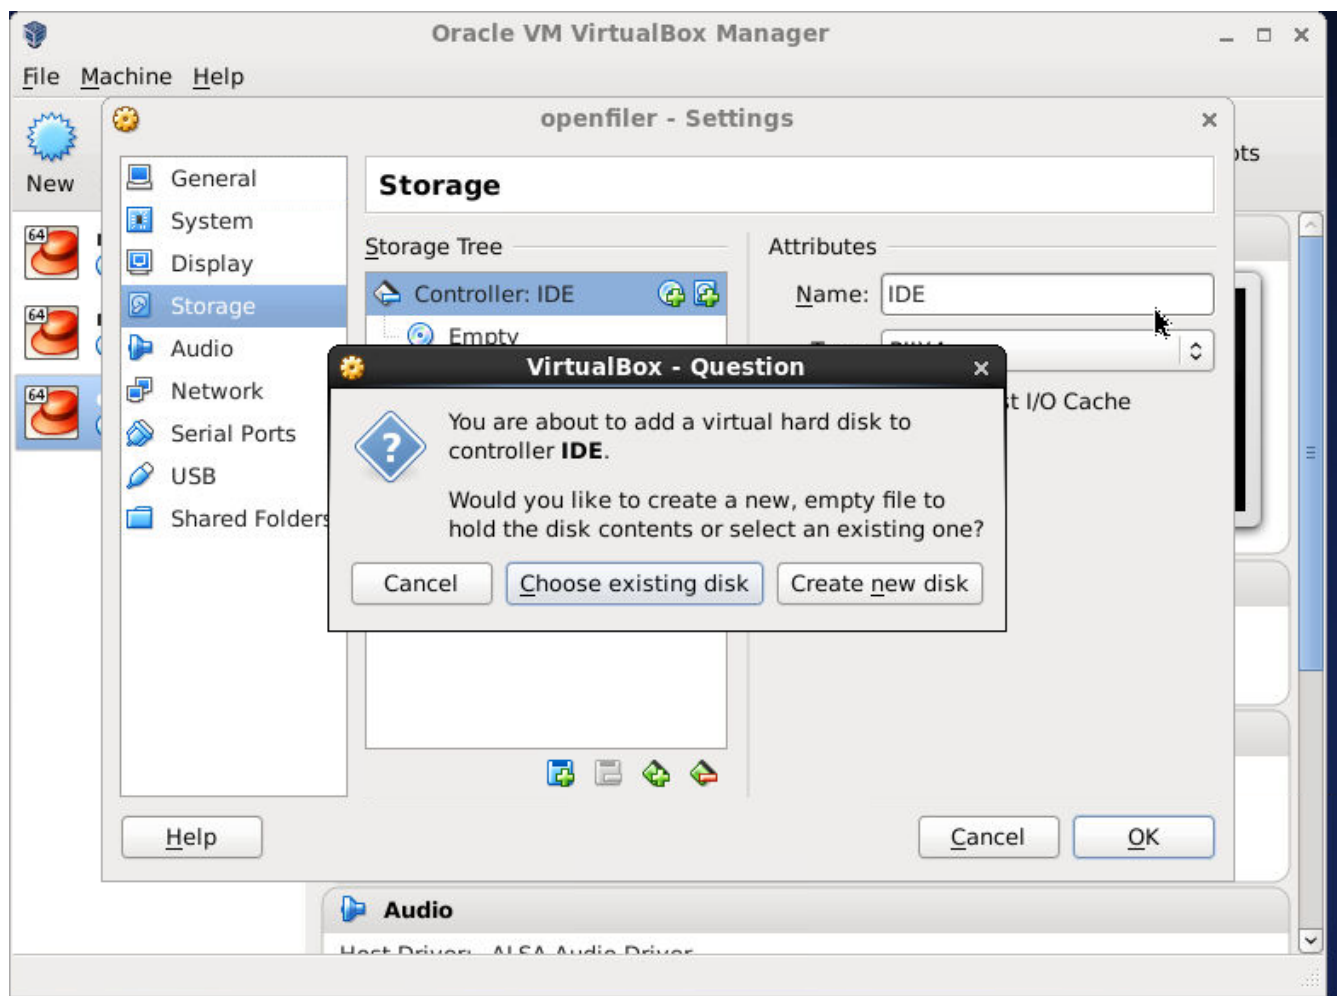








Click “Add Hard Disk”

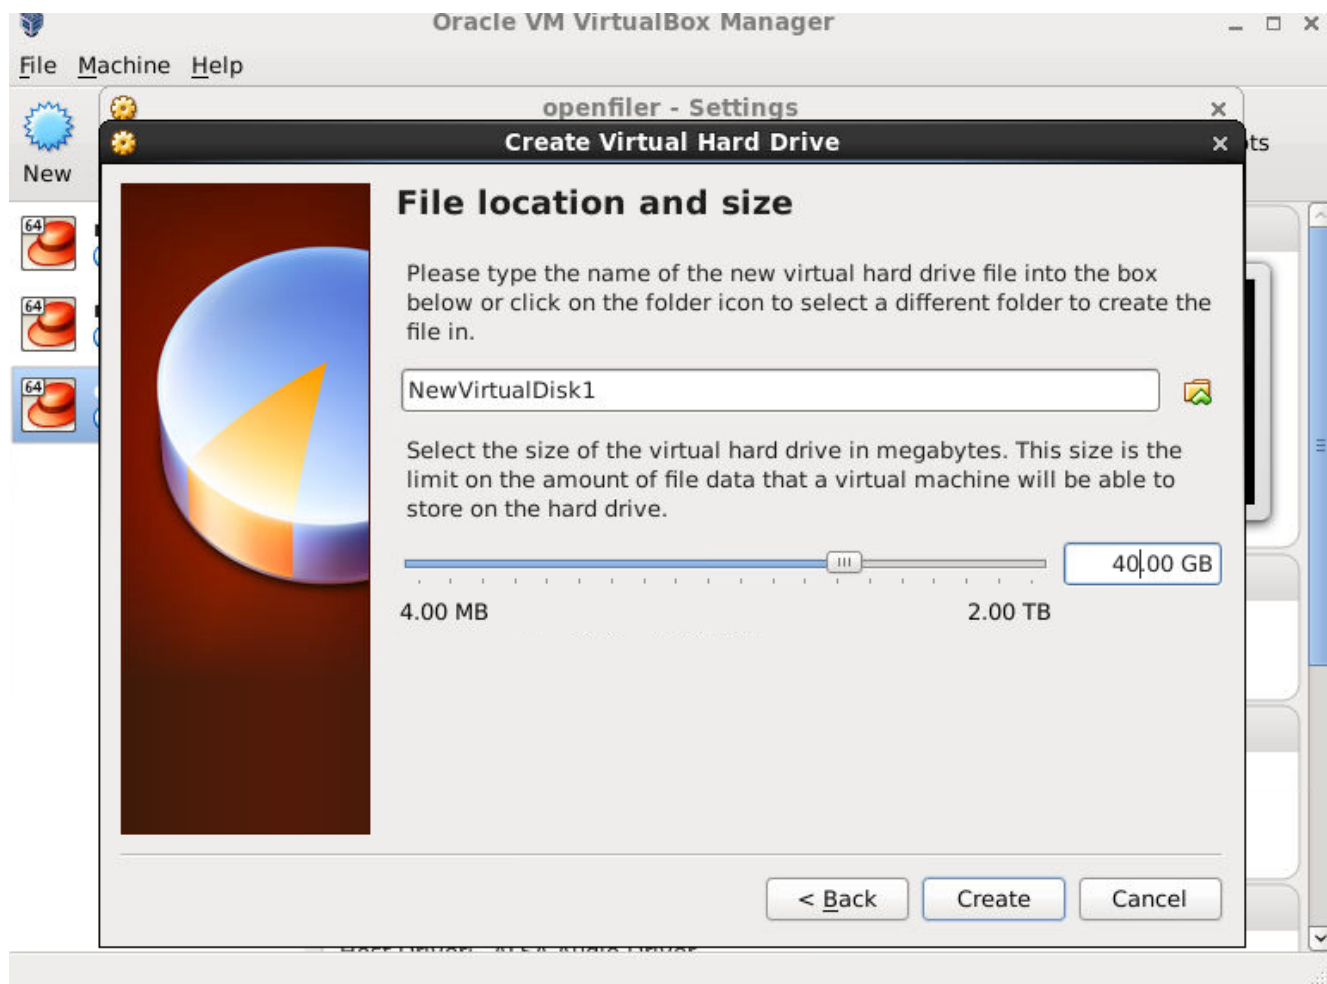


Create New Disk




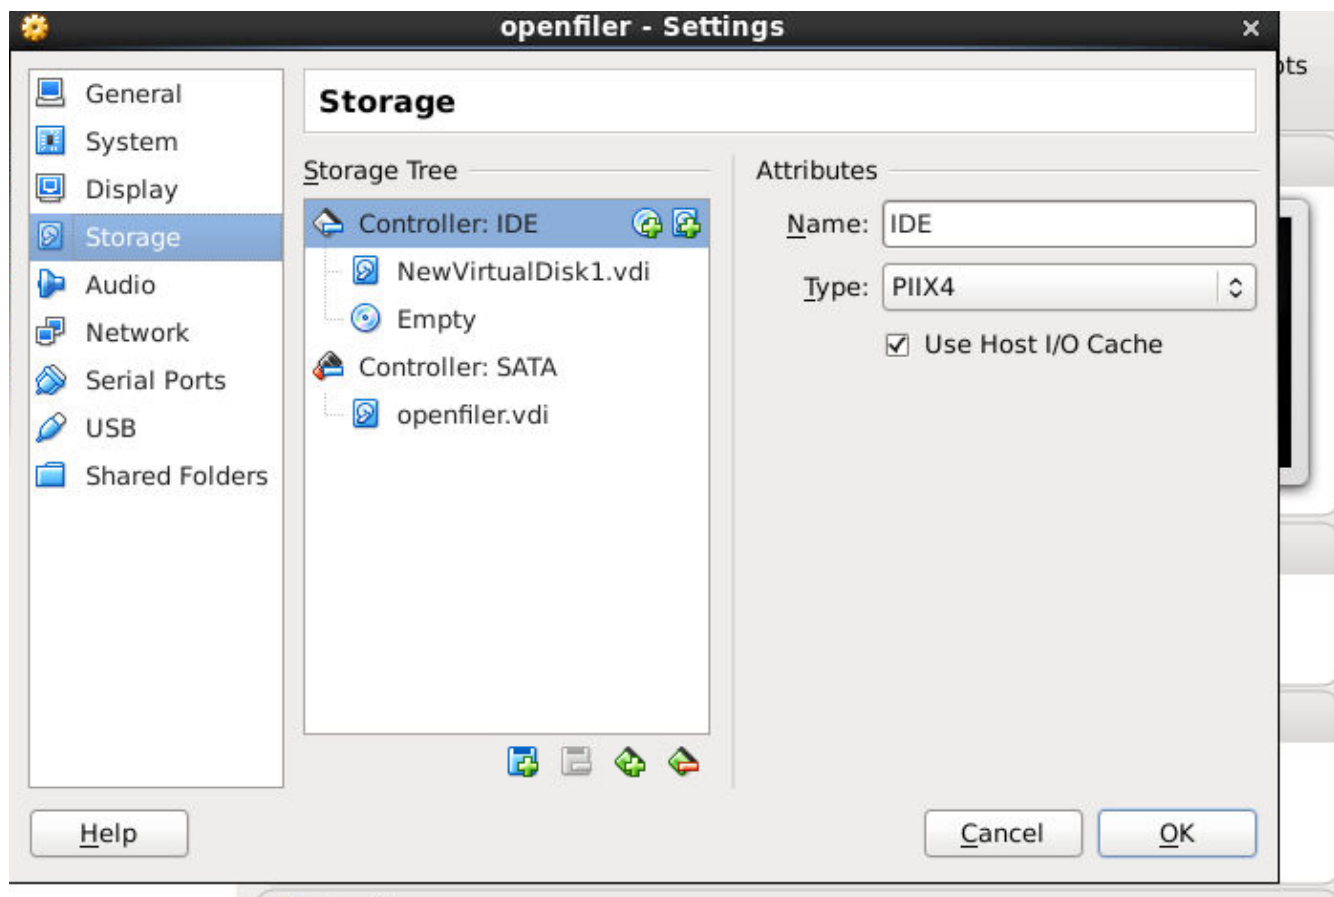
Please select “Fixed Disk” see below:





Click “Create”





Please click “Ok”

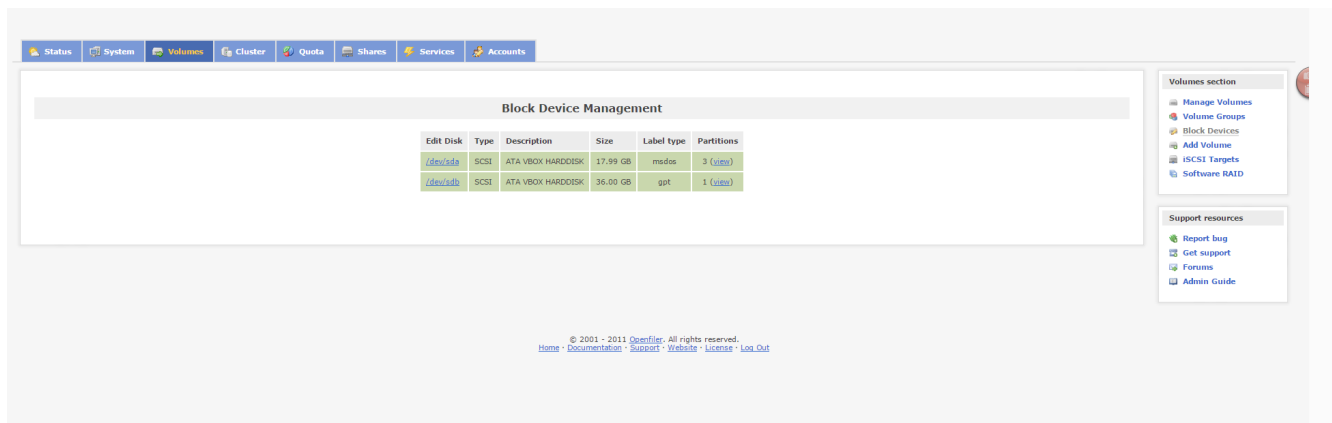
As you may observe an additional drive has been allocated. Now we will configure this as an iSCSI target in openfiler.

NOTE: You may do well to study up on:

- 1)Hardware virtualization
- 2)Virtualization software
- 3)Virtual Machines
- 4)RHEL/CentOS network startup scripts
- 5)iSCSI
- 6)Oracle database
- 7)VCS!!!

Let us start openfiler once more

go to <https://192.168.0.191:446>



Click on volumes tab and then on the right side on “block devices”

NOTE: PLEASE DO NOT TOUCH THE DISK WITH THREE PARTITIONS, THAT IS THE OPENFILER BOOT DRIVE Now, we click on /dev/sdb

NOTE: PLEASE DO NOT TOUCH THE DISK WITH THREE PARTITIONS, THAT IS THE OPENFILER BOOT DRIVE Now, we click on /dev/sdb

Click Create

FROM RIGHT PLANE, CLICK ON “VOLUME GROUPS”



## Volume Group Management

Volume Group Name	Size	Allocated	Free	Members	Add physical storage	Delete VG
-------------------	------	-----------	------	---------	----------------------	-----------

### Create a new volume group



Valid characters for volume group name: **A-Z a-z 0-9 \_ + -**

Volume group name (no spaces)

Select physical volumes to add

☐

/dev/sdb1

34.33 GB

Add volume group

give a volume group name: VOL0 AND CHECK /DEV/SDB1

## Volume Group Management

Volume Group Name	Size	Allocated	Free	Members	Add physical storage	Delete VG
vol0	34.31 GB	0 bytes	34.31 GB	<a href="#">View member PVs</a>	All PVs are used	<a href="#">Delete</a>

IN RIGHT PLANE SELECT “ADD VOLUME”


Please select a volume group to create a volume in.

vol0

Change

Block storage statistics for volume group "vol0"

Total Space	Used Space	Free Space
35979264 bytes (35136 MB)	0 bytes (0 MB)	35979264 bytes (35136 MB)



Create a volume in "vol0"

Volume Name (\*no spaces\*. Valid characters [a-z,A-Z,0-9]):

lun0

Volume Description:

Required Space (MB):

35136

Filesystem / Volume type:

block (iSCSI,FC,etc)

Create

NOTE: we have in vol0, created a volume named lun0, maxed out on the space allocated to it and chosen iSCSI as the Filesystem/Volume type

Click Create

Select Volume Group

Please select a volume group to display.

vol0

Change

Volumes in volume group "vol0" (35136 MB)

A pie chart representing the storage usage of volume group "vol0". The chart is a solid dark green circle, indicating that 100% of the space is used by the volume "lun0". A thin white line marks the 12 o'clock position. Below the chart, the text "lun0 (100%)" is displayed.

Volume name	Volume description	Volume size	File system type	File system size	FS used space	FS free space	Delete	Properties	Snapshots
lun0	lun0	35136 MB	iSCSI	Not applicable	Not applicable	Not applicable	Delete	Edit	Create

0 MB allocated to snapshots

0 MB of free space left

M

Vi

BI

Au

iS

Si

Supp

Rb

Gr

Fr

Ar

NOW GO TO RIGHT PLANE AND SELECT ISCSI TARGETS

StatusSystemVolumesClusterQuotaSharesServicesAccounts

Target ConfigurationLUN MappingNetwork ACLCHAP Authentication

Add new iSCSI Target

Target IQN

Add

iqn.2006-01.com.openfiler:tsn.a12c203227bc

Add

Select iSCSI Target

Please select an iSCSI target to display and/or edit.

iqn.2006-01.com.openfiler:tsn.57b9916dc11a

Change

Settings for target: iqn.2006-01.com.openfiler:tsn.57b9916dc11a

Target Attribute

Attribute Value

Volumes section

Manage Volumes

Volume Groups

Block Devices

Add Volume

iSCSI Targets

Software RAID

Support resources

Report bug

Get support

Forums

Admin Guide

Please select the tab called “LUN Mapping”

Target ConfigurationLUN MappingNetwork ACLCHAP Authentication

LUNs mapped to target: iqn.2006-01.com.openfiler:tsn.57b9916dc11a

No LUNs mapped to this target

Map New LUN to Target: "iqn.2006-01.com.openfiler:tsn.57b9916dc11a"

Name	LUN Path	R/W Mode	SCSI Serial No.	SCSI Id.	Transfer Mode	Map LUN
lun0	/dev/vol0/lun0	write-thru	sTr1uI-n74R-C7ya	sTr1uI-n74R-C7ya	blockio	Map

Click on “Map”

Target ConfigurationLUN MappingNetwork ACLCHAP Authentication

LUNs mapped to target: iqn.2006-01.com.openfiler:tsn.57b9916dc11a

LUN Id.	LUN Path	R/W Mode	SCSI Serial No.	SCSI Id.	Transfer Mode	Unmap LUN
0	/dev/vol0/lun0	write-thru	sTr1uI-n74R-C7ya	sTr1uI-n74R-C7ya	blockio	Unmap

All mappable iSCSI LUNs for this target have been mapped.

If the steps outlined here are not clear, then, please refer to this web page:

<http://dl.faraznetwork.ir/Files/Openfiler%20Configuration.pdf>

You may ignore the section on “Configuring Network Settings”.

Now start node1 and node 2 and on both nodes launch putty sessions to them. On EACH node do this:

```
[root@node1 ~]# yum install iscsi-initiator-utils.x86_64
```

On any node, say, node2, do this:

```
[root@node2 ~]# service iscsid start
[root@node2 ~]# chkconfig iscsid on
[root@node2 ~]# chkconfig iscsi on
[root@node2 ~]# iscsiadm -m discovery -t sendtargets -p 192.168.0.191
192.168.0.191:3260,1 iqn.2006-01.com.openfiler:tsn.d5db5d22c74d
192.168.0.191:3260,1 iqn.2006-01.com.openfiler:tsn.e5b215212686
192.168.0.191:3260,1 iqn.2006-01.com.openfiler:tsn.57b9916dc11a
192.168.0.191:3260,1 iqn.2006-01.com.openfiler:tsn.d3e4680d2625
192.168.0.191:3260,1 iqn.2006-01.com.openfiler:tsn.b795fec580c1
```

```
[root@node2 ~]# iscsiadm -m node -T iqn.2006-01.com.openfiler:tsn.d5db5d22c74d -p 192.168.0.191
--login
```

NOTE: You may use any of the identifiers, your's will be different than mine.

