

A decorative graphic consisting of several overlapping blue circles of varying sizes and shades, arranged in a diagonal line from the top right towards the bottom right. The circles are set against a white background with thin blue lines that intersect at the top left and bottom right corners, creating a triangular frame.

Step-by-Step Oracle11gR2 Two Node RAC running Solaris10 x86 with ISCSI(ZFS) on Vmware

This article (runbook) would help you in setting up Oracle 11gR2 RAC setup on your home x86 machine in a virtualized environment, and can help you to a certain extent in exploring new features of 11gr2 Real Application Clusters on Solaris.

Ragesh Moyan
3/13/2010

Sno	Table of Contents [TOC]	Page
1	Objective & Audience	3
2	Hardware/Software used	3
3	Setup Diagram	4
4	Abbreviations used in this article	5
5	IP Addresses for all Nodes	5
6	High Level execution steps	5
7	Create Utility virtual machine (vmutil)	6
2	Setup DNS server on Utility VM (vmutil)	17
3	Setup ZFS ISCSI target on Utility VM (vmutil)	20
4	Create Virtual machines vmsol1 & vmsol2	22
5	Setup Link aggregation for private Interconnect (fault tolerance).	27
6	Pre-requisites for 11gR2 Grid Infrastructure & Database	28
7	Configure Shared Storage (ZFS ISCSI Targets)	38
8	Install Oracle 11gR2 Grid Infrastructure	47
9	Install Oracle 11gR2 RDBMS Binaries	57
10	Create Oracle 11gR2 RAC Database	60
11	Few Commands & Checks for RAC Administration.	65
12	Create a Vmware Virtual Team for vmutil,vmsol1& vmsol2 [<i>optional</i>]	66
13	Issues & Troubleshooting & Learnings	68
14	References	68

Disclaimer: The below article should be used for educational purposes only

1. Objective: Step-by-Step Oracle 11gR2 two nodes Real Application Cluster running Solaris10 with ISCSI (ZFS) on Vmware Workstation.

Audience: DBA and RAC enthusiasts with understanding of basic Virtualization, Networking Solaris 10 & Oracle 11g Database Administration.

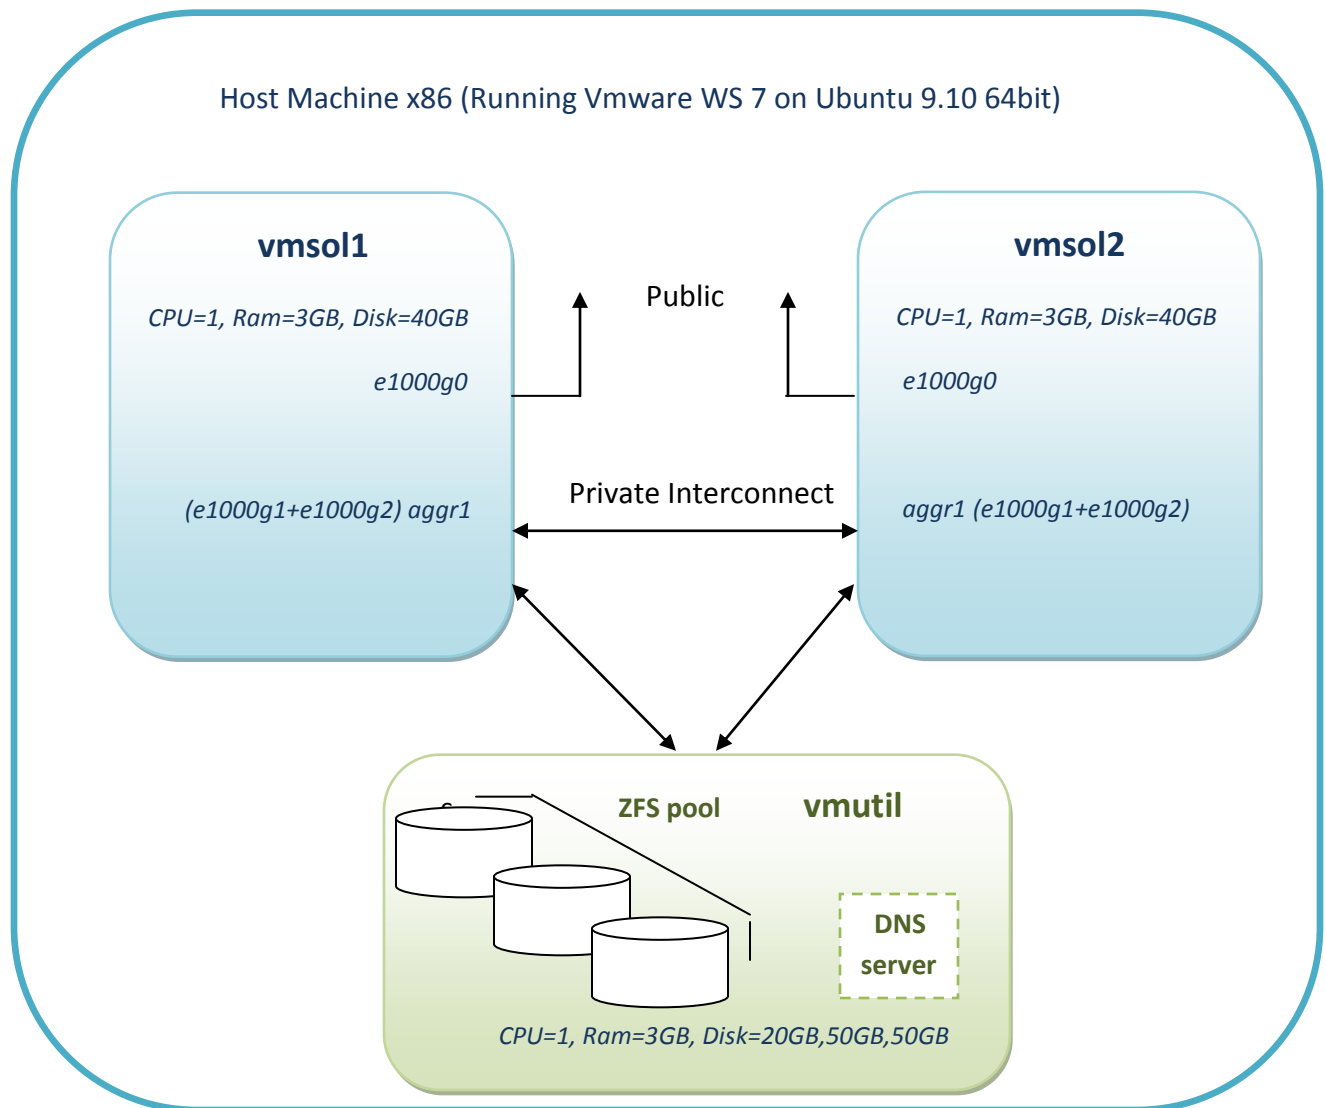
2. Hardware/Software used:

On Host machine (x86 architecture)	
CPU	2.4 GHZ Intel Core2 Duo
RAM	8 GB
Disk Space	160 GB
Operating System	Ubuntu 9.10 (64bit)
Virtualization software	Vmware Workstation 7.0

On Vmware Solaris Guest Machines (Virtual machines)		
Solaris 10 Update 8(64bit)	Free	http://www.sun.com/software/solaris/get.jsp
VMware Workstation 7 for Linux	Trial	http://downloads.vmware.com/d/info/desktop_downloads/vmware_workstation/7_0
Oracle 11gR2 Grid Infrastructure	Free	http://www.oracle.com/technology/software/products/database/index.html
Oracle 11gR2 Database Binaries	Free	http://www.oracle.com/technology/software/products/database/index.html
h2n script for DNS configuration	Free	ftp://ftp.hpl.hp.com/pub/h2n/h2n.tar.gz

Here you can also use Vmware server (free) or Virtual Box (free) for desktop virtualization. But I used Vmware WS 7 for its Snapshot management ability and easier navigation. For shared storage you can go for openfiler2.3 (free), if you have enough CPU and additional RAM (256 MB) for a fourth virtual machine in the current setup.

3. Setup Diagram



Overview of the setup:

The current setup uses 3 virtual machines (vmutil, vmsol1 & vmsol2) as shown above in the setup diagram all running Solaris10 64bit.

vmutil: is a Utility VM which hosts a DNS server (for SCAN IP's for 11gR2 Grid Infra) and also works as a iSCSI target (shared storage) created from ZFS (Zettabyte File System) Inbuilt iSCSI support, serving iSCSI shares for two RAC nodes (vmsol1 & vmsol2).

vmsol1 & vmsol2: These are 2 virtual machines which will host Oracle 11gR2 Grid & RAC database using shared storage from Utility VM. We would be using ASM as the storage option for the RAC database and also have OCR & voting disks on ASM.