Logging in to [iface: default, target: iqn.2006-01.com.openfiler:tsn.d5db5d22c74d, portal: 192.168.0.191,3260] (multiple)

Login to [iface: default, target: iqn.2006-01.com.openfiler:tsn.d5db5d22c74d, portal: 192.168.0.191,3260] successful.

```
[root@node2 ~]# fdisk -l
```

```
Disk /dev/sda: 68.7 GB, 68719476736 bytes
255 heads, 63 sectors/track, 8354 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x000c60df
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1	*	1	64	512000	83	Linux

Partition 1 does not end on cylinder boundary.

/dev/sda2		64	8355	66595840	8e	Linux LVM
-----------	--	----	------	----------	----	-----------

```
Disk /dev/mapper/vg_node1-lv_root: 53.7 GB, 53687091200 bytes
255 heads, 63 sectors/track, 6527 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000000
```

Disk /dev/mapper/vg\_node1-lv\_swap: 4227 MB, 4227858432 bytes  
255 heads, 63 sectors/track, 514 cylinders  
Units = cylinders of 16065 \* 512 = 8225280 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x00000000

Disk /dev/mapper/vg\_node1-lv\_home: 10.3 GB, 10276044800 bytes  
255 heads, 63 sectors/track, 1249 cylinders  
Units = cylinders of 16065 \* 512 = 8225280 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x00000000

Disk /dev/sdb: 36.8 GB, 36842766336 bytes  
64 heads, 32 sectors/track, 35136 cylinders  
Units = cylinders of 2048 \* 512 = 1048576 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x00000000

[root@node2 ~]#

As you can see we now have our SAN(openfiler) disk as /de/sdb

Doing the same on node1 shows us:

```
[root@node1 ~]# service iscsid start
[root@node1 ~]# chkconfig iscsid on
[root@node1 ~]# chkconfig iscsi on
[root@node1 ~]# iscsiadm -m discovery -t sendtargets -p 192.168.0.191
192.168.0.191:3260,1 iqn.2006-01.com.openfiler:tsn.d5db5d22c74d
192.168.0.191:3260,1 iqn.2006-01.com.openfiler:tsn.e5b215212686
192.168.0.191:3260,1 iqn.2006-01.com.openfiler:tsn.57b9916dc11a
192.168.0.191:3260,1 iqn.2006-01.com.openfiler:tsn.d3e4680d2625
192.168.0.191:3260,1 iqn.2006-01.com.openfiler:tsn.b795fec580c1
[root@node1 ~]# iscsiadm -m node -T iqn.2006-01.com.openfiler:tsn.d5db5d22c74d -p 192.168.0.191
--login
Logging in to [iface: default, target: iqn.2006-01.com.openfiler:tsn.d5db5d22c74d, portal:
192.168.0.191,3260] (multiple)
Login to [iface: default, target: iqn.2006-01.com.openfiler:tsn.d5db5d22c74d, portal:
192.168.0.191,3260] successful.
[root@node1 ~]# fdisk -l
```

Disk /dev/sda: 68.7 GB, 68719476736 bytes  
255 heads, 63 sectors/track, 8354 cylinders  
Units = cylinders of 16065 \* 512 = 8225280 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x000c60df

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1	*	1	64	512000	83	Linux

Partition 1 does not end on cylinder boundary.

/dev/sda2		64	8355	66595840	8e	Linux LVM
-----------	--	----	------	----------	----	-----------

Disk /dev/mapper/vg\_node1-lv\_root: 53.7 GB, 53687091200 bytes  
255 heads, 63 sectors/track, 6527 cylinders  
Units = cylinders of 16065 \* 512 = 8225280 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x00000000

Disk /dev/mapper/vg\_node1-lv\_swap: 4227 MB, 4227858432 bytes  
255 heads, 63 sectors/track, 514 cylinders  
Units = cylinders of 16065 \* 512 = 8225280 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x00000000

Disk /dev/mapper/vg\_node1-lv\_home: 10.3 GB, 10276044800 bytes  
255 heads, 63 sectors/track, 1249 cylinders  
Units = cylinders of 16065 \* 512 = 8225280 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x00000000

Disk /dev/sdb: 36.8 GB, 36842766336 bytes  
64 heads, 32 sectors/track, 35136 cylinders  
Units = cylinders of 2048 \* 512 = 1048576 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x00000000

[root@node1 ~]#

So people this is the beauty of shared storage, accessible from both nodes (/dev/sdb is accessible from node1 and from node2).

Now, we will format this lun /dev/sdb

From ANY node do this:

```
[root@node2 ~]# fdisk /dev/sdb
```

WARNING: DOS-compatible mode is deprecated. It's strongly recommended to switch off the mode (command 'c') and change display units to sectors (command 'u').

Command (m for help): m

Command action

- a toggle a bootable flag
- b edit bsd disklabel
- c toggle the dos compatibility flag
- d delete a partition
- l list known partition types
- m print this menu
- n add a new partition
- o create a new empty DOS partition table
- p print the partition table
- q quit without saving changes
- s create a new empty Sun disklabel
- t change a partition's system id
- u change display/entry units
- v verify the partition table
- w write table to disk and exit
- x extra functionality (experts only)

Command (m for help): n

Command action

- e extended
- p primary partition (1-4)

p

Partition number (1-4): 1

First cylinder (1-35136, default 1):

Using default value 1

Last cylinder, +cylinders or +size {K,M,G} (1-35136, default 35136):

Using default value 35136

Command (m for help): w

The partition table has been altered!

Calling ioctl() to re-read partition table.

Syncing disks.

```
[root@node2 ~]#
```



```
[root@node2 ~]# fdisk -l
```

Disk /dev/sda: 68.7 GB, 68719476736 bytes  
255 heads, 63 sectors/track, 8354 cylinders  
Units = cylinders of 16065 \* 512 = 8225280 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x000c60df

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1	*	1	64	512000	83	Linux

Partition 1 does not end on cylinder boundary.

/dev/sda2		64	8355	66595840	8e	Linux LVM
-----------	--	----	------	----------	----	-----------

Disk /dev/mapper/vg\_node1-lv\_root: 53.7 GB, 53687091200 bytes  
255 heads, 63 sectors/track, 6527 cylinders  
Units = cylinders of 16065 \* 512 = 8225280 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x00000000

Disk /dev/mapper/vg\_node1-lv\_swap: 4227 MB, 4227858432 bytes  
255 heads, 63 sectors/track, 514 cylinders  
Units = cylinders of 16065 \* 512 = 8225280 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x00000000

Disk /dev/mapper/vg\_node1-lv\_home: 10.3 GB, 10276044800 bytes  
255 heads, 63 sectors/track, 1249 cylinders  
Units = cylinders of 16065 \* 512 = 8225280 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x00000000

Disk /dev/sdb: 36.8 GB, 36842766336 bytes  
64 heads, 32 sectors/track, 35136 cylinders  
Units = cylinders of 2048 \* 512 = 1048576 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0xa958974b

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		1	35136	35979248	83	Linux

```
[root@node2 ~]#
```

So, now we are free to install VCS cluster software on these two nodes and then install Oracle on the shared storage (so that when one node crashes the other node picks up the oracle database).

Kinda late in the day, but your file /etc/sysconfig/network should look like this:

```
[root@node2 yum.repos.d]# cat /etc/sysconfig/network
NETWORKING=yes
HOSTNAME=node2.mydomain.com
GATEWAY=192.168.0.1
[root@node2 yum.repos.d]#
```

Please substitute your router's IP for the gateway, followed by:

```
[root@node1 ~]# vi /etc/sysconfig/network
[root@node1 ~]# service network restart
Shutting down interface eth0: Device state: 3 (disconnected)
                               [ OK ]
Shutting down interface eth1:           [ OK ]
Shutting down interface eth2:           [ OK ]
Shutting down loopback interface:       [ OK ]
Bringing up loopback interface:         [ OK ]
Bringing up interface eth0: Active connection state: activated
Active connection path: /org/freedesktop/NetworkManager/ActiveConnection/1
                               [ OK ]
[root@node1 ~]#
```

Now we will edit the hosts file on both nodes and make entries for the nodes in there. This is what the /etc/hosts file will look on node1

```
[root@node1 ~]# cat /etc/hosts
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1      localhost localhost.localdomain localhost6 localhost6.localdomain6
192.168.0.189 node1.mydomain.com node1
192.168.0.190 node2.mydomain.com node2
```

It should be identical on node2

```
[root@node2 yum.repos.d]# cat /etc/hosts
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
192.168.0.189 node1.mydomain.com node1
192.168.0.190 node2.mydomain.com node2
[root@node2 yum.repos.d]#
```

One more thing, we need to setup the private interconnects for LLT. To do so, we will edit the files:  
/etc/sysconfig/network-scripts/ifcfg-eth1  
/etc/sysconfig/network-scripts/ifcfg-eth1

and give them PRIVATE IP's, something like:

node1: 10.10.10.10 and 10.10.10.11

node2: 10.10.10.12 and 10.10.10.13

NOTE: YOU DO NOT NEED TO ASSIGN IP ADDRESSES TO THE PRIVATE INTERCONNECTS  
PLEASE AVOID WRITING ANY IP ADDRESSES IN THE CORRESPONDING ifcfg-ethX FILES.

Please refer to this [symantec](http://www.symantec.com/connect/forums/configure-llt-and-gab-veritas-cluster) note:

<http://www.symantec.com/connect/forums/configure-llt-and-gab-veritas-cluster>

the files will look like this:

on node 2:

```
[root@node2 yum.repos.d]# cat /etc/sysconfig/network-scripts/ifcfg-eth1
DEVICE=eth1
#HWADDR=08:00:27:61:F4:AE
TYPE=Ethernet
UUID=905581ef-c586-485c-be3c-65bb81e0b52a
ONBOOT=yes
NM_CONTROLLED=no
BOOTPROTO=static
IPADDR=10.10.10.12
[root@node2 yum.repos.d]#
```

```
[root@node2 yum.repos.d]# cat /etc/sysconfig/network-scripts/ifcfg-eth2
DEVICE=eth2
#HWADDR=08:00:27:57:5A:6E
TYPE=Ethernet
UUID=d368dbec-c647-4695-a9de-40a1a9c6f1ef
ONBOOT=yes
NM_CONTROLLED=no
BOOTPROTO=static
IPADDR=10.10.10.13
[root@node2 yum.repos.d]#
```

```
[root@node2 ~]# service NetworkManager stop
```

```
Stopping NetworkManager daemon:
```

```
[ OK ]
```

```
[root@node2 ~]# chkconfig NetworkManager off
```

```
[root@node2 ~]#
```

PLEASE NOTE: I have commented out the MAC/HWADDR and changed NM\_CONTROLLED to no  
Next:

```
[root@node2 yum.repos.d]# ifup eth1
Error: No suitable device found: no device found for connection 'System eth1'.
[root@node2 yum.repos.d]# ifup eth2
Error: No suitable device found: no device found for connection 'System eth2'.
[root@node2 yum.repos.d]# vi /etc/sysconfig/network-scripts/ifcfg-eth2
[root@node2 yum.repos.d]#
```

OOPS we have an issue, never mind, we will resolve it:

```
[root@node2 yum.repos.d]# cd /etc/udev/rules.d
[root@node2 rules.d]# ls -al
total 52
drwxr-xr-x. 2 root root 4096 Apr  1 20:27 .
drwxr-xr-x. 4 root root 4096 Apr  1 20:19 ..
-rw-r--r--. 1 root root 1652 Aug 25 2010 60-fprint-autosuspend.rules
-rw-r--r--. 1 root root 1060 Jun 29 2010 60-pcmcia.rules
-rw-r--r--. 1 root root  316 Aug 11 2014 60-raw.rules
-rw-r--r--. 1 root root  530 Apr  1 20:27 70-persistent-cd.rules
-rw-r--r--. 1 root root 1245 Apr  1 20:56 70-persistent-net.rules
-rw-r--r--. 1 root root  40 Sep  9 2014 80-kvm.rules
-rw-r--r--. 1 root root  320 Jun 25 2014 90-alsa.rules
-rw-r--r--. 1 root root  83 Jun 19 2014 90-hal.rules
-rw-r--r--. 1 root root 2486 Jun 30 2010 97-bluetooth-serial.rules
-rw-r--r--. 1 root root  308 Aug 26 2014 98-kexec.rules
-rw-r--r--. 1 root root  54 Nov  3 2011 99-fuse.rules
[root@node2 rules.d]# rm -f 70-persistent-net.rules
[root@node2 rules.d]# reboot
```

TRY AGAIN

On node1

```
[root@node1 ~]# cat /etc/sysconfig/network-scripts/ifcfg-eth1
DEVICE=eth1
#HWADDR=08:00:27:61:F4:AE
TYPE=Ethernet
UUID=905581ef-c586-485c-be3c-65bb81e0b52a
ONBOOT=yes
NM_CONTROLLED=no
BOOTPROTO=static
IPADDR=10.10.10.10
[root@node1 ~]# cat /etc/sysconfig/network-scripts/ifcfg-eth2
DEVICE=eth2
#HWADDR=08:00:27:61:F4:AE
TYPE=Ethernet
UUID=905581ef-c586-485c-be3c-65bb81e0b52a
ONBOOT=yes
```

```

NM_CONTROLLED=no
BOOTPROTO=static
IPADDR=10.10.10.11
[root@node1 ~]#

[root@node1 ~]# service NetworkManager stop
Stopping NetworkManager daemon:
[root@node1 ~]# chkconfig NetworkManager off
[root@node1 ~]#
[root@node1 ~]# cd /etc/udev/rules.d
[root@node1 rules.d]# ls -al
total 52
drwxr-xr-x. 2 root root 4096 Apr  1 20:27 .
drwxr-xr-x. 4 root root 4096 Apr  1 20:19 ..
-rw-r--r--. 1 root root 1652 Aug 25  2010 60-fprint-autosuspend.rules
-rw-r--r--. 1 root root 1060 Jun 29  2010 60-pcmcia.rules
-rw-r--r--. 1 root root  316 Aug 11  2014 60-raw.rules
-rw-r--r--. 1 root root  530 Apr  1 20:27 70-persistent-cd.rules
-rw-r--r--. 1 root root  750 Apr  1 20:14 70-persistent-net.rules
-rw-r--r--. 1 root root   40 Sep  9  2014 80-kvm.rules
-rw-r--r--. 1 root root  320 Jun 25  2014 90-alsa.rules
-rw-r--r--. 1 root root   83 Jun 19  2014 90-hal.rules
-rw-r--r--. 1 root root 2486 Jun 30  2010 97-bluetooth-serial.rules
-rw-r--r--. 1 root root  308 Aug 26  2014 98-kexec.rules
-rw-r--r--. 1 root root   54 Nov  3  2011 99-fuse.rules
[root@node1 rules.d]# rm -f 70-persistent-net.rules
[root@node1 rules.d]# reboot
[root@node1 rules.d]#
Broadcast message from root@node1.mydomain.com
(/dev/pts/1) at 8:11 ...

```

The system is going down for reboot NOW!

Now we have downloaded the SFHA(Veritas Storage Foundation High Availability) package on node1, we will now commence installing VCS. The download link for this is:  
[https://www4.symantec.com/Vrt/offer?a\\_id=24928](https://www4.symantec.com/Vrt/offer?a_id=24928)

```

[root@node1 ~]# cd /home/testuser/Downloads/
[root@node1 Downloads]# ls -al
total 3673040
drwxr-xr-x. 2 testuser testuser 4096 Apr  2 07:29 .
drwx----- 26 testuser testuser 4096 Apr  1 20:49 ..
-rw-r--r-- 1 testuser testuser 1673544724 Apr  2 07:20 linuxamd64_12102_database_1of2.zip
-rw-r--r-- 1 testuser testuser 1014530602 Apr  2 07:16 linuxamd64_12102_database_2of2.zip
-rw-r--r-- 1 testuser testuser 1073090232 Apr  2 07:36 VRTS_SF_HA_Solutions_6.1_RHEL.tar.gz
[root@node1 Downloads]#

```

```
[root@node1 Downloads]# gzip -d VRTS_SF_HA_Solutions_6.1_RHEL.tar.gz
[root@node1 Downloads]#
```

```
root@node1 Downloads]# tar xvf VRTS_SF_HA_Solutions_6.1_RHEL.tar
```

<snip>

```
./dvd1-redhatlinux/rhel6_x86_64/docs/dynamic_multipathing/
./dvd1-redhatlinux/rhel6_x86_64/docs/dynamic_multipathing/dmp_admin_61_lin.pdf
./dvd1-redhatlinux/rhel6_x86_64/docs/dynamic_multipathing/dmp_install_61_lin.pdf
./dvd1-redhatlinux/rhel6_x86_64/docs/dynamic_multipathing/dmp_notes_61_lin.pdf
./dvd1-redhatlinux/rhel6_x86_64/docs/getting_started.pdf
./dvd1-redhatlinux/rhel6_x86_64/docs/readme_first.txt
./dvd1-redhatlinux/rhel6_x86_64/docs/sfha_solutions/
./dvd1-redhatlinux/rhel6_x86_64/docs/sfha_solutions/sfhas_db2_admin_61_unix.pdf
./dvd1-redhatlinux/rhel6_x86_64/docs/sfha_solutions/sfhas_oracle_admin_61_unix.pdf
./dvd1-redhatlinux/rhel6_x86_64/docs/sfha_solutions/sfhas_replication_admin_61_lin.pdf
./dvd1-redhatlinux/rhel6_x86_64/docs/sfha_solutions/sfhas_smartio_solutions_61_lin.pdf
./dvd1-redhatlinux/rhel6_x86_64/docs/sfha_solutions/sfhas_solutions_61_lin.pdf
./dvd1-redhatlinux/rhel6_x86_64/docs/sfha_solutions/sfhas_tshoot_61_lin.pdf
./dvd1-redhatlinux/rhel6_x86_64/docs/sf_cluster_file_system_ha/
```

<snip>

```
[root@node1 Downloads]# cd dvd1-redhatlinux/
[root@node1 dvd1-redhatlinux]# ls -al
total 20
drwxr-xr-x  5 root  root  4096 Oct 29 2013 .
drwxr-xr-x  3 testuser testuser 4096 Apr  2 07:49 ..
drwxrwxr-x 18 root  root  4096 Oct 29 2013 rhel5_x86_64
drwxrwxr-x 18 root  root  4096 Oct 29 2013 rhel6_x86_64
drwxrwxr-x  8 root  root  4096 Oct 29 2013 symantec_ha_console
[root@node1 dvd1-redhatlinux]# cd rhel6_x86_64/
[root@node1 rhel6_x86_64]# ls -al
total 104
drwxrwxr-x 18 root root 4096 Oct 29 2013 .
drwxr-xr-x  5 root root 4096 Oct 29 2013 ..
```

```

drwxrwxr-x 3 root root 4096 Oct 29 2013 applicationha
drwxrwxr-x 4 root root 4096 Oct 29 2013 cluster_server
-rw-r--r-- 1 root root 951 Oct 29 2013 copyright
drwxr-xr-x 11 root root 4096 Oct 29 2013 docs
drwxrwxr-x 3 root root 4096 Oct 29 2013 dynamic_multipathing
drwxrwxr-x 3 root root 4096 Oct 29 2013 file_system
-rwxr-xr-x 1 root root 7165 Oct 29 2013 installer
drwxr-xr-x 6 root root 4096 Oct 29 2013 perl
drwxrwxr-x 3 root root 4096 Oct 29 2013 rpms
drwxr-xr-x 7 root root 4096 Oct 29 2013 scripts
drwxrwxr-x 4 root root 4096 Oct 29 2013 storage_foundation
drwxrwxr-x 4 root root 4096 Oct 29 2013 storage_foundation_cluster_file_system_ha
drwxrwxr-x 4 root root 4096 Oct 29 2013 storage_foundation_for_oracle_rac
drwxrwxr-x 4 root root 4096 Oct 29 2013 storage_foundation_high_availability
drwxrwxr-x 4 root root 4096 Oct 21 2013 VII
drwxrwxr-x 3 root root 4096 Oct 29 2013 volume_manager
-rwxr-xr-x 1 root root 18708 Oct 29 2013 webinstaller
drwxrwxr-x 2 root root 4096 Oct 29 2013 windows
drwxr-xr-x 6 root root 4096 Oct 29 2013 xprtl
[root@node1 rhel6_x86_64]#

```

```
[root@node1 rhel6_x86_64]# ./installer
```

Copyright (c) 2013 Symantec Corporation. All rights reserved. Symantec, the Symantec Logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The Licensed Software and Documentation are deemed to be "commercial computer software" and "commercial computer software documentation" as defined in FAR Sections 12.212 and DFARS Section 227.7202.

Logs are being written to /var/tmp/installer-201504020751raw while installer is in progress.

Symantec Product	Version Installed on node1	Licensed
------------------	----------------------------	----------