4. Abbreviations used in this article:

VM : Virtual machine **bothnodes** : vmsol1 & vmsol2
ZFS : Zettabyte File System **allnodes** : vmutil, vmsol1 & vmsol2

5. IP Address for the Nodes:

IP Addr	Hostname	**FQDN	Description
192.168.1.60	vmsol1	vmsol1.mydom.com	Public IP addr (rac node1)
192.168.1.61	vmsol1-vip	vmsol1-vip.mydom.com	Virtual IP addr(rac node1)
172.168.1.60	vmsol1-priv	vmsol1-priv.mydom.com	Private Interconnect(rac node1)
192.168.1.70	vmsol2	vmsol2.mydom.com	Public IP addr (rac node2)
192.168.1.71	vmsol2-vip	vmsol2-vip.mydom.com	Virtual IP addr(rac node2)
172.168.1.70	vmsol2-priv	vmsol2-priv.mydom.com	Private Interconnect(rac node2)
192.168.1.80 192.168.1.81 192.168.1.82	vmcluster-scan	vmcluster-scan.mydom.com	3 Scan IP addr associated with the cluster (vmcluster) registered with DNS server.
192.168.1.96	vmutil	vmutil.mydom.com	DNS server & ISCSI Target.

****FQDN: Fully Qualified Domain Name**

6. High Level Execution Steps

High level steps involved in Setup	
Create Utility virtual machine (vmutil)	Click
Setup DNS server on Utility VM (vmutil)	Click
Setup ZFS ISCSI target on Utility VM (vmutil)	Click
Create Virtual machines vmsol1 & vmsol2	Click
Setup Link aggregation for private Interconnect (fault tolerance).	Click
Pre-requisites for 11gR2 Grid Infrastructure & Database	Click
Configure Shared Storage (ZFS ISCSI Targets)	Click
Install Oracle 11gR2 Grid Infrastructure	Click
Install Oracle 11gR2 RDBMS Binaries	Click
Create Oracle 11gR2 RAC Database	Click
Few Commands & Checks for RAC Administration.	Click
Create a VMware Virtual Team for vmutil,vmsol1& vmsol2 [optional]	Click
Issues & Trouble Shooting	Click

Have included screen shots.

Let's begin

7. Create Utility Virtual Machine (vmutil)

Specification of Utility server(vmutil)	
Memory(RAM)	800MB
Processors	1
Hard Disks(scsi)	3 disks(20GB,50GB,50GB)
Network Adaptor	1 (Bridged)
CD/DVD(IDE)	1
Display	AutoDetect(Default)


If you are already familiar with creation of virtual machine you can look at the screen shot at the end of of Step 7.1 and create it accordingly. For those new to VMware WS, below steps should get you going with vmutil VM.

"=>" arrow symbol indicates a **Screen** or a **Window** and **Text in ""double quotes** would indicate a value to be put in or the label text. Also unless specified, values would be default one.

7.1 Create the virtual Machine for vmutil

Goto VMware Workstation **Home** TAB on Host machine.

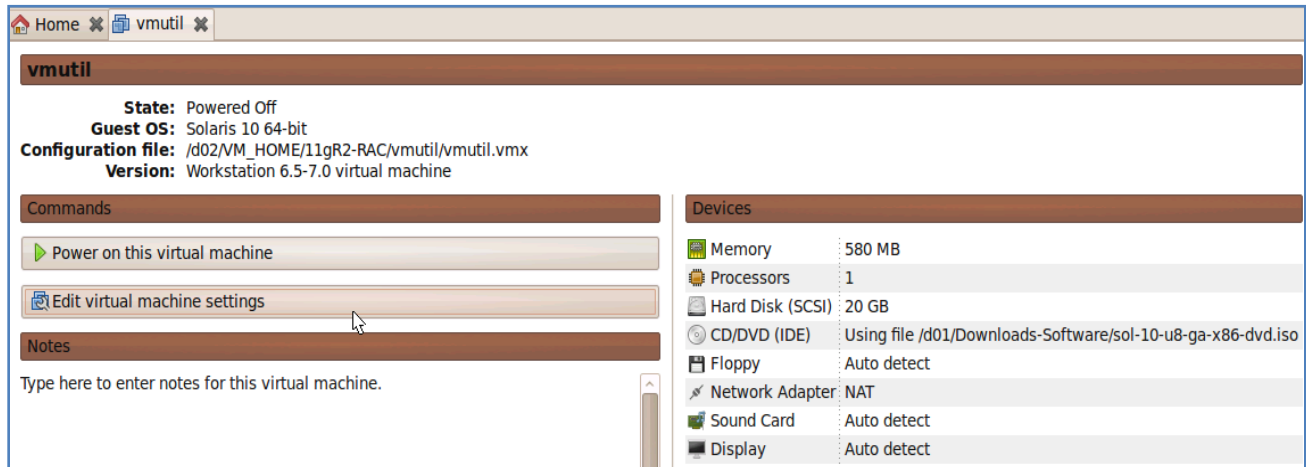
Click "**Create New Virtual Machine**" button



- => Choose radio button "**Typical**"
Click "**Next**"
- => Check radio "**Installer disc image file (ISO):**"
Use "**Browse**" button to include the Solaris10 ISO (here sol-10-u8-ga-x86-dvd.iso)
Click "**Next**"
- => Choose radio "**Sun Solaris**"
Drop down under "**Version**" choose "**Solaris 10 64 bit**"
Click "**Next**"
- => Fill in Virtual Machine details
Name: **vmutil**
Location: **<absolutePath>/vmutil** [use Browse button to choose location]
Click "**Next**"

- => Disk Size
 - From drop down for "**Maximum disk Size: 20 GB**"
 - Choose radio "**Split virtual disk into 2GB files**"
 - Click "**Next**"
- => Uncheck "**Automatically power on this virtual machine after creation**"
- Click "**Finish**"

Now you will get a New Tab **vmutil** with all selected in above steps will look like below.



Click on the "**Edit virtual machine settings**" button

- => Increase memory to "**800MB**"
 - Processors = 1 [default] (Here though Host is on dual core still choose 1 to avoid Performance degradation, if quad core choose 2)
- Choose "**Floppy**" and "**Sound card**" and use "**Remove**" button to remove them one by one.
- Choose "**Network Adaptor**" and choose the radio "**Bridged**" on the right.
- **We need to add two additional Hard Disks (SCSI) disks of size 50GB each.**
- Click on "**Add**" button on the bottom of the window
 - => Select "**Hard Disk**"
 - Click "**Next**"
 - => Choose radio "**Create a new virtual disk**"
 - Click "**Next**"
 - => Choose radio under Virtual Disk type as "**SCSI (recommended)**"
 - [Careful here] Leave the "**Mode unchecked**" [default]
 - [Reason is we would be taking snapshots for these disks]
 - Click "**Next**"
 - => Enter under "**Disk Size**"
 - Maximum disk size in (GB): **50**
 - Choose radio "**Split virtual disk into 2GB files**"
 - Click "**Next**"

- => Enter or use **Browse** button for below
Disk File name : <Path for vmutil dir on host>/vmutil-ZFS-Disk01-0.vmdk
Click "**Finish**"

***You have a new Hard Disks (SCSI) disk of 50 GB appear*

Add another Hard Disk(SCSI) of 50GB (**vmutil-ZFS-Disk02-0.vmdk**) using the above steps
Use "**Save**" Button to save the configuration.

Now you should be ideally back on the Main tab of **vmutil** and your configuration for vmutil should like below.

Devices	
Memory	800 MB
Processors	1
Hard Disk (SCSI)	20 GB
Hard Disk 2 (SCSI)	50 GB
Hard Disk 3 (SCSI)	50 GB
CD/DVD (IDE)	Using file /d01/Downloads-Software/sol-10-u8-ga-x86-dvd.iso
Network Adapter	Bridged
Display	Auto detect

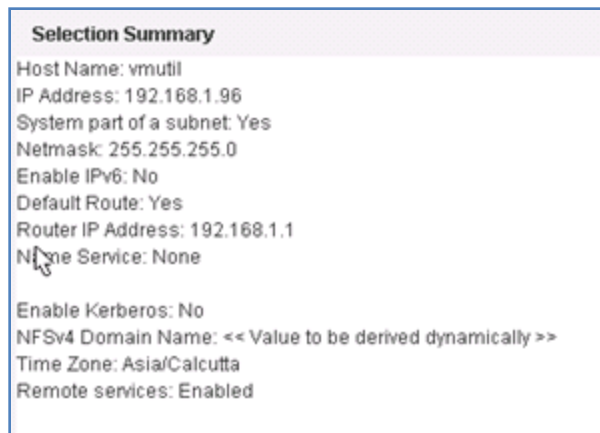
7.2 Install Guest OS - Solaris 10 (64 bit)

"=>" arrow symbol indicates a **Screen** or a **Window** and **Text** in **""double quotes** would indicate a value to be put in or the label text. Also unless specified, values would be default one.

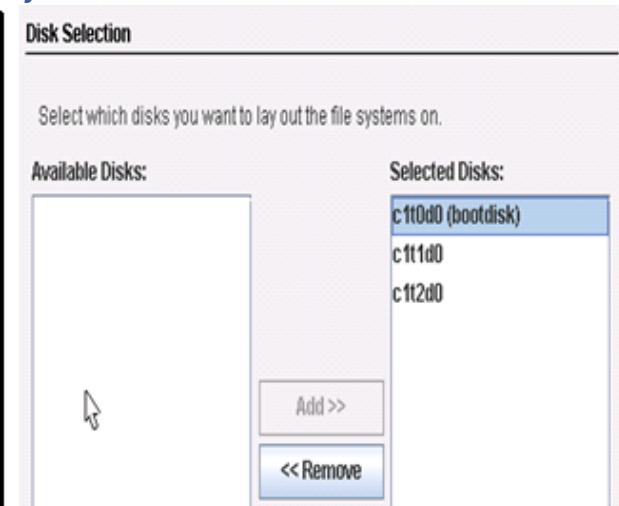
Click on "**Power on this virtual machine**" in **vmutil** tab and power on the vmutil VM.
The installation will start in a CUI mode and the first option you have to choose starts from below step

- => Choose "**1. Solaris Interactive (default)**"
- => "**Configure Keyboard Layout**" choose "**US-English**"(which is default) and
Press "**F2**" to continue
- => Next press "**Enter**" key to start interactive installation mode
- => If Solaris detects your display it will prompt a window to confirm press "**Enter**" key.
- => "**Select a Language**" enter "**0**" and press "**Enter**" key
- => You will get "**Welcome**" screen Click "**Next**" to continue
- => "**Network Connectivity**" choose radio "**Networked**"
Click "**Next**"
- => "**DHCP for e1000g0**" Choose radio "**No**"
Click "**Next**"
- => Enter "**Hostname for e1000g0**" in text box in lowercase as "**vmutil**"
Click "**Next**"

- => Enter "**IP Address for e1000g0**" in text box as "**192.168.1.96**"
Click "**Next**"
- => Enter "**Netmask for e1000g0**" in text box as "**255.255.255.0**"
Click "**Next**"
- => "**Enable IPv6 for e1000g0**" choose radio "**No**"
Click "**Next**"
- => "**Set the Default Route for e1000g0**" choose radio "**Specify one**"
Click "**Next**"
- => "**Set the Default Route for e1000g0**" enter under text box under label "**Route IP Address for e1000g0**" as "**192.168.1.1**"
Click "**Next**"
- => "**Enable Kerberos**" choose radio "**No**"
Click "**Next**"
- => "**Name Service**" choose radio "**None**"
Click "**Next**"
- => "**NFSv4 Domain Name**" choose the radio "**Use the NFSv4 domain derived by the system**"
Click "**Next**"
- => "**Time Zone**" choose radio "**Geographic Continent/Country/Region**"
Click "**Next**"
- => "**Continent and Country**" Choose the continent and country you want to set from the expandable tree menu. I choose "**Asia**" under it "**India**"
Click "**Next**"
- => "**Date and Time**" Vmware would sync the time with your current host machine time. So left it default.
Click "**Next**"
- => "**Root Password**" enter root password twice in the 2 text boxes.
Click "**Next**"
- => "**Enabling Remote Services**" choose radio "**Yes**"
Click "**Next**"
- => "**Confirm Information**" press button "**Confirm**"



- => **"Welcome Screen"**
Click **"Next"**
- => **"Installer Options"**
Choose both radio buttons as **"Yes"**
Click **"Next"**
=> **"Notice"**
Click **"OK"**
- => **"Specify Media"**
Choose radio under label **"Media"** as **"CD/DVD"**
Click **"Next"**
- => **"License"** check **"Accept"**
Click **"Next"**
- => **"Select Type of Install"**
Choose the radio **"Custom Install"**
Click **"Next"**
- => **"Select Software Localizations"** [I have not chosen any]
Click **"Next"**
- => **"Select System Locale"** [leave default as English (POSIX C)(C)]
Click **"Next"**
- => **"Additional Products"**
Choose **"None"**
Click **"Next"**
- => **"Select Solaris Software Group"**
For **"Entire Group Plus OEM"** choose radio **"Default Packages"**
Click **"Next"**
- => **"Disk Selection"**
Use the **"Add"** button to add Available Disks **"c1t1d0 & c1t2d0"** to Selected Disks
as shown below in 2 images
Click **"Next"**

Before**After**

=> "Select Disks for fdisk Partition Customization"

You will see all three disks checked.

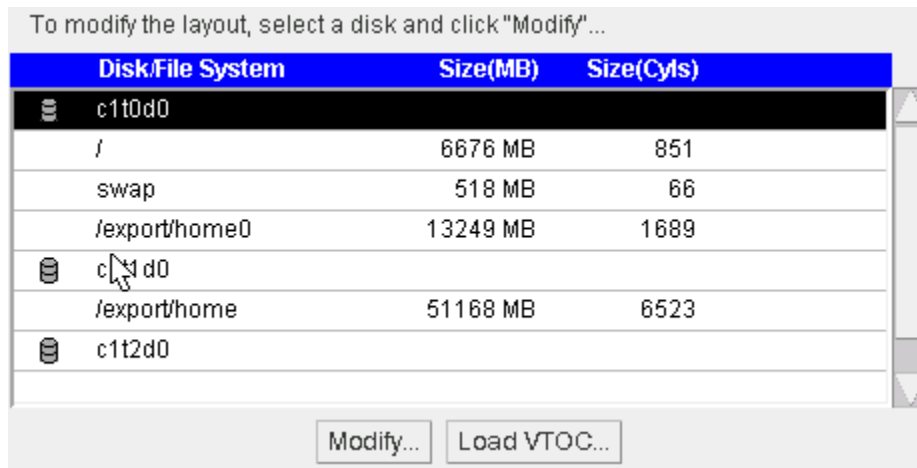
Click "Next"

=> "Customize fdisk Partitions-- Disk c1t0d0" [leave all default]

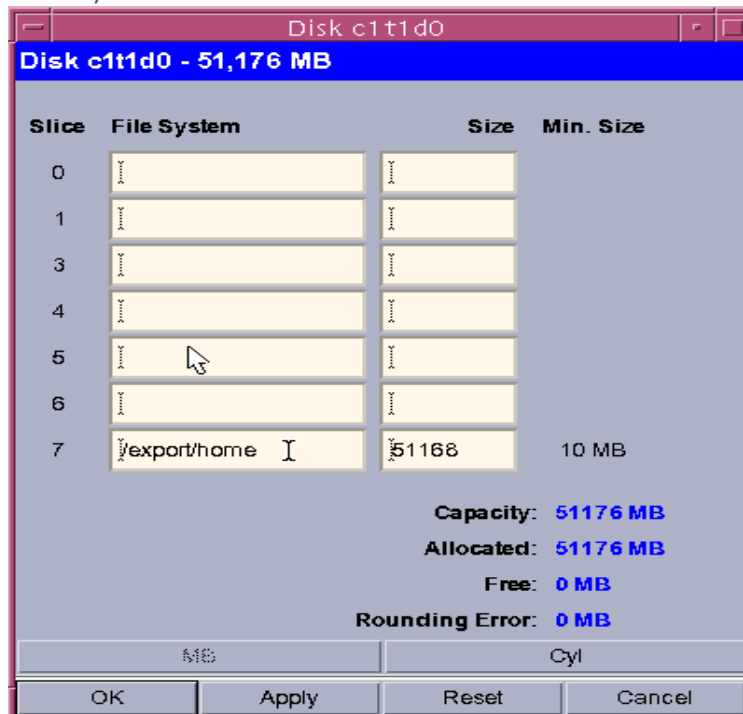
Click "Next"

** Leave Default for the remaining "Customize fdisk Partitions--" and Click "Next"

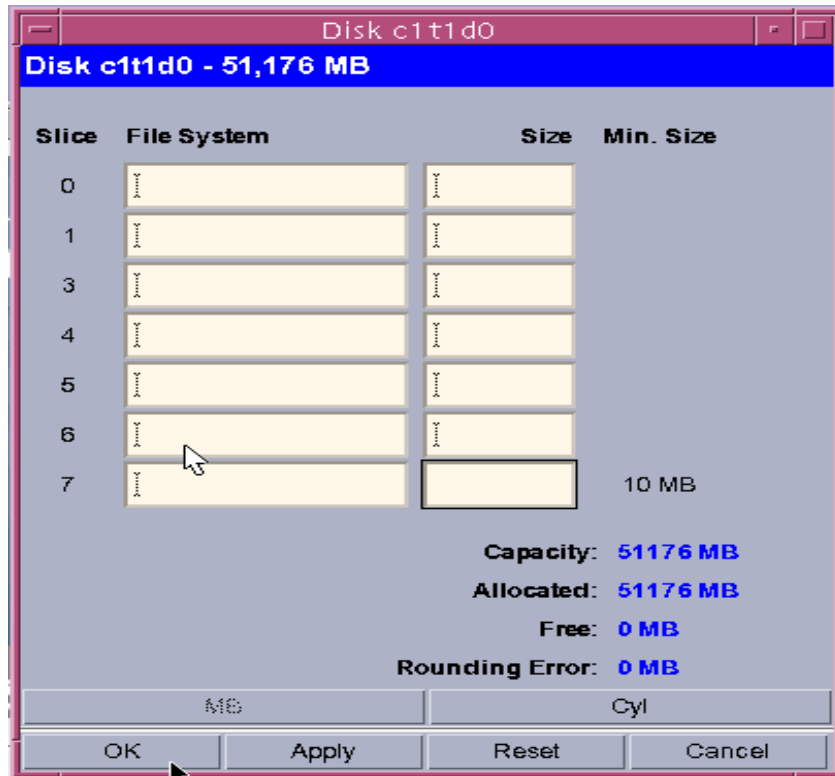
=> "Lay Out File Systems" [Will look like below]



Now Click "c1t1d0" disk and click "Modify" button (it will pop-up a smaller window).