```
=====
```

Symantec Licensing Utilities (VRTSvlic) are not installed due to which products and licenses are not discovered.

Use the menu below to continue.

Task Menu:

- |                                     |                      |
|-------------------------------------|----------------------|
| P) Perform a Pre-Installation Check | I) Install a Product |
| C) Configure an Installed Product   | G) Upgrade a Product |

- O) Perform a Post-Installation Check    U) Uninstall a Product
- L) License a Product                      S) Start a Product
- D) View Product Descriptions            X) Stop a Product
- R) View Product Requirements           ?) Help

Enter a Task: [P,I,C,G,O,U,L,S,D,X,R,?]

Choose option "I"

- 1) Symantec Dynamic Multi-Pathing (DMP)
- 2) Symantec Cluster Server (VCS)
- 3) Symantec Storage Foundation (SF)
- 4) Symantec Storage Foundation and High Availability (SFHA)
- 5) Symantec Storage Foundation Cluster File System HA (SFCFSHA)
- 6) Symantec Storage Foundation for Oracle RAC (SF Oracle RAC)
- 7) Symantec ApplicationHA (ApplicationHA)
- b) Back to previous menu

Select a product to install: [1-7,b,q]

- 1) Symantec Dynamic Multi-Pathing (DMP)
- 2) Symantec Cluster Server (VCS)
- 3) Symantec Storage Foundation (SF)
- 4) Symantec Storage Foundation and High Availability (SFHA)
- 5) Symantec Storage Foundation Cluster File System HA (SFCFSHA)
- 6) Symantec Storage Foundation for Oracle RAC (SF Oracle RAC)
- 7) Symantec ApplicationHA (ApplicationHA)
- b) Back to previous menu

Select a product to install: [1-7,b,q] 4

- 1) Install minimal required rpms - 586 MB required
- 2) Install recommended rpms - 858 MB required
- 3) Install all rpms - 889 MB required
- 4) Display rpms to be installed for each option

Select the rpms to be installed on all systems? [1-4,q,?] (2) 3

Enter the 64 bit RHEL6 system names separated by spaces: [q,?] node1 node2

Either ssh or rsh needs to be set up between the local system and node2 for communication

Would you like the installer to setup ssh or rsh communication automatically between the systems?  
Superuser passwords for the systems will be asked. [y,n,q,?] (y) y

Enter the superuser password for system node2:

- 1) Setup ssh between the systems
- 2) Setup rsh between the systems



b) Back to previous menu

Select the communication method [1-2,b,q,?] (1)

Setting up communication between systems. Please wait.

Logs are being written to /var/tmp/installer-201504020833SGj while installer is in progress

Verifying systems: 100%

Estimated time remaining: (mm:ss) 0:00

8 of 8

Checking system  
communication ..... Done

Checking release  
compatibility ..... Done

Checking installed  
product ..... Done

Checking prerequisite patches and  
rpms ..... Done

Checking platform  
version ..... Failed

Checking file system free  
space ..... Done

Checking product  
licensing ..... Done

Performing product  
prechecks ..... Done

System verification checks completed

The following errors were discovered on the systems:

CPI ERROR V-9-0-0 System node1 is running Kernel Release 2.6.32-504.el6.x86\_64. Symantec Storage Foundation and High Availability Solutions 6.1 is not supported on Kernel Release 2.6.32-504.el6.x86\_64 without additional patches. Visit Symantec SORT web site to download the following required patches and follow the patch instructions to install the patches:

SFHA 6.1.1 release for RHEL6: <https://sort.symantec.com/patch/detail/8572>

SFHA 6.1.1 patch for RHEL6 U6: <https://sort.symantec.com/patch/detail/9327>

CPI ERROR V-9-0-0 System node2 is running Kernel Release 2.6.32-504.el6.x86\_64. Symantec Storage Foundation and High Availability Solutions 6.1 is not supported on Kernel Release 2.6.32-504.el6.x86\_64 without additional patches. Visit Symantec SORT web site to download the following required patches and follow the patch instructions to install the patches:

SFHA 6.1.1 release for RHEL6: <https://sort.symantec.com/patch/detail/8572>

SFHA 6.1.1 patch for RHEL6 U6: <https://sort.symantec.com/patch/detail/9327>

ssh is configured in password-less mode on node2

Do you want to cleanup the communication for the systems node2? [y,n,q] (n)

Instead of VCS 6.1, let us try VCS 6.2

[root@node1 rhel6\_x86\_64]# ./installer

Symantec Storage Foundation and High Availability

Solutions 6.2 Install Program

Copyright (c) 2014 Symantec Corporation. All rights reserved. Symantec, the Symantec Logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The Licensed Software and Documentation are deemed to be "commercial computer software" and "commercial computer software documentation" as defined in FAR Sections 12.212 and DFARS Section 227.7202.

Logs are being written to /var/tmp/installer-201504021652DRd while installer is in progress.

Symantec Storage Foundation and High Availability

Solutions 6.2 Install Program

Symantec Product	Version Installed on node1	Licensed
------------------	----------------------------	----------

Veritas File System	none	no
---------------------	------	----

Symantec Dynamic Multi-Pathing	none	no
Veritas Volume Manager	none	no
Symantec Cluster Server	none	no
Symantec ApplicationHA	none	no
Symantec Storage Foundation	none	no
Symantec Storage Foundation and High Availability	none	no
Symantec Storage Foundation Cluster File System HA	none	no
Symantec Storage Foundation for Oracle RAC	none	no

Task Menu:

- |                                      |                        |
|--------------------------------------|------------------------|
| P) Perform a Pre-Installation Check  | I) Install a Product   |
| C) Configure an Installed Product    | G) Upgrade a Product   |
| O) Perform a Post-Installation Check | U) Uninstall a Product |
| L) License a Product                 | S) Start a Product     |
| D) View Product Descriptions         | X) Stop a Product      |
| R) View Product Requirements         | ?) Help                |

Enter a Task: [P,I,C,G,O,U,L,S,D,X,R,?] I

## Symantec Storage Foundation and High Availability

### Solutions 6.2 Install Program

- 1) Symantec Dynamic Multi-Pathing (DMP)
- 2) Symantec Cluster Server (VCS)
- 3) Symantec Storage Foundation (SF)
- 4) Symantec Storage Foundation and High Availability (SFHA)
- 5) Symantec Storage Foundation Cluster File System HA (SFCFSHA)
- 6) Symantec Storage Foundation for Oracle RAC (SF Oracle RAC)
- 7) Symantec ApplicationHA (ApplicationHA)
- b) Back to previous menu

Select a product to install: [1-7,b,q] 4

This Symantec product may contain open source and other third party materials that are subject to a separate license. See the applicable Third-Party Notice at <http://www.symantec.com/about/profile/policies/eulas>

Do you agree with the terms of the End User License Agreement as specified in the storage\_foundation\_high\_availability/EULA/en/EULA\_SFHA\_Ux\_6.2.pdf file present on media?  
[y,n,q,?] y

## Symantec Storage Foundation and High

### Availability 6.2 Install Program

- 1) Install minimal required rpms - 827 MB required
- 2) Install recommended rpms - 1113 MB required

- 3) Install all rpms - 1114 MB required
- 4) Display rpms to be installed for each option

Select the rpms to be installed on all systems? [1-4,q,?] (2) 3

Enter the system names separated by spaces: [q,?] node1 node2

Availability 6.2 Install Program

Symantec Storage Foundation and High

node1 node2

Logs are being written to /var/tmp/installer-201504021652DRd while installer is in progress

Verifying systems: 100%

Estimated time remaining: (mm:ss) 0:00  
8 of 8

Checking system  
communication ..... Done

Checking release  
compatibility ..... Done

Checking installed  
product ..... Done

Checking prerequisite patches and  
rpms ..... Done

Checking platform  
version ..... Done

Checking file system free  
space ..... Done

Checking product  
licensing ..... Done

Performing product  
prechecks ..... Done

System verification checks completed successfully

Availability 6.2 Install Program

Symantec Storage Foundation and High

node1 node2

The following Symantec Storage Foundation and High Availability rpms will be uninstalled on all systems:

Rpm	Version	Rpm Description
VRTSvlic	3.02.61.010	Licensing
VRTSperl	5.16.1.6	Perl Redistribution

Press [Enter] to continue:

Symantec Storage Foundation and High  
Availability 6.2 Install Program  
node1 node2

The following Symantec Storage Foundation and High Availability rpms will be installed on all systems:

Rpm	Version	Rpm Description
VRTSperl	5.16.1.26	Perl Redistribution
VRTSvlic	3.02.62.003	Licensing
VRTSspt	6.2.0.000	Software Support Tools
VRTSvxvm	6.2.0.000	Volume Manager Binaries
VRTSaslapm	6.2.0.000	Volume Manager - ASL/APM
VRTSob	3.4.703	Enterprise Administrator Service
VRTSvxfs	6.2.0.000	File System
VRTSfsadv	6.2.0.000	File System Advanced Solutions
VRTSfssdk	6.2.0.000	File System Software Developer Kit
VRTSilt	6.2.0.000	Low Latency Transport
VRTSgab	6.2.0.000	Group Membership and Atomic Broadcast
VRTSvxfen	6.2.0.000	I/O Fencing
VRTSamf	6.2.0.000	Asynchronous Monitoring Framework
VRTSvc	6.2.0.000	Cluster Server
VRTScps	6.2.0.000	Cluster Server - Coordination Point Server
VRTSvcag	6.2.0.000	Cluster Server Bundled Agents
VRTSvcldr	6.2.0.000	Cluster Server Disk Reservation Modules
VRTSvcsea	6.2.0.000	Cluster Server Enterprise Agents
VRTSdbed	6.2.0.000	Storage Foundation Databases
VRTSodm	6.2.0.000	Oracle Disk Manager
VRTSsfmh	6.1.0.100	Storage Foundation Managed Host
VRTSvbs	6.2.0.000	Virtual Business Service
VRTSvcswiz	6.2.0.000	Cluster Server Wizards
VRTSsfcp62	6.2.0.000	Storage Foundation Installer

Press [Enter] to continue:

Symantec Storage Foundation and High  
Availability 6.2 Install Program

node1 node2

Logs are being written to /var/tmp/installer-201504021652DRd while installer is in progress

Uninstalling SFHA: 100%

Estimated time remaining: (mm:ss) 0:00

4 of 4

Performing SFHA preremove

tasks .....  
..... Done

Uninstalling

VRTSvlic .....  
..... Done

Uninstalling

VRTSperl .....  
..... Done

Performing SFHA postremove

tasks .....  
..... Done

Symantec Storage Foundation and High Availability Uninstall completed successfully

Availability 6.2 Install Program

Symantec Storage Foundation and High

node1 node2

Logs are being written to /var/tmp/installer-201504021652DRd while installer is in progress

Installing SFHA: 100%

Estimated time remaining: (mm:ss) 0:00

26 of 26

Performing SFHA preinstall

tasks .....  
..... Done

Installing VRTSperl

rpm .....  
..... Done

Installing VRTSvlic

rpm .....  
..... Done

Installing VRTSspt

rpm .....  
..... Done

Installing VRTSvxvm

rpm .....  
..... Done  
    Installing VRTSaslapm  
rpm .....  
..... Done  
    Installing VRTSob  
rpm .....  
..... Done  
    Installing VRTSvxf  
rpm .....  
..... Done  
    Installing VRTSfsadv  
rpm .....  
..... Done  
    Installing VRTSfssdk  
rpm .....  
..... Done  
    Installing VRTSIlt  
rpm .....  
..... Done  
    Installing VRTSgab  
rpm .....  
..... Done  
    Installing VRTSvxfen  
rpm .....  
..... Done  
    Installing VRTSamf  
rpm .....  
..... Done  
    Installing VRTSvc  
rpm .....  
..... Done  
    Installing VRTScps  
rpm .....  
..... Done  
    Installing VRTSvcag  
rpm .....  
..... Done  
    Installing VRTSvcdr  
rpm .....  
..... Done  
    Installing VRTSvcsea  
rpm .....  
..... Done  
    Installing VRTSdbed  
rpm .....  
..... Done  
    Installing VRTSodm  
rpm .....

```

..... Done
  Installing VRTSsfmh
rpm .....
..... Done
  Installing VRTSvbs
rpm .....
..... Done
  Installing VRTSvcswiz
rpm .....
..... Done
  Installing VRTSsfcp162
rpm .....
..... Done
  Performing SFHA postinstall
tasks .....
..... Done

```

Symantec Storage Foundation and High Availability Install completed successfully

```

                                     Symantec Storage Foundation and High
Availability 6.2 Install Program                                     node1 node2

```

To comply with the terms of Symantec's End User License Agreement, you have 60 days to either:

- \* Enter a valid license key matching the functionality in use on the systems
- \* Enable keyless licensing and manage the systems with a Management Server. For more details visit <http://go.symantec.com/sfhakeyless>. The product is fully functional during these 60 days.

- 1) Enter a valid license key
- 2) Enable keyless licensing and complete system licensing later

How would you like to license the systems? [1-2,q] (2)

- 1) SF Standard HA
- 2) SF Enterprise HA
- b) Back to previous menu

Select product mode to license: [1-2,b,q,?] (1) 2

Would you like to enable replication? [y,n,q] (n)

Would you like to enable the Global Cluster Option? [y,n,q] (n)

Registering SFHA license

SFHA vxkeyless key (SFHAENT) successfully registered on node1

SFHA vxkeyless key (SFHAENT) successfully registered on node2



Would you like to configure SFHA on node1 node2? [y,n,q] (n) y

## I/O Fencing

It needs to be determined at this time if you plan to configure I/O Fencing in enabled or disabled mode, as well as help in determining the number of network interconnects (NICS) required on your systems.

If you configure I/O Fencing in

enabled mode, only a single NIC is required, though at least two are recommended.

A split brain can occur if servers within the cluster become unable to communicate for any number of reasons. If I/O Fencing is not enabled, you run the risk of data corruption should a split brain occur.

Therefore, to avoid data

corruption due to split brain in CFS environments, I/O Fencing has to be enabled.

If you do not enable I/O Fencing, you do so at your own risk

See the Administrator's Guide for more information on I/O Fencing

Do you want to configure I/O Fencing in enabled mode? [y,n,q,?] (y)

To configure VCS, answer the set of questions on the next screen.

When [b] is presented after a question, 'b' may be entered to go back to the first question of the configuration set.

When [?] is presented after a question, '?' may be entered for help or additional information about the question.

Following each set of questions, the information you have entered will be presented for confirmation. To repeat the set of questions and correct any previous errors, enter 'n' at the confirmation prompt.

No configuration changes are made to the systems until all configuration questions are completed and confirmed.

Press [Enter] to continue:

To configure VCS for SFHA the following information is required:

- A unique cluster name

- One or more NICs per system used for heartbeat links

- A unique cluster ID number between 0-65535

- One or more heartbeat links are configured as private links

- You can configure one heartbeat link as a low-priority link

All systems are being configured to create one cluster.

Enter the unique cluster name: [q,?] mycluster

- 1) Configure the heartbeat links using LLT over Ethernet
- 2) Configure the heartbeat links using LLT over UDP
- 3) Configure the heartbeat links using LLT over RDMA
- 4) Automatically detect configuration for LLT over Ethernet
- b) Back to previous menu

How would you like to configure heartbeat links? [1-4,b,q,?] (4)

On Linux systems, only activated NICs can be detected and configured automatically.

Press [Enter] to continue:

Availability 6.2 Install Program	Symantec Storage Foundation and High
	node1 node2

Logs are being written to /var/tmp/installer-201504021652DRd while installer is in progress

Configuring LLT links: 100%

Estimated time remaining: (mm:ss) 0:00  
4 of 4

Checking system NICs on

```

node1 .....
..... 4 NICs found
    Checking system NICs on
node2 .....
..... 4 NICs found
    Checking network
links .....
..... 3 links found
    Setting link
priority .....
..... Done

```

Enter a unique cluster ID number between 0-65535: [b,q,?] (38221) 77

The cluster cannot be configured if the cluster ID 77 is in use by another cluster. Installer can perform a check to determine if the cluster ID is duplicate. The check will take less than a minute to complete.

Would you like to check if the cluster ID is in use by another cluster? [y,n,q] (y) n

#### Cluster information verification:

```

Cluster Name:   mycluster
Cluster ID Number: 77

```

Private Heartbeat NICs for node1:

```
link1=eth1
```

```
link2=eth2
```

Low-Priority Heartbeat NIC for node1:

```
link-lowpri1=eth0
```

Private Heartbeat NICs for node2:

```
link1=eth1
```

```
link2=eth2
```

Low-Priority Heartbeat NIC for node2:

```
link-lowpri1=eth0
```

Is this information correct? [y,n,q,?] (y)

The following data is required to configure the Virtual IP of the Cluster:

A public NIC used by each system in the cluster

A Virtual IP address and netmask

Do you want to configure the Virtual IP? [y,n,q,?] (n) y

Active NIC devices discovered on node1: eth0 eth1 eth2 virbr0

Enter the NIC for Virtual IP of the Cluster to use on node1: [b,q,?] (eth0)

Is eth0 to be the public NIC used by all systems? [y,n,q,b,?] (y)

Enter the Virtual IP address for the Cluster: [b,q,?] 192.168.0.200

Enter the NetMask for IP 192.168.0.200: [b,q,?] (255.255.255.0)

Availability 6.2 Install Program

Symantec Storage Foundation and High  
node1 node2

Cluster Virtual IP verification:

NIC: eth0

IP: 192.168.0.200

NetMask: 255.255.255.0

Is this information correct? [y,n,q] (y)

Availability 6.2 Install Program

Symantec Storage Foundation and High  
node1 node2

Symantec recommends to run Symantec Cluster Server in secure mode.

Running VCS in Secure Mode guarantees that all inter-system communication is encrypted, and users are verified with security credentials.

When running VCS in Secure Mode, NIS and system usernames and passwords are used to verify identity. VCS usernames and passwords are no longer utilized when a cluster is running in Secure Mode.

Would you like to configure the VCS cluster in secure mode? [y,n,q,?] (y) n

Fencing configuration

- 1) Configure Coordination Point client based fencing
- 2) Configure disk based fencing
- 3) Configure majority based fencing

Select the fencing mechanism to be configured in this Application Cluster: [1-3,q,?] q

Availability 6.2 Install Program

Symantec Storage Foundation and High  
node1 node2

To configure VCS, answer the set of questions on the next screen.

When [b] is presented after a question, 'b' may be entered to go back to the first question of the configuration set.

When [?] is presented after a question, '?' may be entered for help or additional information about the question.

Following each set of questions, the information you have entered will be presented for confirmation. To repeat the set of questions and correct any previous errors, enter 'n' at the confirmation prompt.

No configuration changes are made to the systems until all configuration questions are completed and confirmed.

Press [Enter] to continue:

	Symantec Storage Foundation and High
Availability 6.2 Install Program	node1 node2

To configure VCS for SFHA the following information is required:

- A unique cluster name
- One or more NICs per system used for heartbeat links
- A unique cluster ID number between 0-65535
  
- One or more heartbeat links are configured as private links
- You can configure one heartbeat link as a low-priority link

All systems are being configured to create one cluster.

Enter the unique cluster name: [q,?] mycluster

	Symantec Storage Foundation and High
Availability 6.2 Install Program	node1 node2

- 1) Configure the heartbeat links using LLT over Ethernet
- 2) Configure the heartbeat links using LLT over UDP
- 3) Configure the heartbeat links using LLT over RDMA
- 4) Automatically detect configuration for LLT over Ethernet
- b) Back to previous menu

How would you like to configure heartbeat links? [1-4,b,q,?] (4)

On Linux systems, only activated NICs can be detected and configured automatically.

Press [Enter] to continue:

Symantec Storage Foundation and High

Availability 6.2 Install Program

node1 node2

Logs are being written to /var/tmp/installer-201504021652DRd while installer is in progress

Configuring LLT links: 100%

Estimated time remaining: (mm:ss) 0:00

4 of 4

Checking system NICs on

node1 .....  
..... 4 NICs found

Checking system NICs on

node2 .....  
..... 4 NICs found

Checking network

links .....  
..... 3 links found

Setting link

priority .....  
..... Done

Enter a unique cluster ID number between 0-65535: [b,q,?] (38221) 77

The cluster cannot be configured if the cluster ID 77 is in use by another cluster. Installer can perform a check to determine if the cluster ID is duplicate. The check will take less than a minute to complete.

Would you like to check if the cluster ID is in use by another cluster? [y,n,q] (y) n

Symantec Storage Foundation and High

Availability 6.2 Install Program

node1 node2

Cluster information verification:

Cluster Name: mycluster

Cluster ID Number: 77

Private Heartbeat NICs for node1:

link1=eth1

link2=eth2

Low-Priority Heartbeat NIC for node1:

link-lowpri1=eth0

Private Heartbeat NICs for node2:

link1=eth1  
link2=eth2  
Low-Priority Heartbeat NIC for node2:  
link-lowpri1=eth0

Is this information correct? [y,n,q,?] (y)

Availability 6.2 Install Program  
Symantec Storage Foundation and High  
node1 node2

The following data is required to configure the Virtual IP of the Cluster:

A public NIC used by each system in the cluster  
A Virtual IP address and netmask

Do you want to configure the Virtual IP? [y,n,q,?] (n) y

Active NIC devices discovered on node1: eth0 eth1 eth2 virbr0

Enter the NIC for Virtual IP of the Cluster to use on node1: [b,q,?] (eth0)

Is eth0 to be the public NIC used by all systems? [y,n,q,b,?] (y)

Enter the Virtual IP address for the Cluster: [b,q,?] 192.168.0.200

Enter the NetMask for IP 192.168.0.200: [b,q,?] (255.255.255.0)

Availability 6.2 Install Program  
Symantec Storage Foundation and High  
node1 node2

Cluster Virtual IP verification:

NIC: eth0  
IP: 192.168.0.200  
NetMask: 255.255.255.0

Is this information correct? [y,n,q] (y)

Availability 6.2 Install Program  
Symantec Storage Foundation and High  
node1 node2

Symantec recommends to run Symantec Cluster Server in secure mode.

Running VCS in Secure Mode guarantees that all inter-system communication is encrypted, and users are verified with security credentials.

When running VCS in Secure Mode, NIS and system usernames and passwords are used to verify identity. VCS usernames and passwords are no longer utilized when a cluster is running in Secure Mode.

Would you like to configure the VCS cluster in secure mode? [y,n,q,?] (y) n

CPI WARNING V-9-40-6338 Symantec recommends that you install the cluster in secure mode. This ensures that communication between cluster components is encrypted and cluster information is visible to specified users only.

Are you sure that you want to proceed with non-secure installation? [y,n,q] (n) y

Availability 6.2 Install Program  
Symantec Storage Foundation and High  
node1 node2

The following information is required to add VCS users:

A user name  
A password for the user  
User privileges (Administrator, Operator, or Guest)

Do you wish to accept the default cluster credentials of 'admin/password'? [y,n,q] (y)

Do you want to add another user to the cluster? [y,n,q] (n)

Availability 6.2 Install Program  
Symantec Storage Foundation and High  
node1 node2

VCS User verification:

User: admin Privilege: Administrators

Passwords are not displayed

Is this information correct? [y,n,q] (y)

Availability 6.2 Install Program  
Symantec Storage Foundation and High  
node1 node2

The following information is required to configure SMTP notification:

The domain-based hostname of the SMTP server  
The email address of each SMTP recipient  
A minimum severity level of messages to send to each recipient

Do you want to configure SMTP notification? [y,n,q,?] (n)

Symantec Storage Foundation and High



Availability 6.2 Install Program

node1 node2

The following information is required to configure SNMP notification:

System names of SNMP consoles to receive VCS trap messages

SNMP trap daemon port numbers for each console

A minimum severity level of messages to send to each console

Do you want to configure SNMP notification? [y,n,q,?] (n)

All SFHA processes that are currently running must be stopped

Do you want to stop SFHA processes now? [y,n,q,?] (y)

Symantec Storage Foundation and High

Availability 6.2 Install Program

node1 node2

Logs are being written to /var/tmp/installer-201504021652DRd while installer is in progress

Stopping SFHA: 100%

Estimated time remaining: (mm:ss) 0:00

10 of 10

Performing SFHA prestop

tasks ..... Done

Stopping sfmh-

discovery ..... Done

Stopping

vxdclid ..... Done

Stopping

vxcperv ..... Done

Stopping

had ..... Done

Stopping

CmdServer ..... Done

Stopping

amf ..... Done

Stopping