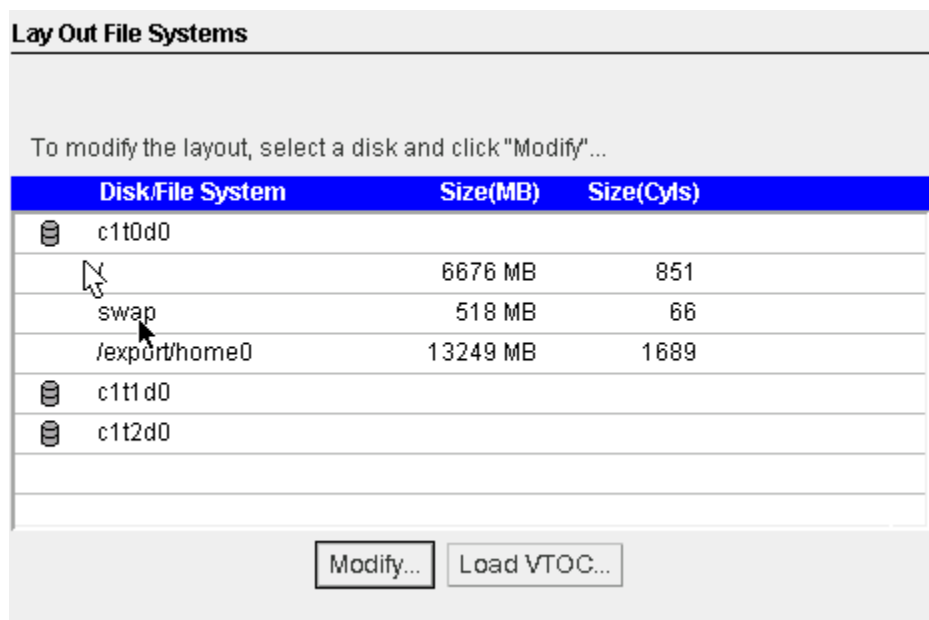


Delete the text "/export/home" and "51168" [Check image below]
Click "OK"



Similarly Choose Disk "**c1t2d0**",
Click "**Modify**" button and do the same steps as for "**c1t1d0**"

After editing both c1t1d0 & c1t2d0 file system layout should look like below.



Now Choose Disk "c1t0d0", Click "**Modify**" button.

Change the "/" filesystem from 6676 MB to 12000 MB

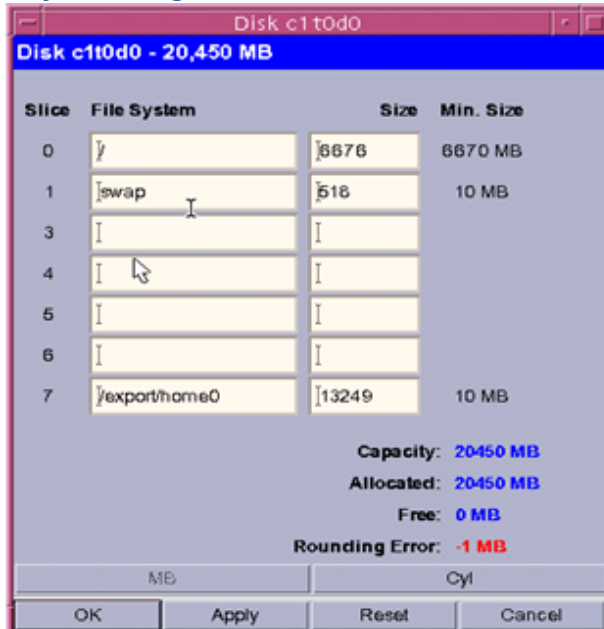
Change the "swap" from 518 MB to 4096 MB

Change /export/home0 to /export/home and size from 13249 MB to 4336 MB

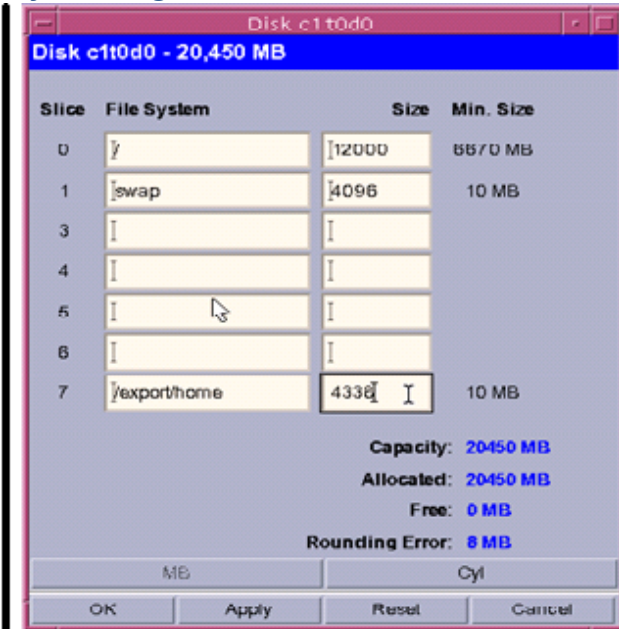
Click "OK"

See the image before and after changes

Before changes



After changes



The Final File System Layout should look like below before start of the install.

Disk/File System	Size(MB)	Size(Cyls)
c1t0d0		
/	12000 MB	1530
swap	4096 MB	523
/export/home	4338 MB	553
c1t1d0		
c1t2d0		

Click "**Next**"

=> "**Ready to Install**"

Click "**Install Now**"

This will kick off the Solaris Installation.

During installation, multiple windows will pop-up for confirmation. Click "**Continue**"

Once the installation completes you should have a "**Welcome to vmutil**" login screen.

7.3 Configure the Desktop on vmutil

- Add Terminal Icon and Set the Screen Resolution
- Install VMware Tools
- Configure the Shared folder (/hgfs)


At this point we already have installed Solaris 10 and we have "Welcome to vmutil" screen.

Add Terminal Icon and Set the Screen Resolution

=> Login as root user in "Welcome to vmutil" screen.

=> Choose "Java Desktop System, Release 3"

Click "OK"



Now Solaris does some start up checks and it would ask for registration.
I chose to Register Later.

Click on "Launch" button on left bottom of the desktop.

Navigation: Launch => Applications => Utilities => Terminal

Right click on Terminal and click on "Add this launcher to panel"



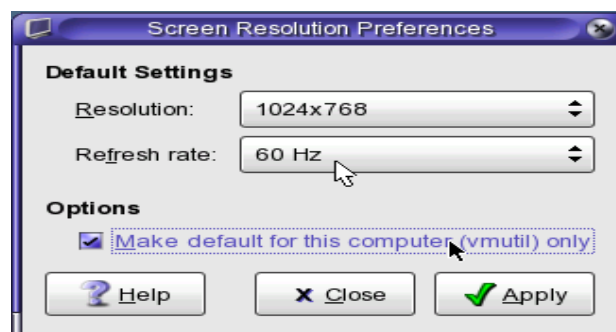
Now you should have the Terminal icon on the Launch panel.

Click on "Launch" button on left bottom of the desktop.

Navigation:


Launch => Preferences => Desktop Preferences => Display => Screen Resolution

Click on Screen Resolution and set the resolution which suits your monitor. I chose to Set Resolution: **1280x1024** and Refresh Rate: **60 Hz** and Click **Apply**.



Install Vmware Tools

Right click on the "vmutil" tab as show below and click on "Install VMware Tools..."



=> A small window pops-up "Question" Click "Install"

=> Now another window "File Browser: vmwaretools" opens up showing "vmware-solaris-tools.tar.gz"

- Copy "vmware-solaris-tools.tar.gz" to "/tmp" location and extract or Right click the ICON of "vmware-solaris-tools.tar.gz" in this window and "Open with" the file using application "File Roller".
- You should see another window with "vmware-tools-distrib" folder and use the "Extract" button on the menu bar to extract "vmware-solaris-tools.tar.gz" on the "Filesystem" (/tmp)
- Open a Terminal window (use the Icon which we created in earlier step on Launch panel)
Go to /tmp:

```
root# cd /tmp/vmware-tools-distrib
```

```
root# ./vmware-install.pl
```

 - ** Choose all default values for vmware-install.pl by pressing <ENTER> key.


Configure the Shared folder (/hgfs)

Right click on the "vmutil" tab and click on "Settings"

=> "Virtual Machine Settings" click on "Options" Tab

Click on "Shared Folders" [Disabled by Default] , on the right side of the screen choose radio "Always Enabled". Use "Add" button to add the Path (folder) from the Host machine (ideally used as Oracle Software Stage area).In this article I have named it as "shared". Click "Save"

** On vmutil Shared mount in host with under /hgfs/<name given to shared folder>



7.4 Set the /etc/hosts on vmutil

Login as "root" on "vmutil" and "vi /etc/hosts"

Current entry in /etc/hosts should like below for "vmutil" and add **FQDN** for vmutil.

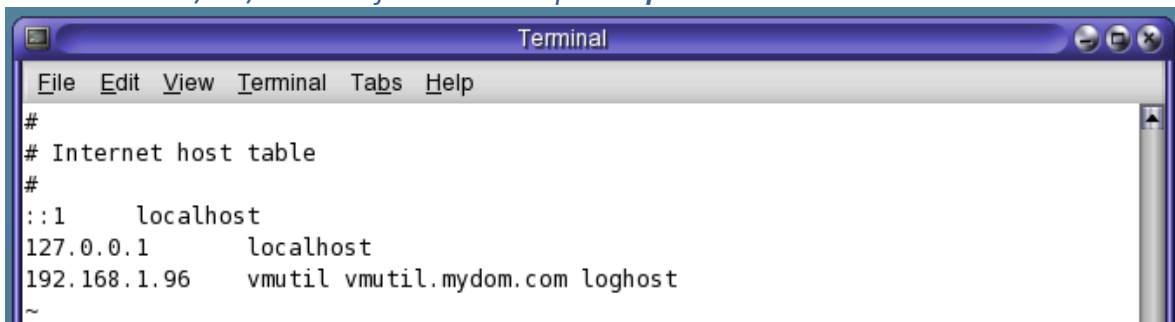
---Current /etc/hosts --

```
192.168.1.96      vmutil loghost
```

--After changes in /etc/hosts--

```
192.168.1.96      vmutil vmutil.mydom.com loghost
```

**** To save the /etc/hosts do force Save & quit :wq!**



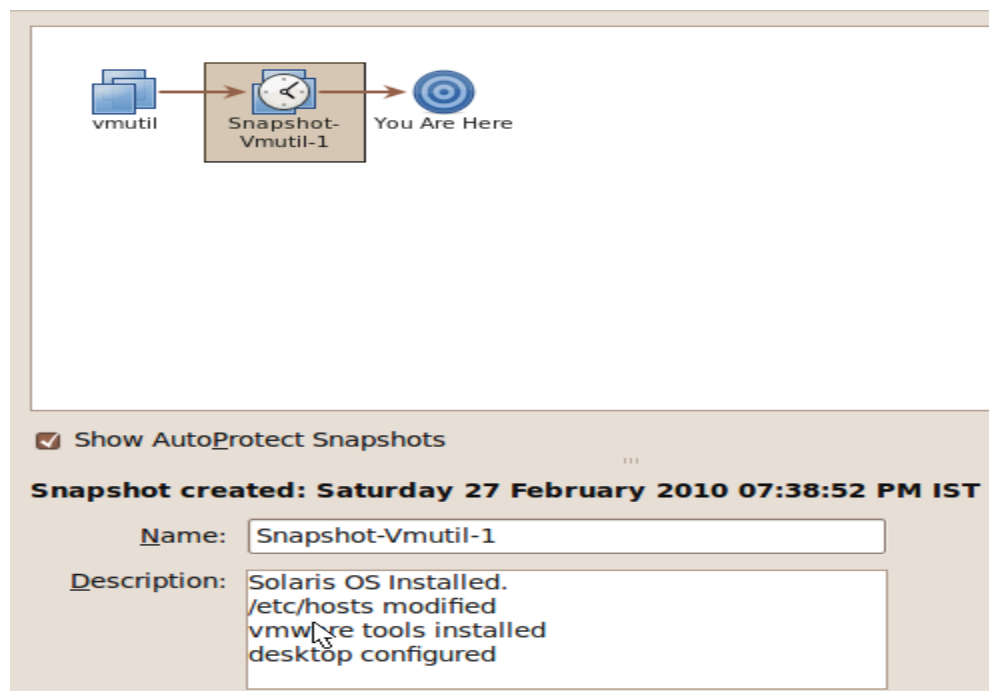
```
Terminal
File Edit View Terminal Tabs Help
#
# Internet host table
#
::1      localhost
127.0.0.1    localhost
192.168.1.96  vmutil vmutil.mydom.com loghost
~
```

7.5 Take a Snapshot of vmutil

Login as root on a Terminal and issue as "**poweroff**" command to Shutdown "vmutil"

After the shutdown, Right Click on the "vmutil" tab and Click on "**Snapshot Manager**".

It will open up a new window, Click "**Take Snapshot**" button and give appropriate **Snapshot Name** and **Description** which you can revert to if required.



8. Setup DNS server on Utility VM (vmutil)

[\[TOC\]](#)

Boot the vmutil server and login as root user

8.1 Add the entries of the IP address and the FQDN names on vmutil to be used in RAC setup.

Open a Terminal and add below entries in /etc/hosts of vmutil.

-----Cut&paste in /etc/hosts on vmutil -----

Public

192.168.1.60 vmsol1 vmsol1.mydom.com

192.168.1.70 vmsol2 vmsol2.mydom.com

VIP

192.168.1.61 vmsol1-vip vmsol1-vip.mydom.com

192.168.1.71 vmsol2-vip vmsol2-vip.mydom.com

Private Interconnect

172.168.1.60 vmsol1-priv vmsol1-priv.mydom.com

172.168.1.70 vmsol2-priv vmsol2-priv.mydom.com

Scan IP's

192.168.1.80 vmcluster-scan vmcluster-scan.mydom.com

192.168.1.81 vmcluster-scan vmcluster-scan.mydom.com

192.168.1.82 vmcluster-scan vmcluster-scan.mydom.com

-----Cut&paste in /etc/hosts on vmutil -----

8.2 Enable DNS service on vmutil

Login as root user in a Terminal on vmutil and follow below steps:

Enable the dns/server service using svcadm command line utility.

```
root# svcadm enable dns/server
```

```
root# svcs -l dns/server #=> Check the status of DNS server
```

```
--Check DNS service----
```

```
# svcs -l dns/server
```

```
fmri      svc:/network/dns/server:default
```

```
enabled   true
```

```
state     maintenance
```

```
next_state none
```

```
state_time Sat Feb 27 19:22:41 2010
```

```
logfile   /var/svc/log/network-dns-server:default.log
```

```
restarter svc:/system/svc/restarter:default
```

```
contract_id
```

```
dependency require_all/none svc:/system/filesystem/local (online)
```

```
dependency require_any/error svc:/network/loopback (online)
```

```
dependency optional_all/error svc:/milestone/network (online)
```

8.3 Configure h2n utility for DNS configuration files (on vmutil only)

- **Download h2n.tar.gz** in the shared folder (shared with host setup in Section 1.3) on vmutil from <ftp://ftp.hpl.hp.com/pub/h2n/h2n.tar.gz>

- **Copy & extract h2n.tar.gz** in /tmp on vmutil

```
root #cp /hgfs/shared/h2n.tar.gz /tmp;cd /tmp;gunzip h2n.tar.gz;tar -xvf h2n.tar
```

**Here shared is the name given to shared folder in Step 7.3

You should now have a directory by name h2n-2.56 in /tmp on vmutil.