```
vxfen ..... Done
  Stopping
gab ..... Done
  Stopping
llt ..... Done
  Stopping
```

Symantec Storage Foundation and High Availability Shutdown completed successfully

Symantec Storage Foundation and High  
Availability 6.2 Install Program  
node1 node2

Logs are being written to /var/tmp/installer-201504021652DRd while installer is in progress

Starting SFHA: 100%

Estimated time remaining: (mm:ss) 0:00  
21 of 21

```
Performing SFHA
configuration ..... Done
  Starting
vxdmp ..... Done
  Starting
vxio ..... Done
  Starting
vxspec ..... Done
  Starting
vxconfigd ..... Done
  Starting
vxesd ..... Done
  Starting
vxrelocd ..... Done
  Starting
vxcached ..... Done
  Starting
vxconfigbackupd .....
```

```

..... Done
  Starting
vxattachd .....
..... Done
  Starting
xprtld .....
..... Done
  Starting
vxportal .....
..... Done
  Starting
fdd .....
..... Done
  Starting
vxcafs .....
..... Done
  Starting
llt .....
..... Done
  Starting
gab .....
..... Done
  Starting
amf .....
..... Done
  Starting
had .....
..... Done
  Starting
CmdServer .....
..... Done
  Starting
vxodm .....
..... Done
  Performing SFHA poststart
tasks .....
..... Done

```

Symantec Storage Foundation and High Availability Startup completed successfully

	Symantec Storage Foundation and High
Availability 6.2 Install Program	
	node1 node2

#### Fencing configuration

- 1) Configure Coordination Point client based fencing
- 2) Configure disk based fencing
- 3) Configure majority based fencing

Select the fencing mechanism to be configured in this Application Cluster: [1-3,q,?] q

```
[root@node1 ~]# export PATH=$PATH:/opt/VRTSvcs/bin
[root@node1 ~]# hastatus -sum
```

-- SYSTEM STATE

-- System	State	Frozen
-----------	-------	--------

A node1	RUNNING	0
---------	---------	---

A node2	RUNNING	0
---------	---------	---

-- GROUP STATE

-- Group	System	Probed	AutoDisabled	State
----------	--------	--------	--------------	-------

B ClusterService	node1	Y	N	ONLINE
------------------	-------	---	---	--------

B ClusterService	node2	Y	N	OFFLINE
------------------	-------	---	---	---------

```
[root@node1 ~]#
```

OUR CLUSTER IS DONE

So, this command, you will use al the time, hastatus -sum. Please note it.

Now to install Oracle as a failover service on this cluster.

```
[root@node1 ~]# fdisk -l
```

Disk /dev/sda: 68.7 GB, 68719476736 bytes

255 heads, 63 sectors/track, 8354 cylinders

Units = cylinders of 16065 \* 512 = 8225280 bytes

Sector size (logical/physical): 512 bytes / 512 bytes

I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk identifier: 0x000c60df

Device	Boot	Start	End	Blocks	Id	System
--------	------	-------	-----	--------	----	--------

/dev/sda1	*	1	64	512000	83	Linux
-----------	---	---	----	--------	----	-------

Partition 1 does not end on cylinder boundary.

/dev/sda2		64	8355	66595840	8e	Linux LVM
-----------	--	----	------	----------	----	-----------

Disk /dev/mapper/vg\_node1-lv\_root: 53.7 GB, 53687091200 bytes

255 heads, 63 sectors/track, 6527 cylinders

Units = cylinders of 16065 \* 512 = 8225280 bytes

Sector size (logical/physical): 512 bytes / 512 bytes

I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk identifier: 0x00000000

Disk /dev/mapper/vg\_node1-lv\_swap: 4227 MB, 4227858432 bytes  
255 heads, 63 sectors/track, 514 cylinders  
Units = cylinders of 16065 \* 512 = 8225280 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x00000000

Disk /dev/mapper/vg\_node1-lv\_home: 10.3 GB, 10276044800 bytes  
255 heads, 63 sectors/track, 1249 cylinders  
Units = cylinders of 16065 \* 512 = 8225280 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x00000000

Disk /dev/sdb: 36.8 GB, 36842766336 bytes  
64 heads, 32 sectors/track, 35136 cylinders  
Units = cylinders of 2048 \* 512 = 1048576 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0xa958974b

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		1	35136	35979248	83	Linux

ON BOTH NODES RUN THIS COMMAND:

```
[root@node1 ~]# vxdctl enable  
[root@node1 ~]#
```

```
[root@node2 ~]# vxdctl enable  
[root@node2 ~]#
```

```
[root@node1 ~]# vxdisk -e list
```

DEVICE	TYPE	DISK	GROUP	STATUS	OS_NATIVE_NAME	ATTR
disk_0	auto:none	-	-	online invalid	sdc	-
sda	auto:LVM	-	-	online invalid	sda	-

```
[root@node1 ~]#
```

So, please note that /dev/sdc is now recognized by veritas as disk\_0

For this training all you need to know about Veritas File System is this:

One or more physical or virtual disks go to create a veritas disk group, a veritas disk group goes into the creation of a veritas volume and finally a veritas file system is created on a veritas volume.

So, one or more disks->veritas disk group->veritas volume->veritas file system. Hence to increase the size of an existing file system we have to add disks to the underlying diskgroup and then increase the volume and filesystem. This is important for WORK and the process is documented at the end of this manual.

```
[root@node1 ~]# vxdiskadm
```

NOTE: This is Veritas Storage Foundation command. We will create a vxfs(veritas file system) on this shared disk and install oracle on it.

## Volume Manager Support Operations

Menu:: VolumeManager/Disk

- 1 Add or initialize one or more disks
- 2 Encapsulate one or more disks
- 3 Remove a disk
- 4 Remove a disk for replacement
- 5 Replace a failed or removed disk
- 6 Mirror volumes on a disk
- 7 Move volumes from a disk
- 8 Enable access to (import) a disk group
- 9 Remove access to (deport) a disk group
- 10 Enable (online) a disk device
- 11 Disable (offline) a disk device
- 12 Mark a disk as a spare for a disk group
- 13 Turn off the spare flag on a disk
- 14 Unrelocate subdisks back to a disk
- 15 Exclude a disk from hot-relocation use
- 16 Make a disk available for hot-relocation use
- 17 Prevent multipathing/Suppress devices from VxVM's view
- 18 Allow multipathing/Unsuppress devices from VxVM's view
- 19 List currently suppressed/non-multipathed devices
- 20 Change the disk naming scheme
- 21 Change/Display the default disk layouts
- 22 Dynamic Reconfiguration Operations
- list List disk information
  
- ? Display help about menu
- ?? Display help about the menuing system
- q Exit from menus

Select an operation to perform: 1

## Add or initialize disks

Menu:: VolumeManager/Disk/AddDisks

Use this operation to add one or more disks to a disk group. You can add the selected disks to an existing disk group or to a new disk group

that will be created as a part of the operation. The selected disks may also be added to a disk group as spares. Or they may be added as nohotuses to be excluded from hot-relocation use. The selected disks may also be initialized without adding them to a disk group leaving the disks available for use as replacement disks.

More than one disk may be entered at the prompt. Here are some disk selection examples:

sda:        add only disk sda  
sdb hdc:    add both disk sdb and hdc  
xyz\_0 :    a single disk (in the enclosure based naming scheme)  
xyz\_ :    all disks on the enclosure whose name is xyz

Select disk devices to add: [<pattern-list>,list,q,?] 1

DEVICE	DISK	GROUP	STATUS
disk_0	-	-	online invalid
sda	-	-	online invalid

Select disk devices to add: [<pattern-list>,list,q,?]

Here is the disk selected. Output format: [Device\_Name]

disk\_0

Continue operation? [y,n,q,?] (default: y) y

You can choose to add this disk to an existing disk group, a new disk group, or leave the disk available for use by future add or replacement operations. To create a new disk group, select a disk group name that does not yet exist. To leave the disk available for future use, specify a disk group name of "none".

Which disk group [<group>,none,list,q,?] newdg

Create a new group named newdg? [y,n,q,?] (default: y)

Create the disk group as a CDS disk group? [y,n,q,?] (default: y)

Use a default disk name for the disk? [y,n,q,?] (default: y)

Add disk as a spare disk for newdg? [y,n,q,?] (default: n)

Exclude disk from hot-relocation use? [y,n,q,?] (default: n)

Add site tag to disk? [y,n,q,?] (default: n)

A new disk group will be created named newdg and the selected disks will be added to the disk group with default disk names.

disk\_0

Continue with operation? [y,n,q,?] (default: y)

The following disk device has a valid partition table, but does not appear to have been initialized for the Volume Manager. If there is data on the disk that should NOT be destroyed you should encapsulate the existing disk partitions as volumes instead of adding the disk as a new disk.

Output format: [Device\_Name]

disk\_0

Encapsulate this device? [y,n,q,?] (default: y) n

disk\_0

Instead of encapsulating, initialize? [y,n,q,?] (default: n) y

Initializing device disk\_0.

Enter desired private region length  
[<privlen>,q,?] (default: 65536)

VxVM NOTICE V-5-2-120

Creating a new disk group named newdg containing the disk device disk\_0 with the name newdg01.

Add or initialize other disks? [y,n,q,?] (default: n)

Volume Manager Support Operations  
Menu:: VolumeManager/Disk



- 1 Add or initialize one or more disks
- 2 Encapsulate one or more disks
- 3 Remove a disk
- 4 Remove a disk for replacement
- 5 Replace a failed or removed disk
- 6 Mirror volumes on a disk
- 7 Move volumes from a disk
- 8 Enable access to (import) a disk group
- 9 Remove access to (deport) a disk group
- 10 Enable (online) a disk device
- 11 Disable (offline) a disk device
- 12 Mark a disk as a spare for a disk group
- 13 Turn off the spare flag on a disk
- 14 Unrelocate subdisks back to a disk
- 15 Exclude a disk from hot-relocation use
- 16 Make a disk available for hot-relocation use
- 17 Prevent multipathing/Suppress devices from VxVM's view
- 18 Allow multipathing/Unsuppress devices from VxVM's view
- 19 List currently suppressed/non-multipathed devices
- 20 Change the disk naming scheme
- 21 Change/Display the default disk layouts
- 22 Dynamic Reconfiguration Operations
- list List disk information
  
- ? Display help about menu
- ?? Display help about the menuing system
- q Exit from menus

Select an operation to perform:q

```
[root@node1 ~]# vxdisk -e list
DEVICE    TYPE      DISK      GROUP    STATUS    OS_NATIVE_NAME  ATTR
disk_0    auto:cdsdisk  newdg01   newdg    online    sdc             -
sda       auto:LVM      -         -        online invalid    sda             -
[root@node1 ~]#
```

AS YOU MAY OBSERVE A NEW DISKGROUP “newdg” HAS BEEN CREATED

Now that we have a diskgroup we will create a volume on top of it and on top of that we will create a VxFS (Veritas File System) and on top of that we will install Oracle.

So, first things first. Creating a volume.

We know that we have a 36GB lun from openfiler (remember our virtual SAN), so we will choose a size that works (some disk space is used up while initializing the disk and some while adding it to the disk group)

```
[root@node1 ~]# vxassist -g newdg make newvol 36G
VxVM vxassist ERROR V-5-1-15315 Cannot allocate space for 75497472 block volume: Not enough
HDD devices that meet specification.
[root@node1 ~]# vxassist -g newdg make newvol 35G
VxVM vxassist ERROR V-5-1-15315 Cannot allocate space for 73400320 block volume: Not enough
HDD devices that meet specification.
[root@node1 ~]# vxassist -g newdg make newvol 34G
[root@node1 ~]#
```

So 34G size works!!

Now to create a VxFS on this volume:

```
[root@node1 vx]# mkfs.vxfs /dev/vx/rdisk/newdg/newvol
version 10 layout
71303168 sectors, 35651584 blocks of size 1024, log size 65536 blocks
rcq size 4096 blocks
largefiles supported
maxlink supported
[root@node1 vx]#
```

Now we will add this volume(and associated filesystem) into Veritas Cluster Server (VCS) control.

ON BOTH NODES DO THIS:

```
[root@node2 ~]# mkdir /ORACLE
[root@node2 ~]#
```

```
[root@node1 vx]# mkdir /ORACLE
[root@node1 vx]#
```

```
[root@node1 vx]# hagr -add new_servicegroup
VCS NOTICE V-16-1-10136 Group added; populating SystemList and setting the Parallel attribute
recommended before adding resources
NOTE: we are adding a new service group to the VCS configuration. This will not be a parallel service
group. We need to tell VCS which systems this service group will run on.
```

```
[root@node1 vx]# hagr -modify new_servicegroup Parallel 0
NOTE: we told VCS that this is not a parallel service group.
```

```
[root@node1 vx]# hagr -modify new_servicegroup SystemList node1 0 node2 1
NOTE: we told VCS that the service group which we created will run on node 1 if it is available and
then on node2
```

```
[root@node1 vx]# hagr -modify new_servicegroup AutoStartList node1
```

NOTE: we told VCS that the service group is set to automatically start on node1

```
[root@node1 vx]#
```

```
[root@node1 vx]# hares -add diskgroup_resource DiskGroup new_servicegroup
```

NOTE: we are now adding resources to the service group which we created. The first resource is the diskgroup, we are naming our resource “diskgroup\_resource” and it is of the type “DiskGroup”.

VCS NOTICE V-16-1-10242 Resource added. Enabled attribute must be set before agent monitors

```
[root@node1 vx]# hares -modify diskgroup_resource DiskGroup newdg
```

NOTE: we are now telling VCS that this resource corresponds to the real diskgroup named “newdg”

```
[root@node1 vx]# hares -modify diskgroup_resource Critical 0
```

NOTE: We are modifying the “Critical” attribute of the resource. More on “Critical” resources later.

```
[root@node1 vx]# hares -modify diskgroup_resource Enabled 1
```

NOTE: We are enabling monitoring of this resource by VCS.

```
[root@node1 vx]#
```

```
[root@node1 vx]# hares -add volume_resource Volume new_servicegroup
```

NOTE: We are adding a new resource of the type “Volume” to the service group.

VCS NOTICE V-16-1-10242 Resource added. Enabled attribute must be set before agent monitors

```
[root@node1 vx]# hares -modify volume_resource Critical 0
```

NOTE: We are modifying the “Critical” attribute of the resource.

```
[root@node1 vx]# hares -modify volume_resource DiskGroup newdg
```

NOTE: we are telling VCS that this volume is based on the “newdg” DiskGroup. We are modifying the “DiskGroup” attribute of this resource.

```
[root@node1 vx]# hares -modify volume_resource Volume newvol
```

NOTE: We are telling VCS that our Volume resource named “volume\_resource” corresponds to the REAL volume we called “newvol”.

```
[root@node1 vx]# hares -modify volume_resource Enabled 1
```

NOTE: We are enabling monitoring of this resource by VCS.

```
[root@node1 vx]#
```

```
[root@node1 vx]# hares -add mount_resource Mount new_servicegroup
```

NOTE: We are adding a new resource of type “Mount” to the service group.

VCS NOTICE V-16-1-10242 Resource added. Enabled attribute must be set before agent monitors

```
[root@node1 vx]# hares -modify mount_resource Critical 0
```

NOTE: we are setting the value of the “Critical” attribute of this resource to zero.

```
[root@node1 vx]# hares -modify mount_resource MountPoint /ORACLE
```

NOTE: We are defining the “MountPoint” attribute of this resource to be /ORACLE and TELLING

## VCS THIS

```
[root@node1 vx]# hares -modify mount_resource BlockDevice /dev/vx/dsk/newdg/newvol
```

NOTE: We are telling VCS what the block device of the mount point is.

```
[root@node1 vx]# hares -modify mount_resource FSType vxfs
```

NOTE: We are telling VCS that this is a vxfs type of filesystem.

```
[root@node1 vx]# hares -modify mount_resource FsckOpt %-y
```

NOTE: We are telling VCS to fsck this file system.

```
[root@node1 vx]# hares -modify mount_resource Enabled 1
```

NOTE: We are telling VCS to monitor this resource.

```
[root@node1 vx]#
```

```
[root@node1 vx]# hares -modify diskgroup_resource Critical 1
```

```
[root@node1 vx]# hares -modify volume_resource Critical 1
```

```
[root@node1 vx]# hares -modify mount_resource Critical 1
```

### About critical and non-critical resources

The Critical attribute for a resource defines whether a service group fails over when the resource faults. If a resource is configured as non-critical (by setting the Critical attribute to 0) and no resources depending on the failed resource are critical, the service group will not fail over. VCS takes the failed resource offline and updates the group status to ONLINE|PARTIAL. The attribute also determines whether a service group tries to come online on another node if, during the group's online process, a resource fails to come online.

Once all of the resources have been added to the service group, VCS must be told in which order to bring them online. Otherwise, it will try to bring all resources on line simultaneously. The command 'hares -link res2 res1' will create a dependency such that "res1" MUST be online before VCS will attempt to start "res2":

Obviously the diskgroup must be operational before the volume which in turn must be available prior to the mount resource.

Adding dependencies of the resources into VCS, this follows the format:

```
hares -link Resource_Name Resource_it_depends_on
```

```
[root@node1 vx]# hares -link volume_resource diskgroup_resource
```

```
[root@node1 vx]# hares -link mount_resource volume_resource
```

```
[root@node1 vx]#
```

```
[root@node1 vx]# haconf -dump -makero
```

NOTE: We are writing and saving our changes to the VCS configuration file which is /opt/VRTSvcs/conf/config/main.cf

```
[root@node1 vx]#
```

```
[root@node1 vx]# hastatus -sum
```

NOTE: this is the most frequently used VCS command, as the name “hastatus” suggests it tells you about the status of your cluster. The “sum” stands for summary, try omitting the -sum....you will have to Ctrl-c to get back. Please try it.

```
[root@node1 vx]# hagr -switch new_servicegroup -to node2
```

NOTE: We are switching our service group from node1 where it currently is to node2.

```
[root@node1 vx]# hastatus -sum
```

-- SYSTEM STATE

-- System	State	Frozen
-----------	-------	--------

A node1	RUNNING	0
---------	---------	---

A node2	RUNNING	0
---------	---------	---

-- GROUP STATE

-- Group	System	Probed	AutoDisabled	State
----------	--------	--------	--------------	-------

B ClusterService	node1	Y	N	ONLINE
------------------	-------	---	---	--------

B ClusterService	node2	Y	N	OFFLINE
------------------	-------	---	---	---------

B new_servicegroup	node1	Y	N	OFFLINE
--------------------	-------	---	---	---------

B new_servicegroup	node2	Y	N	STARTING PARTIAL
--------------------	-------	---	---	------------------

-- RESOURCES ONLINE

-- Group	Type	Resource	System	IState
----------	------	----------	--------	--------

F new_servicegroup	Volume	volume_resource	node2	W_ONLINE
--------------------	--------	-----------------	-------	----------

```
[root@node1 vx]# hastatus -sum
```

-- SYSTEM STATE

-- System	State	Frozen
-----------	-------	--------

A node1	RUNNING	0
---------	---------	---

A node2	RUNNING	0
---------	---------	---

-- GROUP STATE

-- Group	System	Probed	AutoDisabled	State
----------	--------	--------	--------------	-------

B ClusterService	node1	Y	N	ONLINE
------------------	-------	---	---	--------

B ClusterService	node2	Y	N	OFFLINE
------------------	-------	---	---	---------

B new_servicegroup	node1	Y	N	OFFLINE
--------------------	-------	---	---	---------

B new_servicegroup	node2	Y	N	STARTING PARTIAL
--------------------	-------	---	---	------------------

-- RESOURCES ONLINING

-- Group	Type	Resource	System	IState
F new_servicegroup	Mount	mount_resource	node2	W_ONLINE

[root@node1 vx]# hastatus -sum

-- SYSTEM STATE

-- System	State	Frozen
A node1	RUNNING	0
A node2	RUNNING	0

-- GROUP STATE

-- Group	System	Probed	AutoDisabled	State
B ClusterService	node1	Y	N	ONLINE
B ClusterService	node2	Y	N	OFFLINE
B new_servicegroup	node1	Y	N	OFFLINE
B new_servicegroup	node2	Y	N	ONLINE

[root@node1 vx]#

SO WE HAVE BEEN ABLE TO SWITCH THE VOLUME FROM ONE NODE TO THE OTHER, HENCE IF ONE NODE GOES DOWN(CRASHES) THE VOLUME WILL STILL BE AVAILABLE VIA THE OTHER NODE.

Check node2

[root@node2 ~]# df -kh

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/mapper/vg_node1-lv_root	50G	6.5G	41G	14%	/
tmpfs	940M	224K	939M	1%	/dev/shm
/dev/sda1	477M	65M	387M	15%	/boot
/dev/mapper/vg_node1-lv_home	9.3G	23M	8.8G	1%	/home
tmpfs	4.0K	0	4.0K	0%	/dev/vx
/dev/vx/dsk/newdg/newvol	34G	78M	32G	1%	/ORACLE

[root@node2 ~]#



Now we will install Oracle Database as a failover service on our cluster

We will install Mocha X Server on our PC  
<http://www.mochasoft.dk/freeware/x11.htm>

We will download our database installation files on node1.  
On node 1 are root run the following commands

```
unzip linuxamd64_12102_database_1of2.zip
unzip linuxamd64_12102_database_2of2.zip
```

EDIT THIS FILE (/etc/security/limits.conf) ON BOTH NODES, IT SHOULD ULTIMATELY LOOK LIKE THIS:

```
[root@node1 ~]# cat /etc/security/limits.conf
# /etc/security/limits.conf
#
#Each line describes a limit for a user in the form:
#
#<domain>      <type> <item> <value>
#
#Where:
#<domain> can be:
#   - a user name
#   - a group name, with @group syntax
#   - the wildcard *, for default entry
#   - the wildcard %, can be also used with %group syntax,
#       for maxlogin limit
#
#<type> can have the two values:
#   - "soft" for enforcing the soft limits
#   - "hard" for enforcing hard limits
#
#<item> can be one of the following:
#   - core - limits the core file size (KB)
#   - data - max data size (KB)
#   - fsize - maximum filesize (KB)
#   - memlock - max locked-in-memory address space (KB)
#   - nofile - max number of open file descriptors
#   - rss - max resident set size (KB)
#   - stack - max stack size (KB)
#   - cpu - max CPU time (MIN)
#   - nproc - max number of processes
#   - as - address space limit (KB)
#   - maxlogins - max number of logins for this user
#   - maxsyslogins - max number of logins on the system
#   - priority - the priority to run user process with
#   - locks - max number of file locks the user can hold
#   - sigpending - max number of pending signals
```



```
# - msgqueue - max memory used by POSIX message queues (bytes)
# - nice - max nice priority allowed to raise to values: [-20, 19]
# - rtprio - max realtime priority
#
#<domain> <type> <item> <value>
#
```

```
##*      soft  core      0
##*      hard  rss       10000
#@student  hard  nproc     20
#@faculty  soft  nproc     20
#@faculty  hard  nproc     50
#ftp       hard  nproc      0
#@student  -     maxlogins  4
oracle     soft  nproc     2048
oracle     hard  nproc     20480
oracle     soft  stack     20480
oracle     hard  stack     32768
oracle     soft  nofile    4096
# End of file
oracle hard nofile 65536
```

EDIT THIS FILE (/etc/sysctl.conf) OMN BOTH NODES, IT SHOULD ULTIMATELY LOOK LIKE THIS:

```
[root@node1 ~]# cat /etc/sysctl.conf
# Kernel sysctl configuration file for Red Hat Linux
#
# For binary values, 0 is disabled, 1 is enabled. See sysctl(8) and
# sysctl.conf(5) for more details.

# Controls IP packet forwarding
net.ipv4.ip_forward = 0

# Controls source route verification
net.ipv4.conf.default.rp_filter = 1

# Do not accept source routing
net.ipv4.conf.default.accept_source_route = 0

# Controls the System Request debugging functionality of the kernel
kernel.sysrq = 0

# Controls whether core dumps will append the PID to the core filename.
# Useful for debugging multi-threaded applications.
kernel.core_uses_pid = 1

# Controls the use of TCP syncookies
net.ipv4.tcp_syncookies = 1
```

```
# Disable netfilter on bridges.
net.bridge.bridge-nf-call-ip6tables = 0
net.bridge.bridge-nf-call-iptables = 0
net.bridge.bridge-nf-call-arptables = 0

# Controls the default maximum size of a message queue
kernel.msgmnb = 65536

# Controls the maximum size of a message, in bytes
kernel.msgmax = 65536

# Controls the maximum shared segment size, in bytes
kernel.shmmax = 68719476736

# Controls the maximum number of shared memory segments, in pages
kernel.shmall = 4294967296

fs.suid_dumpable = 1
fs.aio-max-nr = 1048576
fs.file-max = 6815744
#kernel.shmall = 2097152
#kernel.shmmax = 536870912
kernel.shmmni = 4096
kernel.sem = 250 32000 100 128
net.core.rmem_default = 262144
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 1048586
net.ipv4.ip_local_port_range = 9000 65500
[root@node1 ~]#
```

PLEASE RUN THIS COMMAND (sysctl -p) ON BOTH NODES

Create the oracle user and the required groups ON BOTH NODES WITH THE SAME UID/GID,  
please refer to this page:  
<https://docs.oracle.com/database/121/LTDQI/toc.htm#CHDJIAAI>

```
cd database
as user oracle:
export DISPLAY=192.168.0.187:0.0 (note: <MY PC'S IP ADDRESS:0.0>)
```

```
[oracle@node1 database]$ ./runInstaller
```

Please do a cksum verification of the downloaded files against the checksum present on Oracle's site to ensure that your download was not corrupted while downloading. (NOTE: my download was corrupted three times and I supposedly have a pretty good provider)

Please note that you will have to do yum installs of the many packages that are missing, please do so on both nodes(BOTH nodes have to be kept as identical as possible).



Oracle Database 12c Release 1 Installer

Configure Security Updates

Configure Security Updates

Installation Option

Grid Installation Options

Install Type

Typical Installation

Prerequisite Checks

Summary

Install Product

Finish

Provide your email address to be informed of security issues, install the product and initiate configuration manager. [View details.](#)

Email:

Easier for you if you use your My Oracle Support email address/username.

☐ I wish to receive security updates via My Oracle Support.

My Oracle Support Password:

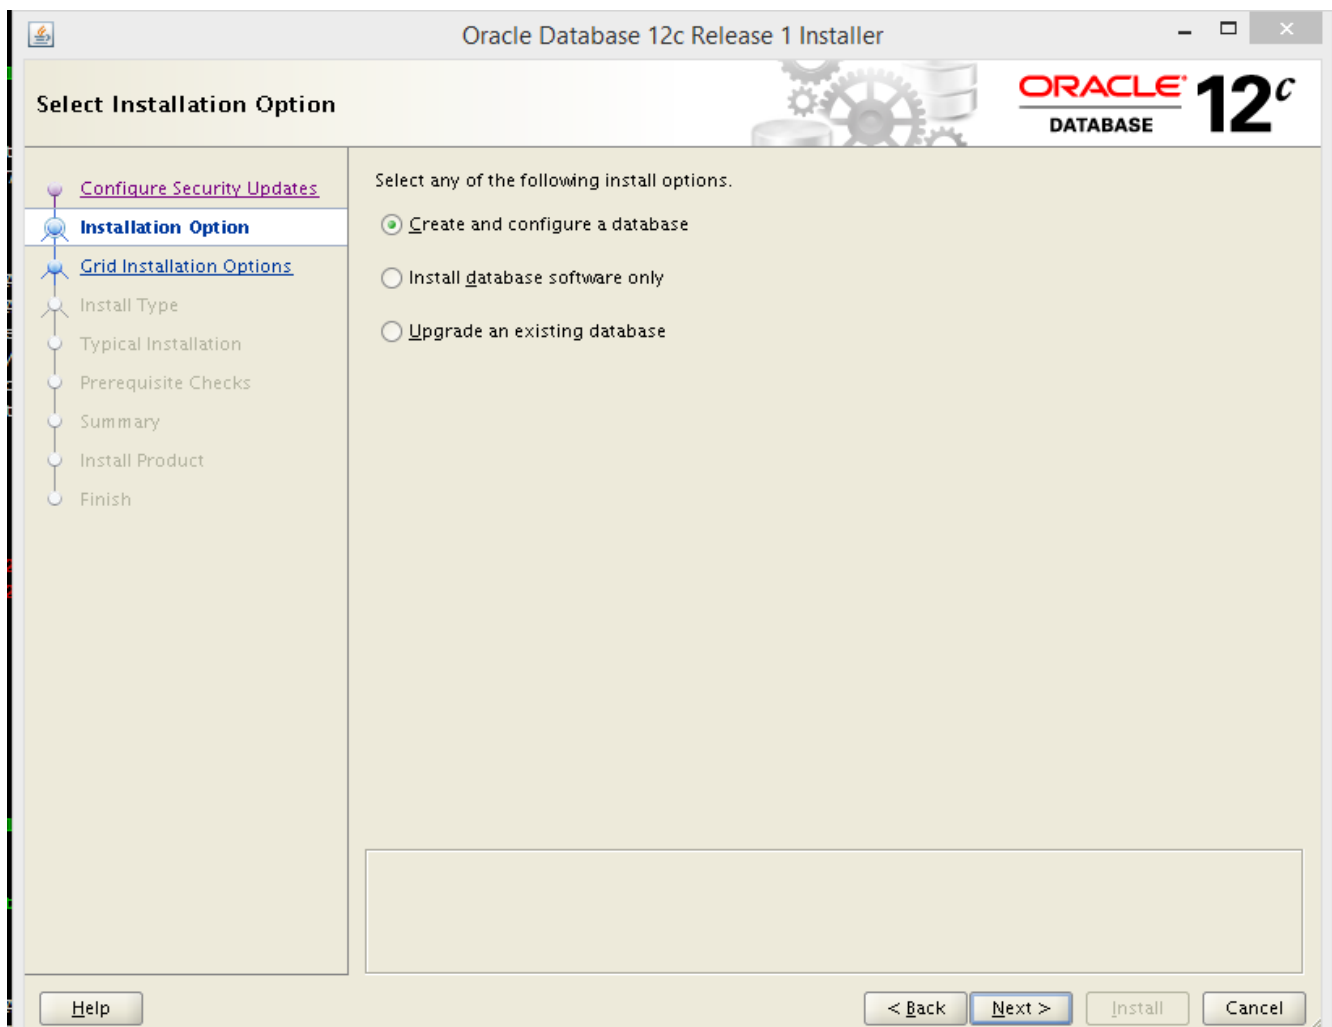
Help

< Back

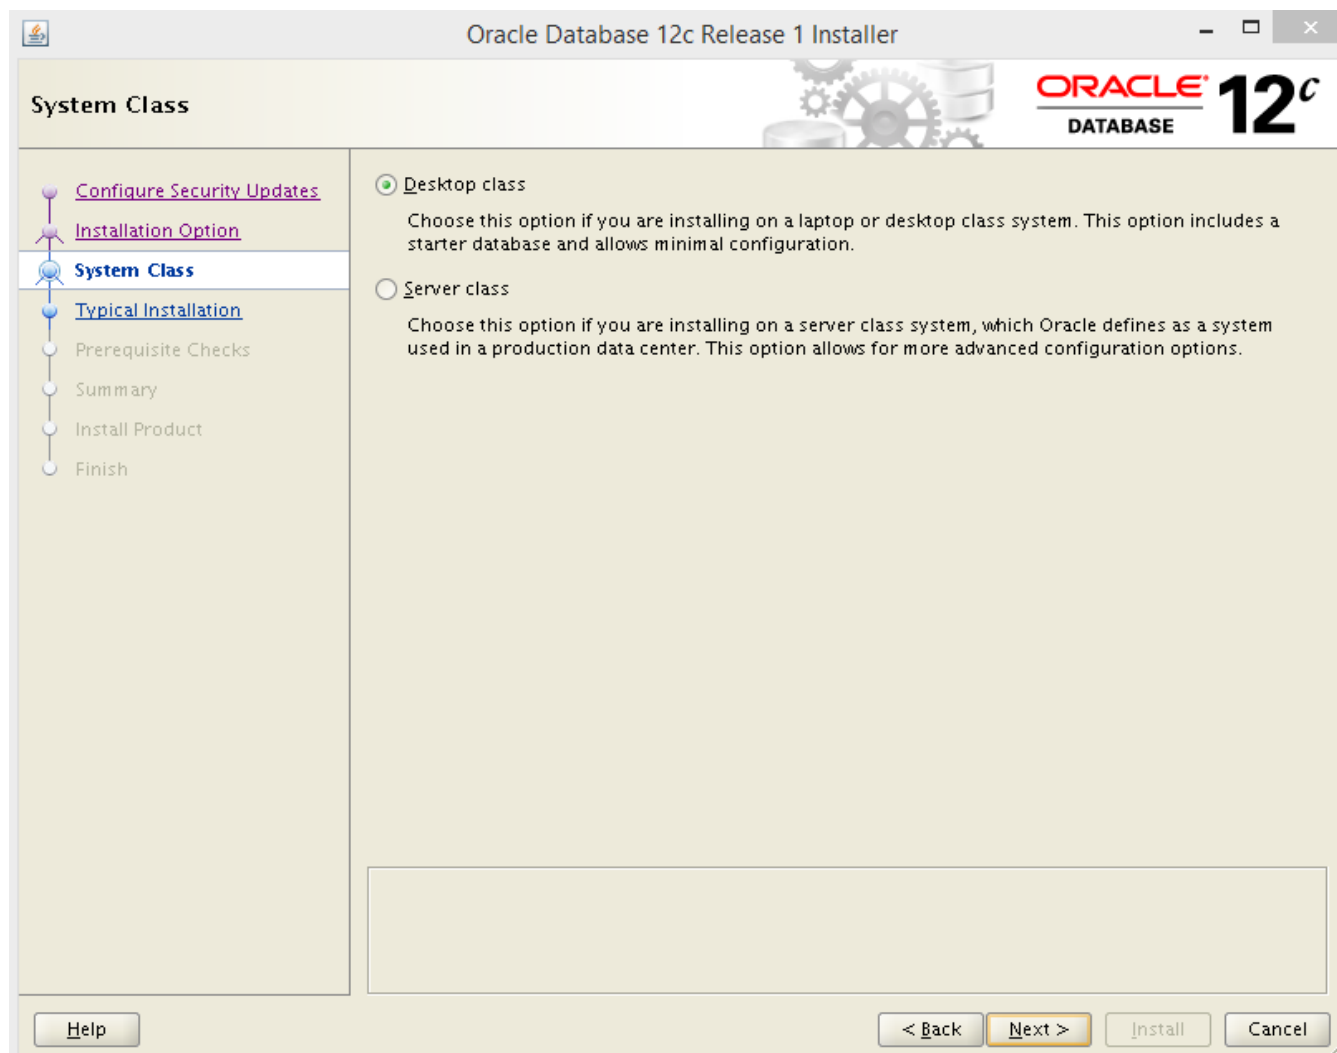
Next >

Install

Cancel



Since this is for our training and testing we will choose desktop class



Oracle Database 12c Release 1 Installer

## Typical Install Configuration

Perform full database installation with basic configuration.

Oracle base:

Software location:

Database file location:

Database edition:

Character set:

OSDBA group:

Global database name:


Administrative password:

Confirm password:

☐ Create as Container database

Pluggable database name:

Messages:

 Administrative password:[INS-30011] The ADMIN password entered does not conform to the Oracle recommended standards.

PLEASE DO not CREATE AS CONTAINER DATABASE

Oracle Database 12c Release 1 Installer

Create Inventory

Configure Security Updates

Installation Option

System Class

Typical Installation

Create Inventory

Prerequisite Checks

Summary

Install Product

Finish

You are starting your first installation on this host. Specify a directory for installation metadata files (for example, install log files). This directory is called the "inventory directory". The installer automatically sets up subdirectories for each product to contain inventory data. The subdirectory for each product typically requires 150 kilobytes of disk space.

Inventory Directory:

Specify an operating system group whose members have write permission to the inventory directory (oralInventory).

oralInventory Group Name:

ORACLE<sup>®</sup> 12<sup>c</sup>

DATABASE

Help

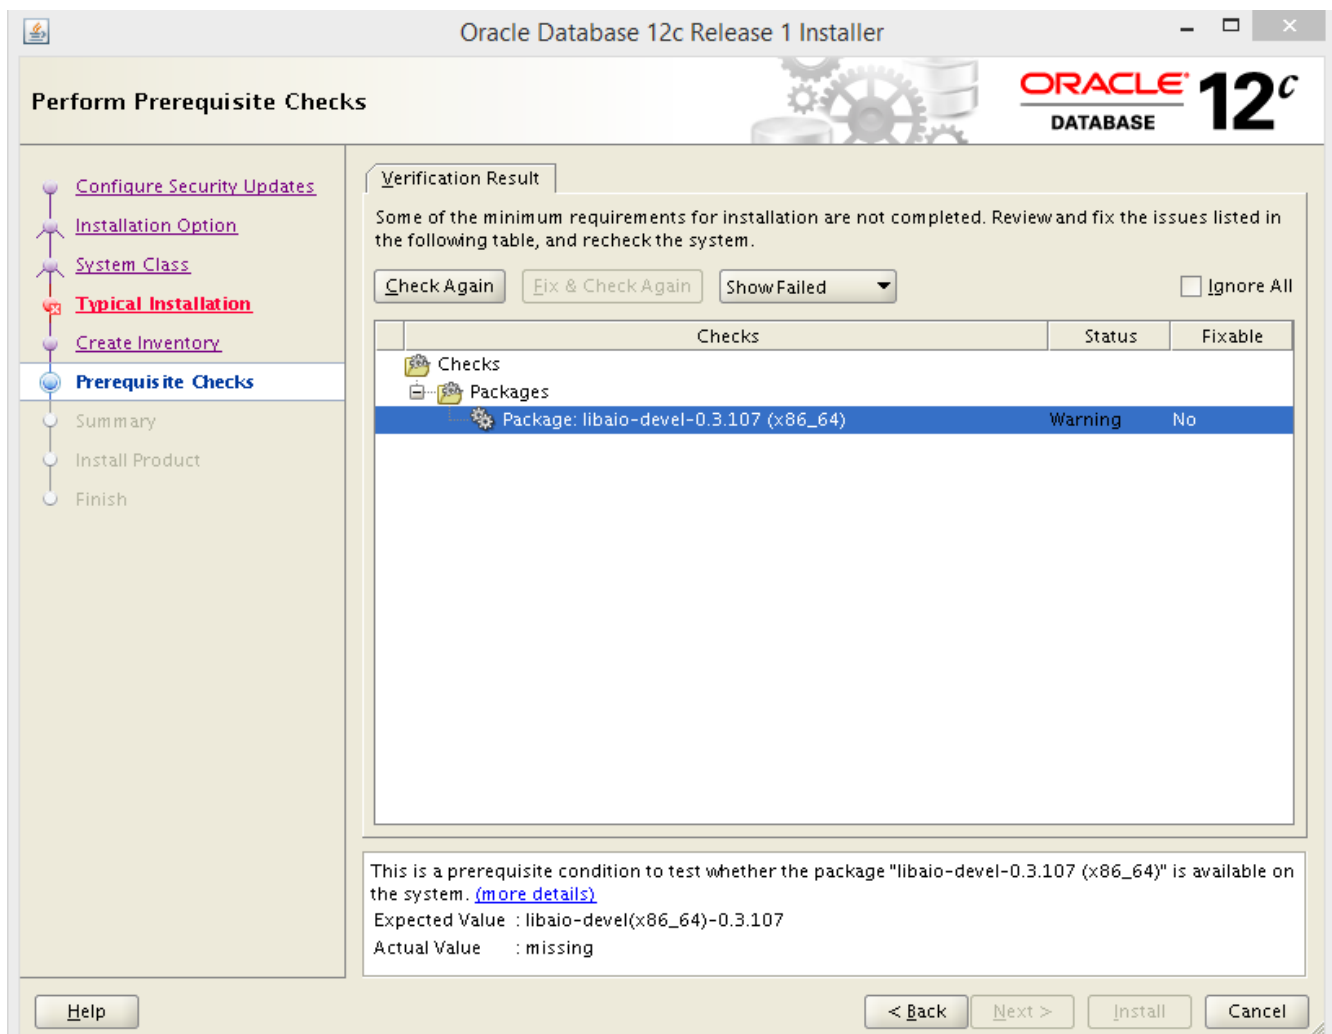
< Back

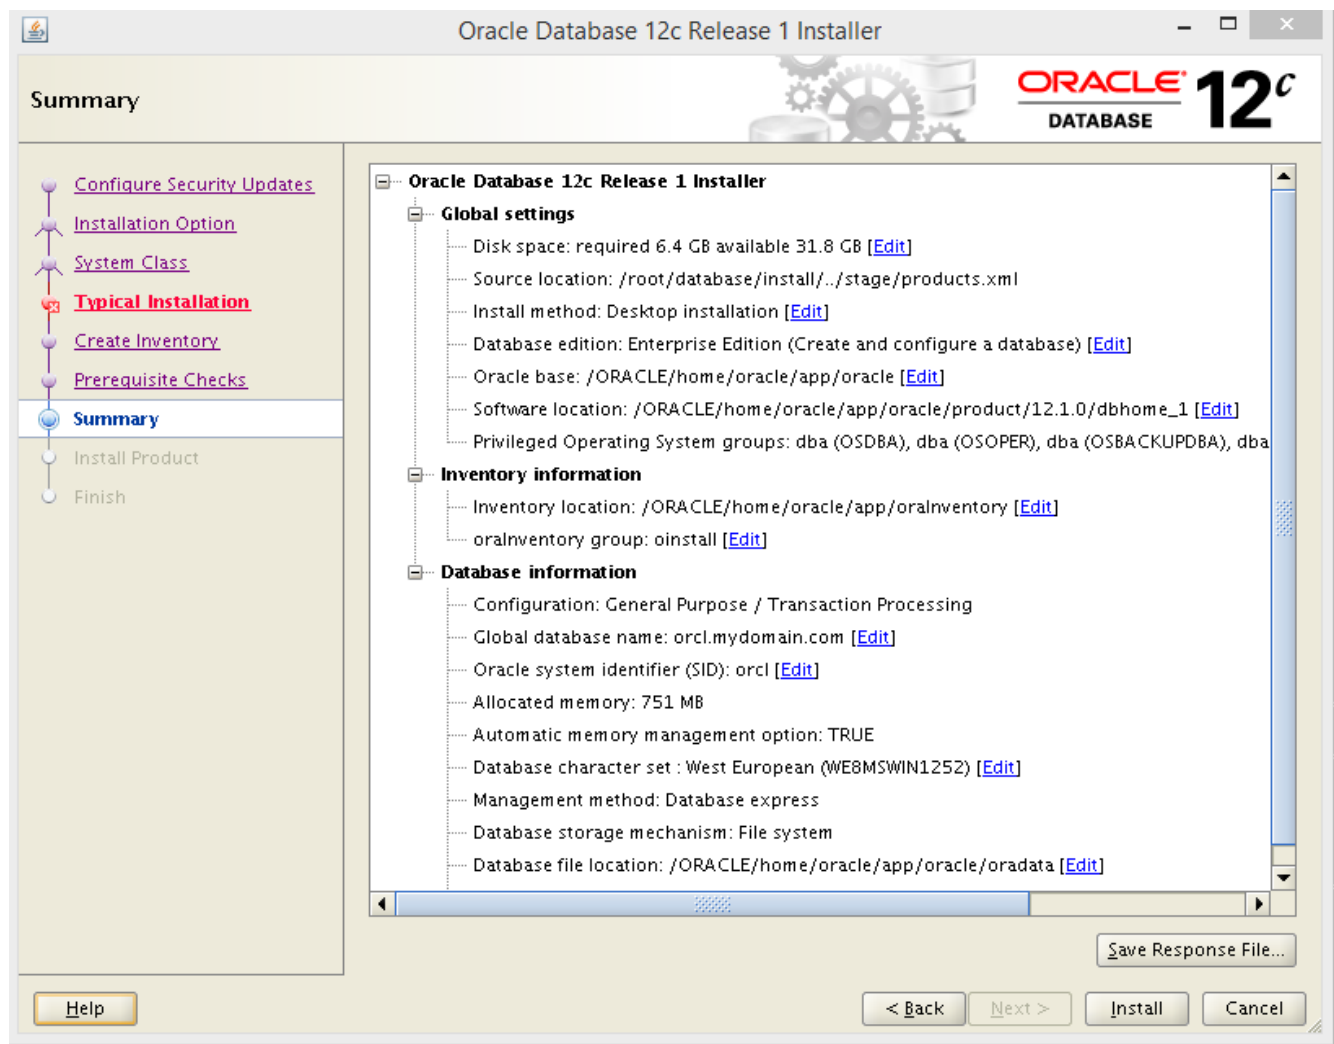
Next >

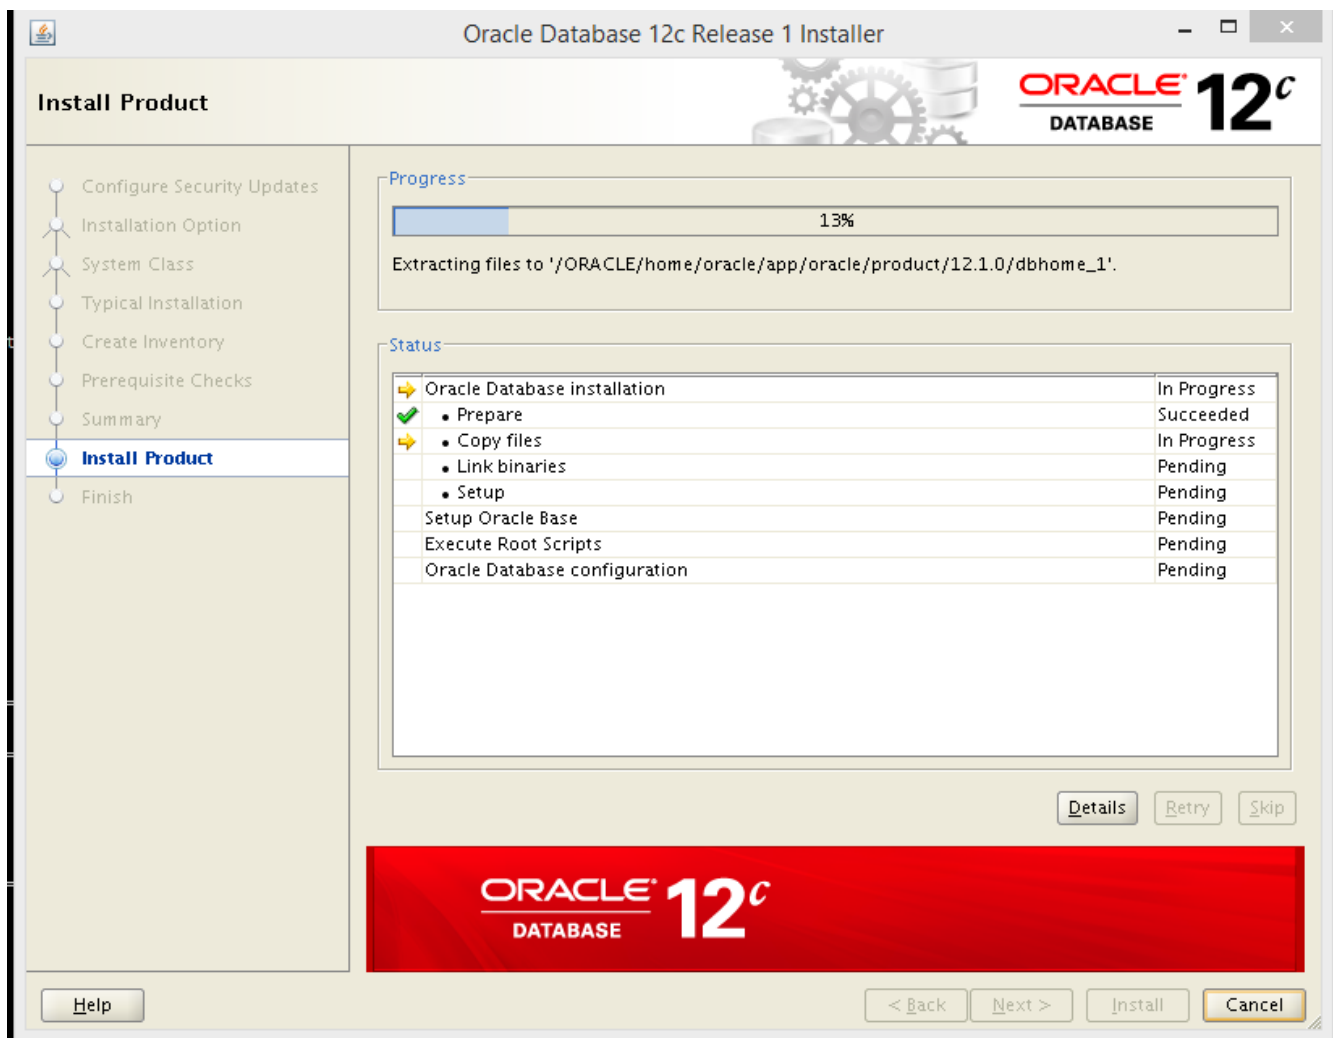
Install

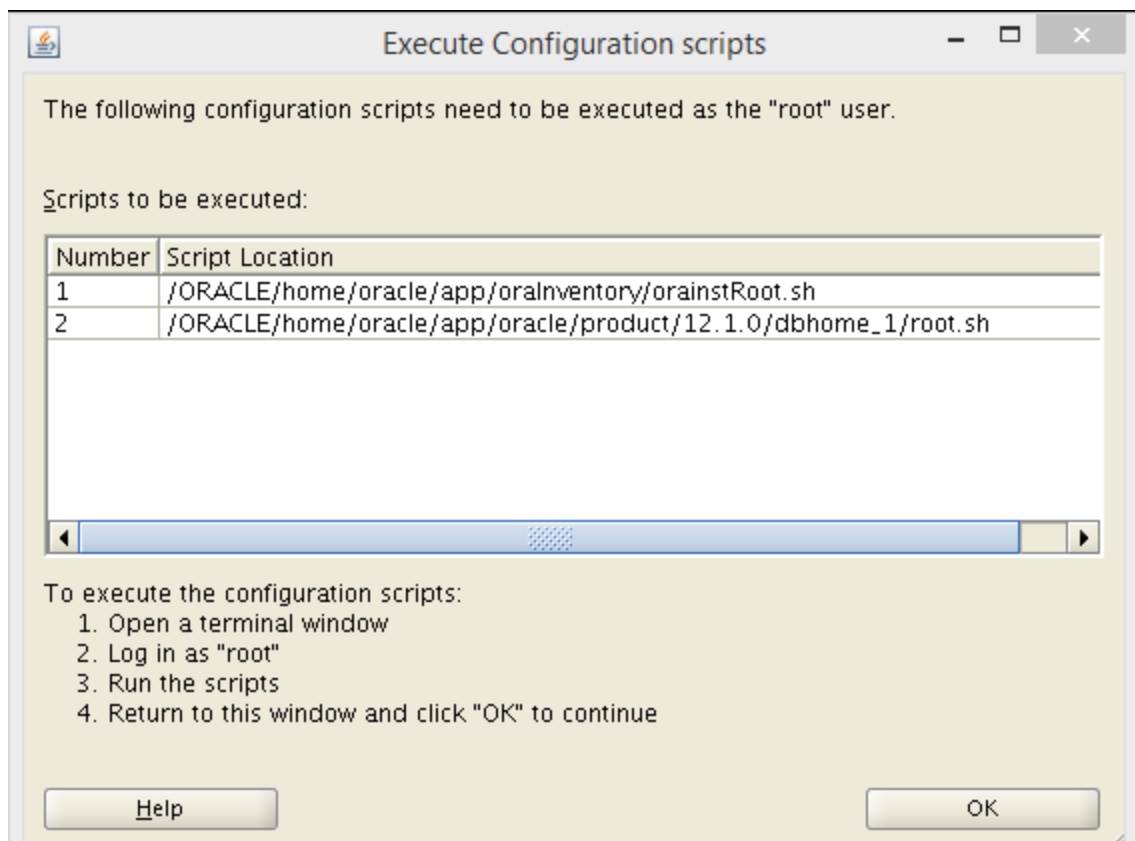
Cancel











On the node we are installing oracle on, please do this:

```
[root@node1 ~]# /ORACLE/home/oracle/app/oraInventory/orainstRoot.sh
```

Changing permissions of /ORACLE/home/oracle/app/oraInventory.

Adding read,write permissions for group.

Removing read,write,execute permissions for world.

Changing groupname of /ORACLE/home/oracle/app/oraInventory to oinstall.

The execution of the script is complete.

```
[root@node1 ~]# /ORACLE/home/oracle/app/oracle/product/12.1.0/dbhome_1/root.sh
```

Performing root user operation.

The following environment variables are set as:

ORACLE\_OWNER= oracle

ORACLE\_HOME= /ORACLE/home/oracle/app/oracle/product/12.1.0/dbhome\_1

Enter the full pathname of the local bin directory: [/usr/local/bin]:

Copying dbhome to /usr/local/bin ...

Copying oraenv to /usr/local/bin ...

Copying coraenv to /usr/local/bin ...

Creating /etc/oratab file...

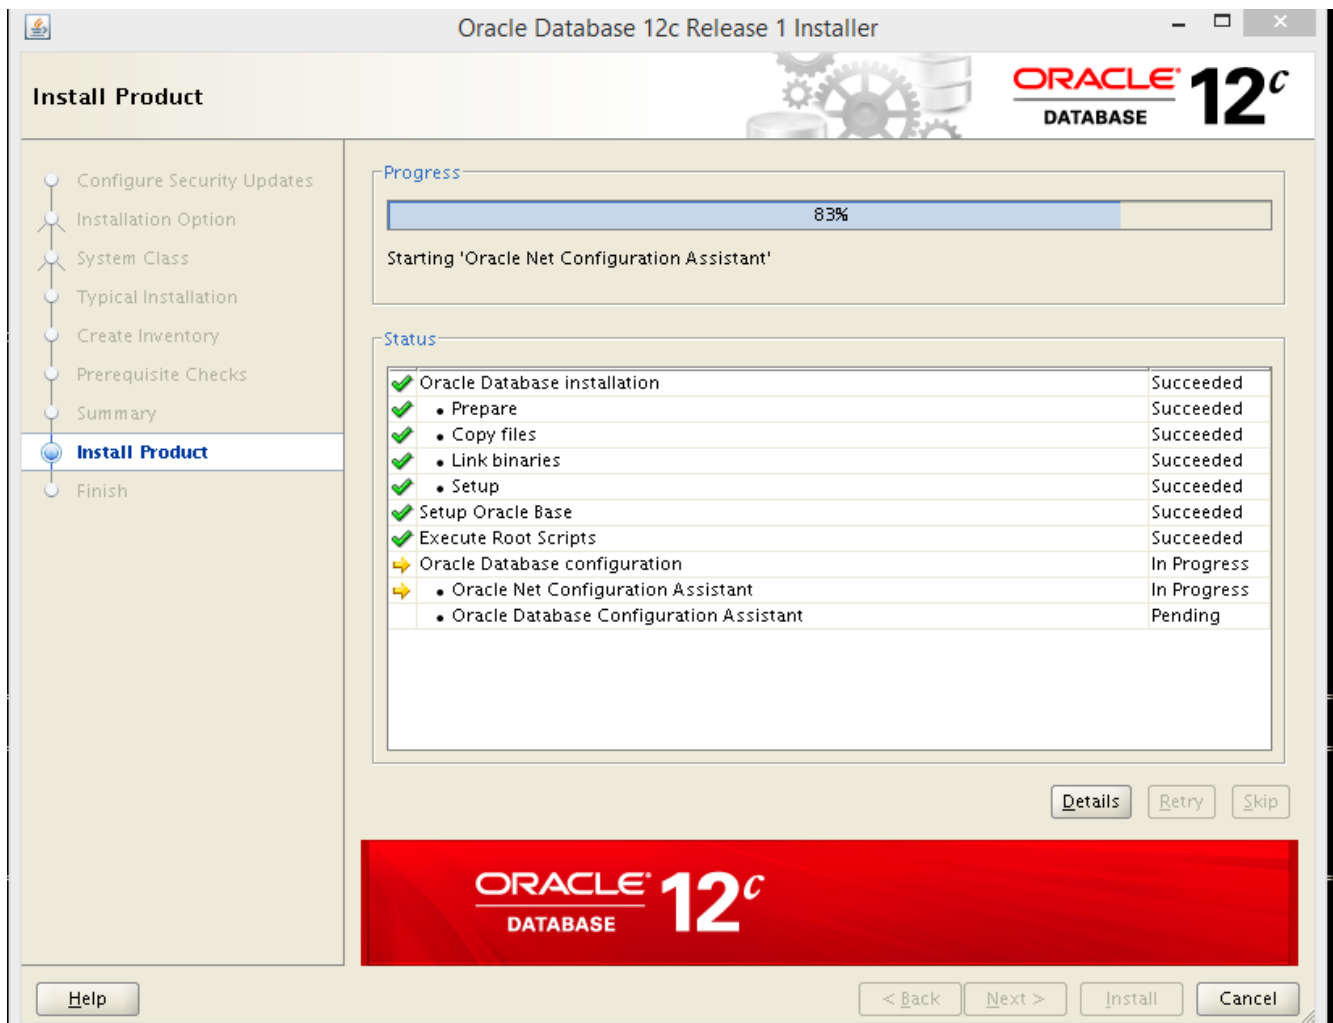
Entries will be added to the /etc/oratab file as needed by

Database Configuration Assistant when a database is created

Finished running generic part of root script.


Now product-specific root actions will be performed.


```
[root@node1 ~]#
```



### Progress

Clone database "orcl.mydomain.com" creation in progress...



	Steps	Status
	Copying database files	In Progress
	Creating and starting Oracle instance	
	Completing Database Creation	


[Activity Log](#)

[Alert Log](#)

[Stop](#)

Progress

Clone database "orcl.mydomain.com" creation completed



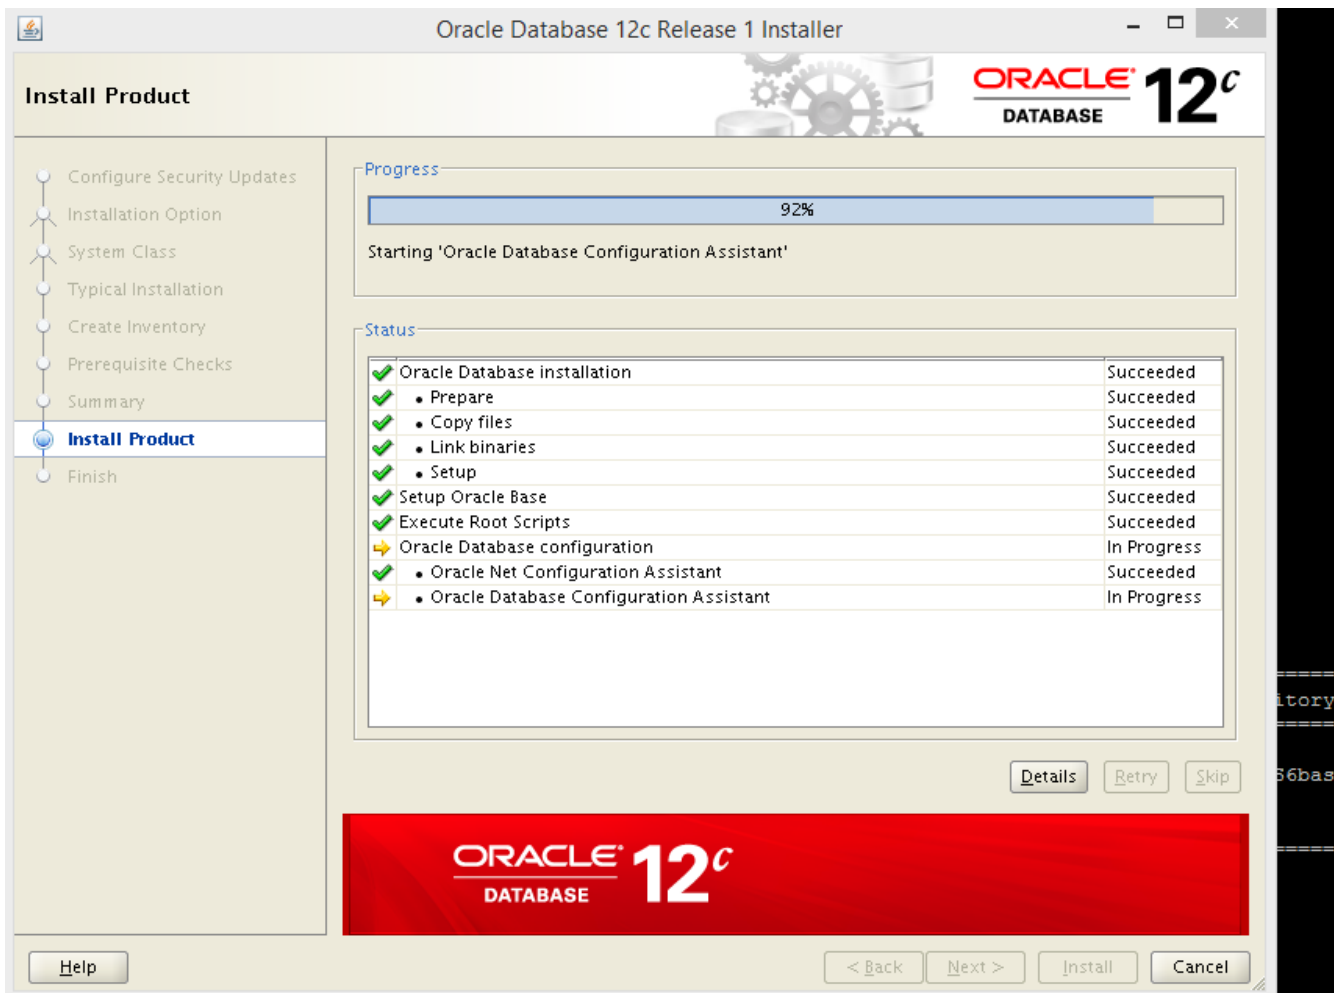
	Steps	Status
✓	Copying database files	Finished
✓	Creating and starting Oracle instance	Finished
✓	Completing Database Creation	Finished

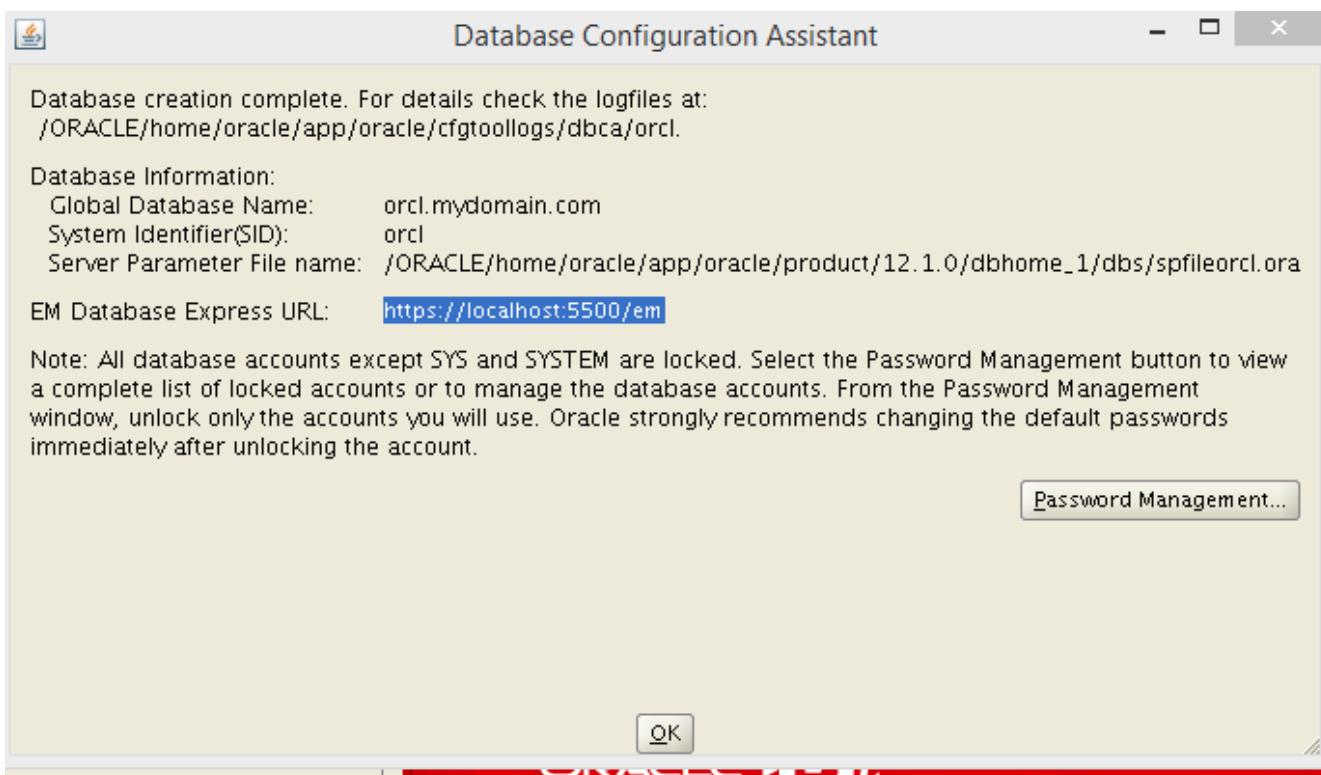
Activity Log

Alert Log

OK









NOW WE WILL BRING THE DATABASE UNDER CLUSTER CONTROL