- **Create /var/named directory & copy h2n** as root on vmutil

```
root # mkdir -p /var/named
```

```
root # cp /tmp/h2n-2.56/h2n /var/named
```

- **Create DNS configuration files on vmutil**

```
root# cd /var/named
```

```
root# ./h2n -d mydom.com -n 192.168.1 172.168.1 -u hostmaster.mydom.com
```

```
# id
uid=0(root) gid=0(root)
# ./h2n -d mydom.com -n 192.168.1 172.168.1 -u hostmaster.mydom.com
Initializing new database files...
Reading host file `/etc/hosts'...
Line 4: Skipping; incorrectly formatted data.
> ::1 localhost
Line 5: Skipping; IP not within range specified by -n/-a options.
> 127.0.0.1 localhost
Writing database files...
Generating boot and conf files...
Checking NS, MX, and other RRs for various improprieties...
Done.
# █
```

Above command will generate the below files in /var/named directory

```
--log--
```

```
boot.cacheonly db.127.0.0 db.192.168.1 h2n named.conf
```

```
conf.cacheonly db.172.168.1 db.mydom named.boot
```

```
--log--
```

- **Download "named.root" file from <ftp://ftp.rs.internic.net/domain> on vmutil.**
on vmutil:

```
root# cp named.root /var/named/db.cache
```

```
root# cp /var/named/named.conf /etc/
```

- **Edit /etc/nsswitch.files** of vmutil to add the entry dns to the host entry as
"hosts: files dns"

```
---log--
passwd: files
group: files
hosts: files dns
ipnodes: files
---log--
```

- **Create or Edit a file called /etc/defaultdomain** and add the single entry
On vmutil:
root# echo "mydom.com" > /etc/defaultdomain

- **Edit /etc/resolv.conf** on vmutil
root# echo "domain mydom.com" >> /etc/resolv.conf
root# echo "nameserver 192.168.1.96" >> /etc/resolv.conf

- **Restart the named service on vmutil.**
root# svcadm disable dns/server;svcadm enable dns/server

- **Verify if DNS is working (forward & reverse lookup)**

```
---log--
# nslookup vmsol1
Server: 192.168.1.96
Address: 192.168.1.96#53

Name: vmsol1.mydom.com
Address: 192.168.1.60

# nslookup vmsol1.mydom.com
Server: 192.168.1.96
Address: 192.168.1.96#53

Name: vmsol1.mydom.com
Address: 192.168.1.60

# nslookup 192.168.1.60
Server: 192.168.1.96
Address: 192.168.1.96#53
```

```
60.1.168.192.in-addr.arpa name = vmsol1.mydom.com.
#
---log----
```

Now we have a fully functional DNS server running on vmutil.

9. Setup ZFS ISCSI target on Utility VM (vmutil)

9.1 Identify the disks to be used on vmutil server for ZFS Iscsi target.

Use format command as root user to identify the disks.

```
root # format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
  0. c1t0d0 <DEFAULT cyl 2607 alt 2 hd 255 sec 63>
     /pci@0,0/pci15ad,1976@10/sd@0,0
  1. c1t1d0 <DEFAULT cyl 6524 alt 2 hd 255 sec 63> ==> used as disk for ZFS volume
     /pci@0,0/pci15ad,1976@10/sd@1,0
  2. c1t2d0 <DEFAULT cyl 6524 alt 2 hd 255 sec 63> ==> used as disk for ZFS volume
     /pci@0,0/pci15ad,1976@10/sd@2,0
Specify disk (enter its number): ^C
root #
```

9.2 Create a ZFS pool

We will create a mirrored zpool with 2 disks (c1t1d0 & c1t2d0)

```
root# zpool create iscsi_target mirror c1t1d0 c1t2d0
--log--
# zpool list
no pools available
# zpool create iscsi_target mirror c1t1d0 c1t2d0
# zpool list
NAME      SIZE  USED  AVAIL  CAP  HEALTH  ALTROOT
iscsi_target 49.8G  78K  49.7G   0%  ONLINE  -
#
---log--
```

9.3 Disable Iscsi initaitors on vmutil

Use "svcs -a|grep -i iscsi" to check the service of iscsi initaitor.

```
root# svcs -a|grep -i iscsi
disabled  20:22:59 svc:/system/iscsitgt:default
online    20:23:04 svc:/network/iscsi/initiator:default <=Need to be disabled
```

```
root# svcadm disable svc:/network/iscsi/initiator:default <=Disable iscsi initiator
```

```
root # svcs -a|grep -i iscsi
disabled  20:22:59 svc:/system/iscsitgt:default
disabled  23:30:55 svc:/network/iscsi/initiator:default <=Disabled iscsi initiator
```

9.4 Create Iscsi volumes in zpool on vmutil.

For this setup I chose **3G volume for CRS, 10G for DATA & 15G for FRA**.
These will be used as ASM disks on the rac nodes.

```
root# zfs create -V 3g iscsi_target/crs01
root# zfs create -V 10g iscsi_target/data01
root# zfs create -V 15g iscsi_target/fra01
```

9.5 Enable Iscsi on ZFS shares created above on vmutil

```
root# zfs set shareiscsi=on iscsi_target/crs01
root# zfs set shareiscsi=on iscsi_target/data01
root# zfs set shareiscsi=on iscsi_target/fra01
```

***Note do not use iscsitadm command if you set shareiscsi property, else you will end up with duplicate target information for same device.*

9.6 How to list & identify Iscsi targets

Use the command as root on vmutil "**iscsitadm list target -v**"

--log--

```
# iscsitadm list target -v
```

```
Target: iscsi_target/crs01
```

```
iSCSI Name: iqn.1986-03.com.sun:02:0d7b372c-c3b2-6041-bd24-b2cea2a34761
```

```
Alias: iscsi_target/crs01 <== CRS volume for ASM
```

```
Connections: 0
```

```
ACL list:
```

```
TPGT list:
```

```
LUN information:
```

```
LUN: 0
```

```
GUID: 0
```

```
VID: SUN
```

```
PID: SOLARIS
```

```
Type: disk
```

```
Size: 3.0G
```

```
Backing store: /dev/zvol/rdisk/iscsi_target/crs01
```

```
Status: online
```

```
Target: iscsi_target/data01
```

```
iSCSI Name: iqn.1986-03.com.sun:02:b5b07821-899e-4eb3-b8fd-c540f4ab05ff
```

```
Alias: iscsi_target/data01 <== Data Volume for ASM
```

```
Connections: 0
```

```
ACL list:
```

```
TPGT list:
```

```
LUN information:
```

```
LUN: 0
```

```
GUID: 0
```

```

VID: SUN
PID: SOLARIS
Type: disk
Size: 10G
Backing store: /dev/zvol/rdisk/iscsi_target/data01
Status: online
Target: iscsi_target/fra01
iSCSI Name: iqn.1986-03.com.sun:02:5b272782-6438-c793-a575-cbf42e8b9503
Alias: iscsi_target/fra01 <== FRA volume for ASM
Connections: 0
ACL list:
TPGT list:
LUN information:
  LUN: 0
  GUID: 0
  VID: SUN
  PID: SOLARIS
  Type: disk
  Size: 15G
  Backing store: /dev/zvol/rdisk/iscsi_target/fra01
  Status: online
#
---log---
```

9.7 Take a Snapshot of vmutil

Shutdown the vmutil and follow Steps as mentioned in Step 7.5 to take a snapshot and give a new name to the Snapshot and meaningful description.

10. Create Virtual Machines vmsol1 & vmsol2

Boot vmutil server before you start Step 10.1


10.1 Create virtual Machines (vmsol1 & vmsol2) with Specifications below

Specification of Utility server(vmsol1/vmsol2)	
Memory(RAM)	3000MB
Processors	1
Hard Disks(scsi)	1 disk(40GB)
Network Adaptor	3 (1 Bridged + 2 Host Only)
CD/DVD(IDE)	1
Display	AutoDetect(Default)