```
[root@node1 init.d]# haconf -makerw
#we make the configuration file writeable
#open another window and run this command, more /opt/VRTSvcs/conf/config/main.cf
```

```
[root@node1 init.d]# hares -add listener_resource Netlsnr new_servicegroup
#we add a resource to the service group. We are adding the resource of VCS resource type NetLsnr and
#we are calling it listener_resource. The service group is our service group which we had created
#earlier, it is called, new_servicegroup
```

VCS NOTICE V-16-1-10242 Resource added. Enabled attribute must be set before agent monitors  
#This message tells us to set the “Enabled” attribute of the resource which we have added to “1” to  
#ensure that VCS starts monitoring this resource, we will do it.

```
[root@node1 init.d]# hares -modify listener_resource Owner oracle
#This is some important configuration required by Oracle. We specify the Owner attribute to be the
#user “oracle”.
```

```
[root@node1 init.d]# hares -modify listener_resource Home
/ORACLE/home/oracle/app/oracle/product/12.1.0/dbhome_1
#Another important Oracle specific attribute, we specify the home directory for user “oracle”. Please
#NOTE that this is on shared storage.
```

```
[root@node1 init.d]#
[root@node1 init.d]# hares -modify listener_resource TnsAdmin
/ORACLE/home/oracle/app/oracle/product/12.1.0/dbhome_1/network/admin/
#Another important Oracle resource attribute, we are modifying this hence the “hares -modify”, we are
#setting it to the location of the admin directory.
```

```
[root@node1 init.d]# hares -modify listener_resource Listener ""
[root@node1 init.d]# hares -modify listener_resource Enabled 1
#we finally asked VCS to monitor this resource by setting “Enabled” to “1”.
```

```
[root@node1 init.d]#
[root@node1 init.d]# hares -add oracledb_resource Oracle new_servicegroup
#Now we add the Oracle database proper as a resource to the service group.
```

VCS NOTICE V-16-1-10242 Resource added. Enabled attribute must be set before agent monitors

```
[root@node1 init.d]# hares -modify oracledb_resource Sid orcl
#This is another oracle specific attribute. WE had chosen “orcl” as the SID during our installation. We
#tell VCS this in this step.
```

```
[root@node1 init.d]# hares -modify oracledb_resource Owner oracle
#Setting owner attribute.
```

```
[root@node1 init.d]# hares -modify oracledb_resource Home  
/ORACLE/home/oracle/app/oracle/product/12.1.0/dbhome_1  
#Setting the location of Oracle home directory.
```

```
[root@node1 init.d]# hares -link listener_resource oracledb_resource  
#NOTE here we establish a link between the Oracle listener resource and the Oracle database resource.  
#NOTE, the format is hares -link Resource_Name Resource_it_depends_on  
#NOTE listener depends on the database
```

```
[root@node1 init.d]# hares -link oracledb_resource mount_resource  
#NOTE oracle database depends on the mount being present. We cannot start the database unless it is  
#present(mounted).
```

```
[root@node1 init.d]# hares -modify oracledb_resource Enabled 1  
#Asking VCS to start monitoring this oracle database resource.
```

```
[root@node1 init.d]# hares -modify listener_resource Critical 1  
#This means that the listener_resource is a "Critical" resource and must be ONLINE for the service  
#group to come up
```

```
[root@node1 init.d]# hares -modify oracledb_resource Critical 1
```

```
[root@node1 init.d]# haconf -dump  
#we write our configuration changes to the config file.  
[root@node1 init.d]#
```

NOTE: We have to copy certain oracle related files to node2, these are already present on node1 because we did our installation of oracle database on node1.

```
[root@node1 init.d]# cat /etc/oratab
#
```

```
# This file is used by ORACLE utilities.  It is created by root.sh
# and updated by either Database Configuration Assistant while creating
# a database or ASM Configuration Assistant while creating ASM instance.
```

```
# A colon, ':', is used as the field terminator.  A new line terminates
# the entry.  Lines beginning with a pound sign, '#', are comments.
```

```
#
```

```
# Entries are of the form:
```

```
# $ORACLE_SID:$ORACLE_HOME:<N|Y>:
```

```
#
```

```
# The first and second fields are the system identifier and home
# directory of the database respectively.  The third field indicates
# to the dbstart utility that the database should , "Y", or should not,
# "N", be brought up at system boot time.
```

```
#
```

```
# Multiple entries with the same $ORACLE_SID are not allowed.
```

```
#
```

```
#
```

```
orcl:/ORACLE/home/oracle/app/oracle/product/12.1.0/dbhome_1:N
```

```
[root@node1 init.d]#
```

so, do this:

```
[root@node1 init.d]# scp /etc/oratab node2:/etc
```

```
root@node2's password:
```

```
oratab
```

```
100% 803 0.8KB/s 00:00
```

```
[root@node1 init.d]# ls -al /etc/oratab
-rw-rw-r-- 1 oracle oinstall 803 Apr  4 17:10 /etc/oratab
[root@node1 init.d]#
```

So, change the permissions on node2 as well:

```
[root@node2 ~]# ls -al /etc/oratab
-rw-r--r-- 1 root root 803 Apr  4 17:57 /etc/oratab
[root@node2 ~]# chmod g+w /etc/oratab
[root@node2 ~]# chown oracle:oinstall /etc/oratab
[root@node2 ~]#
```

Searching for more files that may need to be copied over:

```
[root@node1 init.d]# find / -name *oraenv* -print
/usr/local/bin/coraenv
/usr/local/bin/oraenv
/ORACLE/home/oracle/app/oracle/product/12.1.0/dbhome_1/bin/coraenv
/ORACLE/home/oracle/app/oracle/product/12.1.0/dbhome_1/bin/oraenv
[root@node1 init.d]# ls -al /usr/local/bin/coraenv
-rwxr-xr-x 1 oracle root 6583 Apr  4 17:09 /usr/local/bin/coraenv
[root@node1 init.d]# ls -al /usr/local/bin/oraenv
-rwxr-xr-x 1 oracle root 7012 Apr  4 17:09 /usr/local/bin/oraenv
[root@node1 init.d]#
```

So we need to copy these two files from node1 to node2

/usr/local/bin/coraenv

/usr/local/bin/oraenv

and change the permissions appropriately

```
[root@node1 init.d]# scp /usr/local/bin/oraenv node2:/usr/local/bin
root@node2's password:
Permission denied, please try again.
root@node2's password:
Permission denied, please try again.
root@node2's password:
oraenv                                     100% 7012   6.9KB/s   00:00
[root@node1 init.d]#
```

```
[root@node2 bin]# chmod o+rwx /usr/local/bin/coraenv
[root@node2 bin]# chmod o+rwx /usr/local/bin/oraenv
```

```
[root@node2 bin]# chown oracle:root /usr/local/bin/oraenv
[root@node2 bin]# chown oracle:root /usr/local/bin/coraenv
[root@node2 bin]#
```

Please do the same for /usr/local/bin/dbhome

```
[root@node1 ~]# hastatus -sum
```

```
-- SYSTEM STATE
```

```
-- System      State      Frozen
```

```
A node1        RUNNING    0
```

```
A node2        RUNNING    0
```

```
-- GROUP STATE
```

```
-- Group      System      Probed   AutoDisabled  State
```

```
B ClusterService node1      Y      N      OFFLINE
```

```
B ClusterService node2      Y      N      ONLINE
```

```
B new_servicegroup node1     Y      N      OFFLINE
```

```
B new_servicegroup node2     Y      N      ONLINE
```

test that the cluster works as intended:

```
root@node1 ~]# hagr -switch new_servicegroup -to node1
```

```
[root@node1 ~]# hastatus -sum
```

```
-- SYSTEM STATE
```

```
-- System      State      Frozen
```

```
A node1        RUNNING    0
```



A node2            RUNNING            0

-- GROUP STATE

-- Group	System	Probed	AutoDisabled	State
B ClusterService	node1	Y	N	OFFLINE
B ClusterService	node2	Y	N	ONLINE
B new_servicegroup	node1	Y	N	ONLINE
B new_servicegroup	node2	Y	N	OFFLINE

[root@node1 ~]#

CONGRATULATIONS! We have successfully configured Oracle as a failover service in VCS.

## WORK Related Stuff

At work you will probably be asked to monitor and troubleshoot an existing installation of VCS. The most important thing to remember is to monitor the VCS log files. These are located in /var/VRTSvcs/log/engine\_A.log

For your practice, open two terminal sessions(two putty sessions), run this command:

```
root@node1 ~]# hagr -switch new_servicegroup -to node2
```

In the second putty session, run this command:

```
root@node1 ~]#tail -f /var/VRTSvcs/log/engine_A.log
```

This will show you everything that is happening.

As part of routine maintenance, you might be asked by the DBA to “freeze the cluster” as they do an Oracle upgrade or apply security patches to Oracle etc. Freezing the cluster is an erroneous description of what they want. They really want you to freeze the Oracle service group.

Freezing a service group: We freeze a service group to prevent it from failing over to another system. This freezing process stops all online and offline procedures on the service group.

Unfreeze a frozen service group to perform online or offline operations on the service group.

To freeze a service group (disable online, offline, and failover operations)

Type the following command:

```
hagr -freeze service_group [-persistent]
```

The option -persistent enables the freeze to be remembered when the cluster is rebooted.

To unfreeze a service group (reenable online, offline, and failover operations)

Type the following command:

```
hagr -unfreeze service_group [-persistent]
```

NOTE: Try this by yourself

About critical and non-critical resources

The Critical attribute for a resource defines whether a service group fails over when the resource faults. If a resource is configured as non-critical (by setting the Critical attribute to 0) and no resources depending on the failed resource are critical, the service group will not fail over. VCS takes the failed resource offline and updates the group status to ONLINE|PARTIAL. The attribute also determines whether a service group tries to come online on another node if, during the group's online process, a resource fails to come online.

Please try the following commands:

```
lltconfig
gabconfig -a
cat /etc/sysconfig/llt
cat /etc/sysconfig/gab
cat /etc/llthosts
cat /etc/llttab
cat /etc/gabtab
```

If someone asks you, how to do it, say this: "I have done it with EMC powerpath installed"

here are the steps:

1) Request lun from storage team. Request lun be assigned a specific VNX tag, example "db\_extended".

On both nodes of the cluster run these commands:

Scan for lun's and have the os detect them:

```
2) echo "- - -" > /sys/class/scsi_host/host0/scan
```

```
echo "- - -" > /sys/class/scsi_host/host1/scan
```

```
echo "- - -" > /sys/class/scsi_host/host2/scan
```

```
echo "- - -" > /sys/class/scsi_host/host3/scan
```

3) Write label to disk

Fdisk /dev/sd\* , option w.

4) Have powerpath detect them:

```
powermt config
```

5) Apply and save powerpath configuration:

```
powermt save
```

6) Confirm disk is seen by powerpath. Look for requested lun tag

```
powermt display dev=all
```

7) Have Veritas pick up disk

```
vxdctl enable
```

8)Verify Veritas sees new disk and is named disk properly. Disk name should match emcpower name. If not fix with vxedit.

`vxdisk -e list`

On only the node where the volume group is online, run these commands:

9)Add lun to the Veritas diskgroup:

`vxdiskadm, option 1`

10)Use the vxresize command to resize the volume

`vxresize -x -g mydg myvol <new size>G [newdiskname]`

## WORK RELATED STUFF:

1)INCREASE FILE SYSTEM SIZE OF A VCS CONTROLLED FILESYSTEM

2)FREEZE A CLUSTER FOR MAINTENANCE DONE BY THE DBA ON THE CLUSTER

3)UPGRADE THE VCS VERSION ON A VCS CLUSTER

4)SWITCH A SERVICE GROUP FROM ONE NODE TO ANOTHER NODE

5)OFFLINE AND ONLINE RESOURCES

6)ADD A NODE TO AN EXISTING VCS CLUSTER, SO FOR A TWO NODE CLUSTER, MAKE IT A THREE NODE CLUSTER, ETC.

7)INCREASE/CHANGE THE MONITORING RECEIPTS OF THE CLUSTER

# Veritas Cluster Cheat sheet

[LLT and GAB Commands](#) | [Port Membership](#) | [Daemons](#) | [Log Files](#) | [Dynamic Configuration](#) | [Users](#) | [Resources](#) | [Resource Agents](#) | [Service Groups](#) | [Clusters](#) | [Cluster Status](#) | [System Operations](#) | [Service Group Operations](#) | [Resource Operations](#) | [Agent Operations](#) | [Starting and Stopping](#)

## LLT and GRAB

VCS uses two components, LLT and GAB to share data over the private networks among systems.

These components provide the performance and reliability required by VCS.

LLT	LLT (Low Latency Transport) provides fast, kernel-to-kernel comms and monitors network connections. The system admin configures the LLT by creating a configuration file (llttab) that describes the systems in the cluster and private network connections among them. The LLT runs in layer 2 of the network stack
GAB	GAB (Group membership and Atomic Broadcast) provides the global message ordering required to maintain a synchronised state among the systems, and monitors data such as that required by the VCS heartbeat utility. The system admin configures the driver by creating a configuration file ( gabtab).

## LLT and GAB files

<a href="#">/etc/llthosts</a>	The file is a database, containing one entry per system, that links the LLT system ID with the hosts name. The file is identical on each server in the cluster.
<a href="#">/etc/llttab</a>	The file contains information that is derived during installation and is used by the utility lltdconfig.
<a href="#">/etc/gabtab</a>	The file contains the information needed to configure the GAB driver. This file is used by the gabconfig utility.
<a href="#">/etc/VRTSvcs/conf/config/main.cf</a>	The VCS configuration file. The file contains the information that defines the cluster and its systems.

## Gabtab Entries

<pre>/sbin/gabdiskconf -i /dev/dsk/clt2d0s2 -s 16 -S 1123 /sbin/gabdiskconf -i /dev/dsk/clt2d0s2 -s 144 -S 1124 /sbin/gabdiskhb -a /dev/dsk/clt2d0s2 -s 16 -p a -s 1123 /sbin/gabdiskhb -a /dev/dsk/clt2d0s2 -s 144 -p h -s 1124 /sbin/gabconfig -c -n2</pre>	
---	--

<a href="#">gabdiskconf</a>	-i Initialises the disk region -s Start Block -S Signature
<a href="#">gabdiskhb (heartbeat disks)</a>	-a Add a gab disk heartbeat resource -s Start Block -p Port -S Signature
<a href="#">gabconfig</a>	-c Configure the driver for use -n Number of systems in the cluster.

## LLT and GAB Commands

Verifying that links are active for LLT	lltstat -n
verbose output of the lltstat command	lltstat -nvv   more
open ports for LLT	lltstat -p
display the values of LLT configuration directives	lltstat -c
lists information about each configured LLT link	lltstat -l
List all MAC addresses in the cluster	lltconfig -a list
stop the LLT running	lltconfig -U
start the LLT	lltconfig -c
verify that GAB is operating	gabconfig -a  Note: port a indicates that GAB is communicating, port h indicates VCS is started
stop GAB running	gabconfig -U
start the GAB	gabconfig -c -n <number of nodes>
override the seed values in the gabtab file	gabconfig -c -x

## GAB Port Membership

List Membership	gabconfig -a
Unregister port f	/opt/VRTS/bin/fsclustadm cfsdeinit
Port Function	a gab driver b I/O fencing (designed to guarantee data integrity) d ODM (Oracle Disk Manager) f CFS (Cluster File System) h VCS (VERITAS Cluster Server: high availability daemon) o VCSMM driver (kernel module needed for Oracle and VCS interfacing) q QuickLog daemon v CVM (Cluster Volume Manager) w vxconfigd (module for cvm)

## Cluster daemons

High Availability Daemon	had
Companion Daemon	hashadow
Resource Agent daemon	<resource>Agent
Web Console cluster management daemon	CmdServer

## Cluster Log Files

Log Directory	/var/VRTSvcs/log
primary log file (engine log file)	/var/VRTSvcs/log/engine_A.log

## Starting and Stopping the cluster

"-stale" instructs the engine to treat the local config as stale "-force" instructs the engine to treat a stale config as a valid one	hasstart [-stale -force]
Bring the cluster into running mode from a stale state using the configuration file from a particular server	hasys -force <server_name>
stop the cluster on the local server but leave the application/s running, do not failover the application/s	hastop -local
stop cluster on local server but evacuate (failover) the application/s to another node within the cluster	hastop -local -evacuate
stop the cluster on all nodes but leave the application/s running	hastop -all -force

--	--

## Cluster Status

display cluster summary	hastatus -summary
continually monitor cluster	hastatus
verify the cluster is operating	hasys -display

## Cluster Details

information about a cluster	haclus -display
value for a specific cluster attribute	haclus -value <attribute>
modify a cluster attribute	haclus -modify <attribute name> <new>
Enable LinkMonitoring	haclus -enable LinkMonitoring
Disable LinkMonitoring	haclus -disable LinkMonitoring

## Users

add a user	hauser -add <username>
modify a user	hauser -update <username>
delete a user	hauser -delete <username>
display all users	hauser -display

## System Operations

add a system to the cluster	hasys -add <sys>
delete a system from the cluster	hasys -delete <sys>
Modify a system attributes	hasys -modify <sys> <modify options>
list a system state	hasys -state
Force a system to start	hasys -force
Display the systems attributes	hasys -display [-sys]
List all the systems in the cluster	hasys -list
Change the load attribute of a system	hasys -load <system> <value>
Display the value of a systems nodeid (/etc/llthosts)	hasys -nodeid
Freeze a system (No offlining system, No groups online)	hasys -freeze [-persistent][-evacuate]  Note: main.cf must be in write mode
Unfreeze a system (reenable groups and resource back online)	hasys -unfreeze [-persistent]  Note: main.cf must be in write mode

## Dynamic Configuration

The VCS configuration must be in read/write mode in order to make changes. When in read/write mode the configuration becomes stale, a .stale file is created in \$VCS\_CONF/conf/config. When the configuration is put back into read only mode the .stale file is removed.

Change configuration to read/write mode	haconf -makerw
Change configuration to read-only mode	haconf -dump -makero

Check what mode cluster is running in	<pre>haclus -display  grep -i 'readonly'</pre> <p>0 = write mode 1 = read only mode</p>
Check the configuration file	<pre>hacf -verify /etc/VRTSvcs/conf/config</pre> <p>Note: you can point to any directory as long as it has main.cf and types</p>
convert a main.cf file into cluster commands	<pre>hacf -cftocmd /etc/VRTSvcs/conf/config -dest /tmp</pre>
convert a command file into a main.cf file	<pre>hacf -cmdtoctf /tmp -dest /etc/VRTSvcs/conf/config</pre>

## Service Groups

add a service group	<pre>haconf -makerw hagrp -add groupw hagrp -modify groupw SystemList sun1 1 sun2 2 hagrp -autoenable groupw -sys sun1 haconf -dump -makero</pre>
delete a service group	<pre>haconf -makerw hagrp -delete groupw haconf -dump -makero</pre>
change a service group	<pre>haconf -makerw hagrp -modify groupw SystemList sun1 1 sun2 2 sun3 3 haconf -dump -makero</pre> <p>Note: use the "hagrp -display &lt;group&gt;" to list attributes</p>
list the service groups	<pre>hagrp -list</pre>
list the groups dependencies	<pre>hagrp -dep &lt;group&gt;</pre>
list the parameters of a group	<pre>hagrp -display &lt;group&gt;</pre>
display a service group's resource	<pre>hagrp -resources &lt;group&gt;</pre>
display the current state of the service group	<pre>hagrp -state &lt;group&gt;</pre>
clear a faulted non-persistent resource in a specific grp	<pre>hagrp -clear &lt;group&gt; [-sys] &lt;host&gt; &lt;sys&gt;</pre>
Change the system list in a cluster	<pre># remove the host hagrp -modify grp_zlnrssid SystemList -delete &lt;hostname&gt;  # add the new host (don't forget to state its position) hagrp -modify grp_zlnrssid SystemList -add &lt;hostname&gt; 1  # update the autostart list hagrp -modify grp_zlnrssid AutoStartList &lt;host&gt; &lt;host&gt;</pre>

## Service Group Operations

Start a service group and bring its resources online	<pre>hagrp -online &lt;group&gt; -sys &lt;sys&gt;</pre>
Stop a service group and takes its resources offline	<pre>hagrp -offline &lt;group&gt; -sys &lt;sys&gt;</pre>
Switch a service group from system to another	<pre>hagrp -switch &lt;group&gt; to &lt;sys&gt;</pre>
Enable all the resources in a group	<pre>hagrp -enableresources &lt;group&gt;</pre>
Disable all the resources in a group	<pre>hagrp -disableresources &lt;group&gt;</pre>
Freeze a service group (disable onlining and offlining)	<pre>hagrp -freeze &lt;group&gt; [-persistent]</pre>

	note: use the following to check "hagrp -display <group>   grep T
Unfreeze a service group (enable onlining and offlining)	hagrp -unfreeze <group> [-persistent]  note: use the following to check "hagrp -display <group>   grep T
Enable a service group. Enabled groups can only be brought online	haconf -makerw hagrp -enable <group> [-sys] haconf -dump -makero  Note to check run the following command "hagrp -display   grep Er
Disable a service group. Stop from bringing online	haconf -makerw hagrp -disable <group> [-sys] haconf -dump -makero  Note to check run the following command "hagrp -display   grep Er
Flush a service group and enable corrective action.	hagrp -flush <group> -sys <system>

## Resources

add a resource	haconf -makerw hares -add appDG DiskGroup groupw hares -modify appDG Enabled 1 hares -modify appDG DiskGroup appdg hares -modify appDG StartVolumes 0 haconf -dump -makero
delete a resource	haconf -makerw hares -delete <resource> haconf -dump -makero
change a resource	haconf -makerw hares -modify appDG Enabled 1 haconf -dump -makero  Note: list parameters "hares -display <resource>"
change a resource attribute to be globally wide	hares -global <resource> <attribute> <value>
change a resource attribute to be locally wide	hares -local <resource> <attribute> <value>
list the parameters of a resource	hares -display <resource>
list the resources	hares -list
list the resource dependencies	hares -dep

## Resource Operations

Online a resource	hares -online <resource> [-sys]
Offline a resource	hares -offline <resource> [-sys]
display the state of a resource( offline, online, etc)	hares -state
display the parameters of a resource	hares -display <resource>
Offline a resource and propagate the command to its children	hares -offprop <resource> -sys <sys>
Cause a resource agent to immediately monitor the resource	hares -probe <resource> -sys <sys>
Clearing a resource (automatically initiates the onlining)	hares -clear <resource> [-sys]



## Resource Types

Add a resource type	hatype -add <type>
Remove a resource type	hatype -delete <type>
List all resource types	hatype -list
Display a resource type	hatype -display <type>
List a particular resource type	hatype -resources <type>
Change a particular resource types attributes	hatype -value <type> <attr>

## Resource Agents

add a agent	pkgadd -d . <agent package>
remove a agent	pkgrm <agent package>
change a agent	n/a
list all ha agents	haagent -list
Display agents run-time information i.e has it started, is it running ?	haagent -display <agent_name>
Display agents faults	haagent -display  grep Faults

## Resource Agent Operations

Start an agent	haagent -start <agent_name>[-sys]
Stop an agent	haagent -stop <agent_name>[-sys]

## Interview Questions

### The ultimate Veritas Cluster Server (VCS) interview questions

#### Basics

What are the different service group types ?

Service groups can be one of the 3 type :

1. Failover – Service group runs on one system at a time.

2. Parallel – Service group runs on multiple systems simultaneously.
3. Hybrid – Used in replicated data clusters (disaster recovery setups). SG behaves as Failover within the local cluster and Parallel for the remote cluster.

Where is the VCS main configuration file located ?

The main.cf file contains the configuration of the entire cluster and is located in the directory /etc/VRTSvcs/conf/config.

How to set VCS configuration file (main.cf) ro/rw ?

To set the configuration file in read-only/read-write :