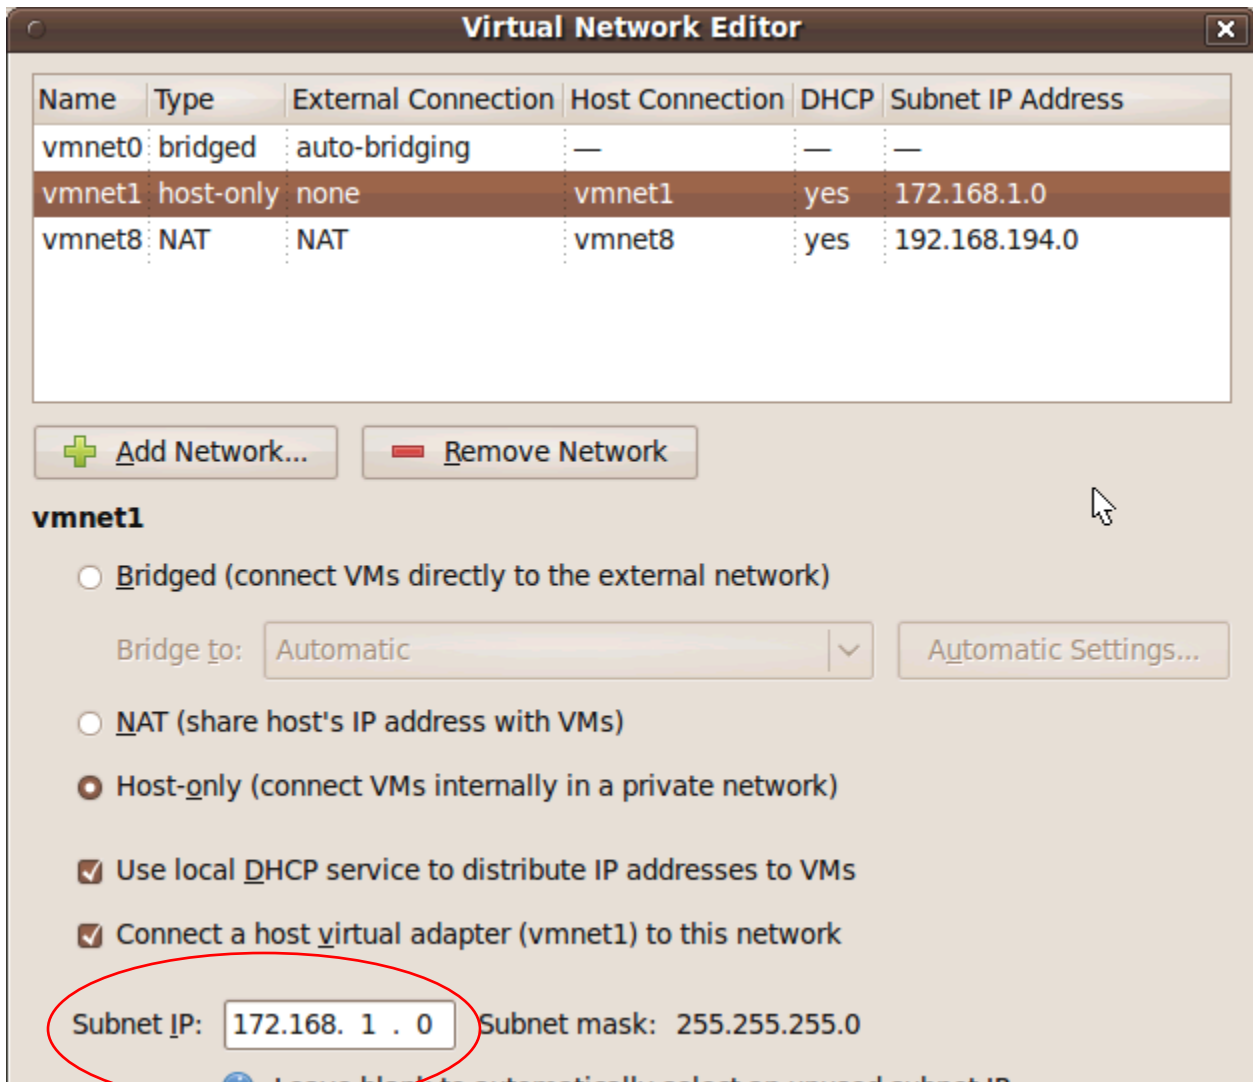
Use **“Edit Virtual Machine Settings”** => **“Add”** button to add additional 2 Network Adaptors as **Host Only**. Before this, let’s set the subnet of **Host only** to **172.168.1.0** for private network.

Set subnet of Host only to 172.168.1.0

Go to VMware WorkStation Home Tab Menu: **Edit =>Virtual Network Editor** (click it)



Here we use **vmnet0** (for bridged interfaces) and **vmnet1** (for Host Only interfaces)
As seen in image below if required edit the text field "**Subnet IP: 172.168.1.0**" and "**Save**".



vmсол1 & vmсол2 configuration should look like in below image.

Devices	
Memory	3000 MB
Processors	1
Hard Disk (SCSI)	40 GB
CD/DVD (IDE)	Using file /d01/Downloads-Software/sol-10-u8-ga-x86-dvd.iso
Network Adapter	Bridged
Network Adapter 2	Host-only
Network Adapter 3	Host-only
Display	Auto detect

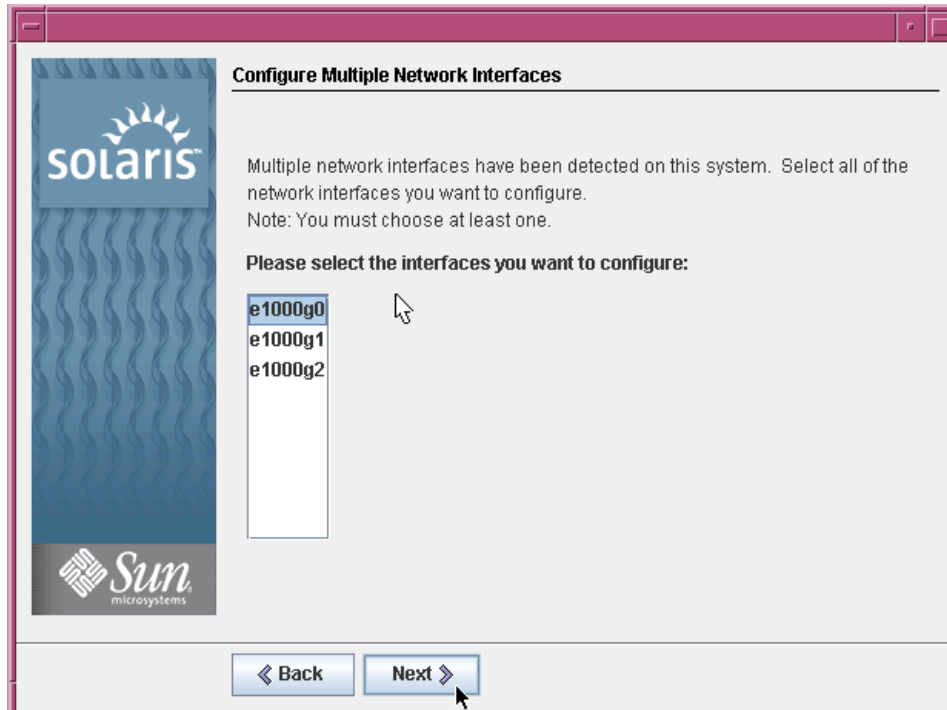
10.2 Install Solaris 10 in both vmсол1 & vmсол2

Follow the same Steps for installation on vmutil except few mentioned below,

=> "Configuring Multiple Network Interfaces" window

Please select the Interfaces you want to configure:

Choose only "e1000g0" and **do not select the remaining two.**



=> **Give IP address for**

For vmsol1: e1000g0 as **192.168.1.60**

For vmsol2: e1000g0 as **192.168.1.70**

Default Route: **192.168.1.1**

Broadcast : **255.255.255.0**

=> **"Name Service"**

Choose radio **"DNS"**

=> **"Domain Name"** give as mydom.com

=> **"DNS Server Address"**

For required Text box give IP: **192.168.1.96**

=> **"DNS Search List"**

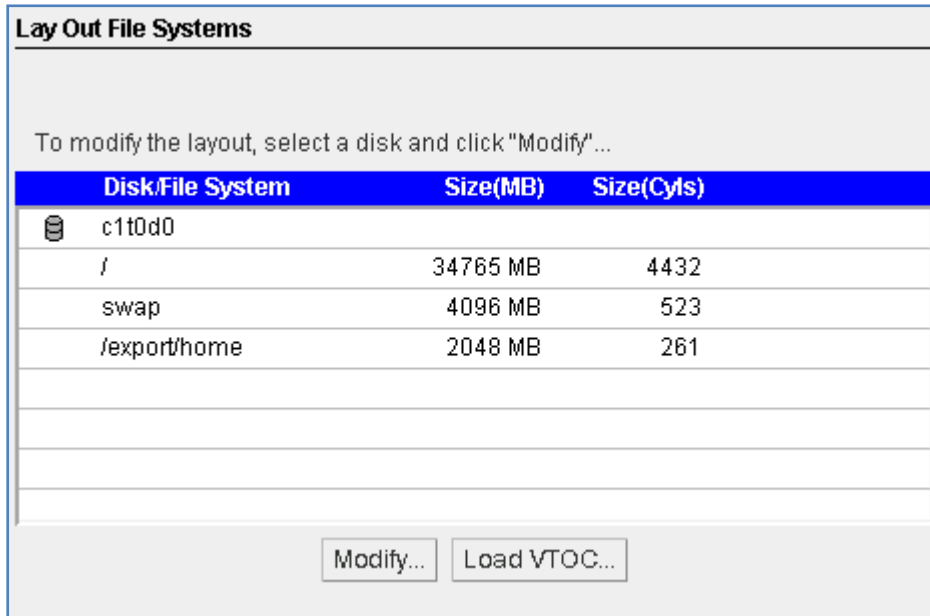
For **Label 1** give value in Textbox as **mydom.com**

Below are the Configurations for vmsol1 & vmsol2.

Selection Summary	Selection Summary
Host Name: vmsol1	Host Name: vmsol2
IP Address: 192.168.1.60	IP Address: 192.168.1.70
System part of a subnet: Yes	System part of a subnet: Yes
Netmask: 255.255.255.0	Netmask: 255.255.255.0
Enable IPv6: No	Enable IPv6: No
Default Route: Yes	Default Route: Yes
Router IP Address: 192.168.1.1	Router IP Address: 192.168.1.1
Name Service: DNS	Name Service: DNS
Domain Name: mydom.com	Domain Name: mydom.com
Server's IP Address:	Server's IP Address:
192.168.1.96	192.168.1.96
Search Domain:	Search Domain:
mydom.com	mydom.com
Enable Kerberos: No	Enable Kerberos: No
NFSv4 Domain Name: << Value to be derived dynamically >>	NFSv4 Domain Name: << Value to be derived dynamically >>
Time Zone: Asia/Calcutta	Time Zone: Asia/Calcutta
Remote services: Enabled	Remote services: Enabled

=> "Lay out File systems"

For vmsol1 and vmsol2



We have Solaris 10 64bit Guest OS on both vmsol1 & vmsol2.