```
# haconf -dump -makero    (Dumps in memory configuration to main.cf and makes it read-only)
```

```
# haconf -makerw          (Makes configuration writable)
```

Where is the VCS engine log file located ?

The VCS cluster engine logs is located at /var/VRTSvcs/log/engine\_A.log. We can either directly view this file or use command line to view it :

```
# hamsg engine_A
```

How to check the complete status of the cluster

To check the status of the entire cluster :

```
# hastatus -sum
```

How to verify the syntax of the main.cf file

To verify the syntax of the main.cf file just mention the absolute directory path to the main.cf file :

```
# hacf -verify /etc/VRTSvcs/conf/config
```

What are the different resource types ?

1. Persistent : VCS can only monitor these resources but can not offline or online them.
2. On-Off : VCS can start and stop On-Off resource type. Most resources fall in this category.
3. On-Only : VCS starts On-Only resources but does not stop them. An example would be NFS daemon. VCS can start the NFS daemon if required, but can not take it offline if the associated service group is take offline.

Explain the steps involved in Offline VCS configuration

1. Save and close the configuration :

```
# haconf -dump -makero
```

2. Stop VCS on all nodes in the cluster :

```
# hastop -all
```

3. Edit the configuration file after taking the backup and do the changes :

```
# cp -p /etc/VRTSvcs/conf/config/main.cf /etc/VRTSvcs/conf/config/main.cf_17march
```

```
# vi /etc/VRTSvcs/conf/config/main.cf
```

4. Verify the configuration file syntax :

```
# hacf -verify /etc/VRTSvcs/conf/config/
```

5. start the VCS on the system with modified main.cf file :

```
# hastart
```

6. start VCS on other nodes in the cluster.

Note : This can be done in another way by just stopping VCS and leaving services running to minimize the downtime. (hastop -all -force

GAB, LLT and HAD

What is GAB, LLT and HAD and whats their functionalities ?

GAB, LLT and HAD forms the basic building blocks of vcs functionality.

LLT (low latency transport protocol) – LLT transmits the heartbeats over the interconnects. It is also used to distribute the inter system communication traffic equally among all the interconnects.

GAB (Group membership services and atomic broadcast) – The group membership service part of GAB maintains the overall cluster membership information by tracking the heartbeats sent over LLT interconnects. The atomic broadcast of cluster membership ensures that every node in the cluster has same information about every resource and service group in the cluster.

HAD (High Availability daemon) – the main VCS engine which manages the agents and service group. It is in turn monitored by a daemon named hashadow.

What are the various GAB ports and their functionalities ?

- a --> gab driver
- b --> I/O fencing (to ensure data integrity)
- d --> ODM (Oracle Disk Manager)
- f --> CFS (Cluster File System)
- h --> VCS (VERITAS Cluster Server: high availability daemon, HAD)
- o --> VCSMM driver (kernel module needed for Oracle and VCS interface)
- q --> QuickLog daemon
- v --> CVM (Cluster Volume Manager)
- w --> vxconfigd (module for cvm)

How to check the status of various GAB ports on the cluster nodes

To check the status of GAB ports on various nodes :

```
# gabconfig -a
```

Whats the maximum number of LLT links (including high and low priority) can a cluster have ?

A cluster can have a maximum of 8 LLT links including high and low priority LLT links.

How to check the detailed status of LLT links ?

The command to check detailed LLT status is :

```
# llstat -nvv
```

What are the various LLT configuration files and their function ?

LLT uses /etc/llttab to set the configuration of the LLT interconnects.

```
# cat /etc/llttab
```

```
set-node node01
```

```
set-cluster 02
```

```
link nxge1 /dev/nxge1 - ether - -
```

```
link nxge2 /dev/nxge2 - ether - -  
link-lowpri /dev/nxge0 - ether - -
```

Here, set-cluster -> unique cluster number assigned to the entire cluster [ can have a value ranging between 0 to (64k – 1) ]. It should be unique across the organization.

set-node -> a unique number assigned to each node in the cluster. Here the name node01 has a corresponding unique node number in the file /etc/llthosts. It can range from 0 to 31.

Another configuration file used by LLT is – /etc/llthosts. It has the cluster-wide unique node number and nodename as follows:

```
# cat /etc/llthosts  
0 node01  
1 node02
```

LLT has an another optional configuration file : /etc/VRTSvcs/conf/sysname. It contains short names for VCS to refer. It can be used by VCS to remove the dependency on OS hostnames.

What are various GAB configuration files and their function ?

The file /etc/gabtab contains the command to start the GAB.

```
# cat /etc/gabtab  
/sbin/gabconfig -c -n 4  
here -n 4 -> number of nodes that must be communicating in order to start VCS.
```

How to start/stop GAB

The commands to start and stop GAB are :

```
# gabconfig -c      (start GAB)  
# gabconfig -U      (stop GAB)
```

How to start/stop LLT

The commands to stop and start LLT are :

```
# lltconfig -c      -> start LLT  
# lltconfig -U      -> stop LLT (GAB needs to stopped first)
```

What's a GAB seeding and why manual GAB seeding is required ?

The GAB configuration file /etc/gabtab defines the minimum number of nodes that must be communicating for the cluster to start. This is called as GAB seeding.

In case we don't have sufficient number of nodes to start VCS [ may be due to a maintenance activity ], but have to do it anyways, then we have to do what is called as manual seeding by firing below command on each of the nodes.

```
# gabconfig -c -x
```

How to start HAD or VCS ?

To start HAD or VCS on all nodes in the cluster, the hastart command need to be run on all nodes individually.

```
# hastart
```

What are the various ways to stop HAD or VCS cluster ?

The command hastop gives various ways to stop the cluster.

```
# hastop -local
```

```
# hstop -local -evacuate
# hstop -local -force
# hstop -all -force
# hstop -all
-local -> Stops service groups and VCS engine [HAD] on the node where it is fired
-local -evacuate -> migrates Service groups on the node where it is fired and stops HAD on the same node only
-local -force -> Stops HAD leaving services running on the node where it is fired
-all -force -> Stops HAD on all the nodes of cluster leaving the services running
-all -> Stops HAD on all nodes in cluster and takes service groups offline
```

## Resource Operations

How to list all the resource dependencies

To list the resource dependencies :

```
# hares -dep
How to enable/disable a resource ?
# hares -modify [resource_name] Enabled 1    (To enable a resource)
# hares -modify [resource_name] Enabled 0    (To disable a resource)
How to list the parameters of a resource
To list all the parameters of a resource :
```

```
# hares -display [resource]
```

## Service group operations

How to add a service group(a general method) ?

In general, to add a service group named SG with 2 nodes (node01 and node02) :

```
haconf -makerw
hagrp -add SG
hagrp -modify SG SystemList node01 0 node02 1
hagrp -modify SG AutoStartList node02
haconf -dump -makero
How to check the configuration of a service group – SG ?
To see the service group configuration :
```

```
# hagrp -display SG
How to bring service group online/offline ?
To online/offline the service group on a particular node :
```

```
# hagrp -online [service-group] -sys [node]    (Online the SG on a particular node)
# hagrp -offline [service-group] -sys [node]    (Offline the SG on particular node)
The -any option when used instead of the node name, brings the SG online/offline based on SG's failover policy.
```

```
# hagrp -online [service-group] -any
# hagrp -offline [service-group] -any
```

How to switch service groups ?

The command to switch the service group to target node :

```
# hagr -switch [service-group] -to [target-node]
```

How to freeze/unfreeze a service group and what happens when you do so ?

When you freeze a service group, VCS continues to monitor the service group, but does not allow it or the resources under it to be taken offline or brought online. Failover is also disabled even when a resource faults. When you unfreeze the SG, it starts behaving in the normal way.

To freeze/unfreeze a Service Group temporarily :

```
# hagr -freeze [service-group]
```

```
# hagr -unfreeze [service-group]
```

To freeze/unfreeze a Service Group persistently (across reboots) :

```
# hagr -freeze -persistent [service-group]
```

```
# hagr -unfreeze [service-group] -persistent
```

Communication failures : Jeopardy, split brain

What's a Jeopardy membership in VCS clusters

When a node in the cluster has only the last LLT link intact, the node forms a regular membership with other nodes with which it has more than one LLT link active and a Jeopardy membership with the node with which it has only one LLT link active.

Jeopardy in VCS cluster

Effects of Jeopardy : (considering example in diagram above)

1. Jeopardy membership formed only for node03
2. Regular membership between node01, node02, node03
3. Service groups SG01, SG02, SG03 continue to run and other cluster functions remain unaffected.
4. If node03 faults or last link breaks, SG03 is not started on node01 or node02. This is done to avoid data corruption, as in case the last link is broken the nodes node02 and node01 may think that node03 is down and try to start SG03 on them. This may lead to data corruption as same service group may be online on 2 systems.
5. Failover due to resource fault or operator request would still work.

How to recover from a Jeopardy membership ?

To recover from Jeopardy, just fix the failed link(s) and GAB automatically detects the new link(s) and the Jeopardy membership is removed from node.

What's a split brain condition ?

Split brain occurs when all the LLT links fail simultaneously. Here systems in the cluster fail to identify whether it is a system failure or an interconnect failure. Each mini-cluster thus formed thinks that it is the only cluster that's active at the moment and tries to start the service groups on the other mini-cluster which it thinks is down. Similar thing happens to the other mini-cluster and this may lead to a simultaneous access to the storage and can cause data corruption.

What is I/O fencing and how it prevents split brain ?

VCS implements I/O fencing mechanism to avoid a possible split-brain condition. It ensure data integrity and data protection. I/O fencing driver uses SCSI-3 PGR (persistent group reservations) to fence off the data in case of a possible split brain scenario.

i:o fencing in vcs

In case of a possible split brain

As show in the figure above assume that node01 has key “A” and node02 has key “B”.

1. Both nodes think that the other node has failed and start racing to write their keys to the coordinator disks.
2. node01 manages to write the key to majority of disks i.e. 2 disks
3. node02 panics
4. node01 now has a perfect membership and hence Service groups from node02 can be started on node01

Whats the difference between MultiNICA and MultiNICB resource types ?

MultiNICA and IPMultiNIC

- supports active/passive configuration.
- Requires only 1 base IP (test IP).
- Does not require to have all IPs in the same subnet.

MultiNICB and IPMultiNICB

- supports active/active configuration.
- Faster failover than the MultiNICA.
- Requires IP address for each interface.

Troubleshooting

How to flush a service group and when its required ?

Flushing of a service group is required when, agents for the resources in the service group seems suspended waiting for resources to be taken online/offline. Flushing a service group clears any internal wait states and stops VCS from attempting to bring resources online.

To flush the service group SG on the cluster node, node01 :

```
# hagrps -flush [SG] -sys node01
```

How to clear resource faults ?

To clear a resource fault, we first have to fix the underlying problem.

1. For persistent resources :

Do not do anything and wait for the next OfflineMonitorInterval (default – 300 seconds) for the resource to become online.

2. For non-persistent resources :

Clear the fault and probe the resource on node01 :

```
# hares -clear [resource_name] -sys node01
```

```
# hares -probe [resource_name] -sys node01
```

How to clear resources with ADMIN\_WAIT state ?

If the ManageFaults attribute of a service group is set to NONE, VCS does not take any automatic action when it detects a resource fault. VCS places the resource into the ADMIN\_WAIT state and waits for administrative intervention.

1. To clear the resource in ADMIN\_WAIT state without faulting service group :

```
# hares -probe [resource] -sys node01
```

2. To clear the resource in ADMIN\_WAIT state by changing the status to OFFLINE|FAULTED :

```
# hagr -clearadminwait -fault [SG] -sys node01
```

How to upgrade VCS software on a running cluster:

Here's a procedure to upgrade VCS or shutdown VCS during hardware maintenance.

1. Open, freeze each Service Group, and close the VCS config.

```
haconf -makerw
hagr -freeze <Service Group> -persistent
haconf -dump makero
```

2. Shutdown VCS but keep services up.

```
hastop -all -force
```

3. Confirm VCS has shut down on each system.

```
gabconfig -a
```

4. Confirm GAB is not running on any disks.

```
gabdisk -l (use this if upgrading from VCS 1.1.x)
```

```
gabdiskhb -l
gabdiskx -l
```

If it is, remove it from the disks on each system.

```
gabdisk -d (use this if upgrading from VCS 1.1.x)
```

```
gabdiskhb -d
gabdiskx -d
```

5. Shutdown GAB and confirm it's down on each system.

```
gabconfig -U
gabconfig -a
```



6. Identify the GAB kernel module number and unload it from each system.

```
modinfo | grep gab  
modunload -i <GAB module number>
```

7. Shutdown LLT. On each system, type:

```
lltconfig -U
```

Enter "y" if any questions are asked.

8. Identify the LLT kernel module number and unload it from each system.

```
modinfo | grep llt  
modunload -i <LLT module number>
```

9. Rename VCS startup and stop scripts on each system.

```
cd /etc/rc2.d  
mv S70llt s70llt  
mv S92gab s92gab  
cd /etc/rc3.d  
mv S99vcs s99vcs  
cd /etc/rc0.d  
mv K10vcs k10vcs
```

10. Make a backup copy of /etc/VRTSvcs/conf/config/[main.cf](#).  
Make a backup copy of /etc/VRTSvcs/conf/config/[types.cf](#).

Starting with VCS 1.3.0, preonline and other trigger scripts must be in /opt/VRTSvcs/bin/triggers. Also, all preonline scripts in previous versions (such as VCS 1.1.2) must now be combined in one preonline script.

11. Remove old VCS packages.

```
pkgrm VRTScsga VRTSvcs VRTSgab VRTSllt VRTSperl VRTSvcs wz
```

If you are upgrading from 1.0.1 or 1.0.2, you must also remove the package VRTSsnmp, and any packages containing a .2 extension, such as VRTScsga.2, VRTSvcs.2, etc.

Also remove any agent packages such as VRTSvcsix (Informix), VRTSvcsnb (NetBackup), VRTSvcssor (Oracle), and VRTSvcssy (Sybase).

Install new VCS packages.

Restore your [main.cf](#) and [types.cf](#) files.

## 12. Start LLT, GAB and VCS.

```
cd /etc/rc2.d
mv s70llt S70llt
mv s92gab S92gab
cd /etc/rc3.d
mv s99vcs S99vcs
cd /etc/rc0.d
mv k10vcs K10vcs
```

```
/etc/rc2.d/S70llt start
/etc/rc2.d/S92gab
/etc/rc3.d/S99vcs start
```

## 13. Check on status of VCS.

```
hastatus
hastatus -sum
```

## 14. Unfreeze all Service Groups.

```
haconf -makerw
hagrp -unfreeze <Service Group> -persistent
haconf -dump -makero
```

How to increase the Filesystem size in a clustered environment with EMC powerpath installed.

- 1) Request lun from storage team. Request lun be assigned a specific VNX tag, example "db\_extended". [NOTE: I am assuming you are using EMC VNC for storage]
- 2) Run df -kh as a baseline.  
On both nodes of the cluster, run these commands:
- 3) Scan for lun's and have the os detect them:  
echo "- - -" > /sys/class/scsi\_host/host0/scan  
echo "- - -" > /sys/class/scsi\_host/host1/scan  
echo "- - -" > /sys/class/scsi\_host/host2/scan  
echo "- - -" > /sys/class/scsi\_host/host3/scan
- 4) Write label to disk  
Fdisk /dev/sd\* , option w.
- 5) Have powerpath detect them:  
powermt config
- 6) Apply and save powerpath configuration:  
powermt save
- 7) Confirm disk is seen by powerpath. Look for requested lun tag  
powermt display dev=all
- 8) Have Veritas pick up disk

`vxctl enable`

9) Verify Veritas sees new disk and is named disk properly. Disk name should match emcpower name. If not fix with `vxedit`.

`vxdisk -e list`

On only the node where the volume group is online, run these commands:

10) Add lun to the Veritas diskgroup:

`vxdiskadm, option 1`

11) Use the `vxresize` command to resize the volume

`vxresize -x -g mydg myvol <new size>G [newdiskname]`