10.3 Install Vmware tools and configure (Resolution, Terminal & Shared Folders) the desktop

Follow steps in **section 7.3** on **both vmsol1 & vmsol2**.

10.4 Setup /etc/hosts for vmsol1 & vmsol2

Login as "root" on "vmsol1" and "vi /etc/hosts"

Current entry in /etc/hosts should like below for "vmsol1" and add

---Current--

```
192.168.1.60      vmsol1 loghost
```

---Current--

--Add below changes--

```
192.168.1.60      vmsol1 vmsol1.mydom.com loghost
```

--Add below changes--

Login as "root" on "vmsol2" and "vi /etc/hosts"

Current entry in /etc/hosts should like below for "vmsol2" and add

---Current--

```
192.168.1.70      vmsol2 loghost
```

---Current--

--Add below changes--

```
192.168.1.70      vmsol2 vmsol2.mydom.com loghost
```

--Add below changes--

10.5 Verify DNS Client on vmsol1 & vmsol2

Check we are able to do a nslookup from vmsol1 & vmsol2
As **root** (Here both nodes points to vmsol1 & vmsol2)

```
bothnodes# nslookup vmsol1
bothnodes# nslookup vmsol2
bothnodes# nslookup vmsol1-priv
bothnodes# nslookup vmsol2-priv
bothnodes# nslookup vmsol1-vip
bothnodes# nslookup vmsol2-vip
bothnodes# nslookup vmcluster-scan
```

10.6 Take a Snapshot of vmsol1 & vmsol2

Shutdown the vmsol1 & vmsol2 and follow Steps as mentioned in Step 7.5 to take a Snapshot and give each snap a new name and a meaningful description.

11. Setup Link Aggregation for Private Interconnect (*fault tolerance*)

Start vmsol1 & vmsol2 (**vmutil is already running**)

Determine interfaces installed on the system

```
root# dladm show-link
```

```
# dladm show-link
e1000g0      type: non-vlan  mtu: 1500      device: e1000g0
e1000g1      type: non-vlan  mtu: 1500      device: e1000g1
e1000g2      type: non-vlan  mtu: 1500      device: e1000g2
# # ifconfig -a
```

Determine which interfaces are already plumbed

```
root# ifconfig -a
```

```
# # ifconfig -a
lo0: flags=2001000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4,VIRTUAL> mtu 8232 index 1
    inet 127.0.0.1 netmask ffffffff
e1000g0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
    inet 192.168.1.70 netmask fffffff0 broadcast 192.168.1.255
    ether 0:c:29:62:a1:7b
```

Create a Link Aggregation for 2 interfaces (e1000g1 & e1000g2)

- Confirm if these are not plumbed (*ifconfig -a* shouldn't show up e1000g1/e1000g2)
- Create link aggregation


```
root# dladm create-aggr -d e1000g1 -d e1000g2 1
```

*This will create link aggregation by name **aggr1***
- Give an IP for new aggregation


```
root# ifconfig aggr1 plumb 172.168.1.60 up [on vmsol1 ]
root# ifconfig aggr1 plumb 172.168.1.70 up [on vmsol2 ]
```

Check status of aggregation created

```
root# dladm show-aggr
```

```
# # dladm show-aggr
key: 1 (0x0001) policy: L4      address: 0:c:29:62:a1:85 (auto)
      device      address      speed      duplex  link  state
      e1000g1     0:c:29:62:a1:85  1000  Mbps   full  up    attached
      e1000g2     0:c:29:62:a1:8f  1000  Mbps   full  up    attached
```

Make the aggregation persistent across reboots

```
root# echo 172.168.1.60 > /etc/hostname.aggr1 [on vmsol1]
```

```
root# echo 172.168.1.70 > /etc/hostname.aggr1 [on vmsol2]
```

```
# ifconfig -a
lo0: flags=2001000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4,VIRTUAL> mtu 8232 index 1
      inet 127.0.0.1 netmask ff000000
e1000g0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
      inet 192.168.1.70 netmask fffffff0 broadcast 192.168.1.255
      ether 0:c:29:62:a1:7b
aggr1: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 3
      inet 172.168.1.70 netmask ffff0000 broadcast 172.168.255.255
      ether 0:c:29:62:a1:85
```

12. Pre-requisites for 11gR2 Grid Infrastructure & Database

12.1 Check/Install required OS version, patches & packages for 11gR2 Grid & Database

- **Version of Operating system** should be **Solaris 10 U6 (5.10-2008.10)** or later.

You can check using the command "**uname -r**" for current article **Solaris10 U8** is used.

- **Check and install required OS packages and Patches.**

Required OS Packages & Patches:

```
SUNWarc
SUNWbtool
SUNWcsl
SUNWhea
SUNWlibC
SUNWlibm
SUNWlibms
SUNWsprt
SUNWtoo
SUNWi1of (ISO8859-1)
SUNWi1cs (ISO8859-15)
SUNWi15cs
SUNWxfnt
119961-05 or later
119964-14 or later
120754-06 or later
139556-08 or later
139575-03 or later
137104-02 or later
```

Check if the above packages are installed as root user

```
bothnodes # pkginfo -i SUNWarc SUNWbtool SUNWhea SUNWlibC SUNWlibm \
  SUNWlibms SUNWsprout SUNWtoo SUNWi1of SUNWi1cs SUNWi15cs \
  SUNWxwfont SUNWcsl
```

```
bothnodes# pkginfo -i SUNWarc SUNWbtool SUNWhea SUNWlibC SUNWlibm
SUNWlibms SUNWsprout SUNWtoo SUNWi1of SUNWi1cs SUNWi15cs SUNWxwfont
SUNWcsl
```

```
system SUNWarc Lint Libraries (usr)
system SUNWbtool CCS tools bundled with SunOS
system SUNWcsl Core Solaris, (Shared Libs)
system SUNWhea SunOS Header Files
system SUNWi1of ISO-8859-1 (Latin-1) Optional Fonts
system SUNWlibC Sun Workshop Compilers Bundled libC
system SUNWlibm Math & Microtasking Library Headers & Lint Files (Usr)
system SUNWlibms Math & Microtasking Libraries (Usr)
system SUNWsprout Solaris Bundled tools
system SUNWtoo Programming Tools
system SUNWxwfont X Window System platform required fonts
ERROR: information for "SUNWi1cs" was not found
ERROR: information for "SUNWi15cs" was not found
```

Here above 2 packages "SUNWi1cs" & "SUNWi15cs" are not installed.

Mount the Solaris 10 ISO on bothnodes using "Settings" ->"CD/DVD(IDE)" for a each vm.

Install the missing packages as root on each vm.

As root on vmsol1/2:

```
bothnodes# pkgadd -d /cdrom/sol_10_1009_x86/Solaris_10/Product SUNWi1cs SUNWi15cs
```

Check & Install (if missing) Patches

To check patches installed use the command as root on each of the node.

```
bothnodes# /usr/sbin/patchadd -p | grep patch_number
```

eg: bothnodes# /usr/sbin/patchadd -p | grep 119964

***In our case we have all the required patches installed with the OS install*

12.2 Disable NTP service

For this article, **CTSS is used as time synchronisation** rather than NTP among the RAC nodes.

To **disable the NTP** service, run the following command as the root user

```
bothnodes# /usr/sbin/svcadm disable ntp
```

Check if any ntp services are running in **legacy_run mode**

```
bothnodes# svcs "*ntp*"
STATE      STIME    FMRI
disabled   8:23:09  svc:/network/ntp:default
```

disabled 8:23:20 svc:/network/ntp4:default

***No ntp services in legacy_run mode.*

12.3 Network settings

Add the IP addresses and FQDN of all hosts (vmsol1/vmsol2) participating in the RAC

On vmsol1:

```
-----Copy & Paste in /etc/hosts-----
# Public [e1000g0]
192.168.1.70      vmsol2 vmsol2.mydom.com

# VIP [e1000g0:1]
192.168.1.61     vmsol1-vip  vmsol1-vip.mydom.com
192.168.1.71     vmsol2-vip  vmsol2-vip.mydom.com

# Private Interconnect [aggr1]
172.168.1.60     vmsol1-priv vmsol1-priv.mydom.com
172.168.1.70     vmsol2-priv vmsol2-priv.mydom.com
-----
```

On vmsol2:

```
-----Copy & Paste in /etc/hosts-----
# Public [e1000g0]
192.168.1.60     vmsol1 vmsol1.mydom.com

# VIP [e1000g0:1]
192.168.1.61     vmsol1-vip  vmsol1-vip.mydom.com
192.168.1.71     vmsol2-vip  vmsol2-vip.mydom.com

# Private Interconnect [aggr1]
172.168.1.60     vmsol1-priv vmsol1-priv.mydom.com
172.168.1.70     vmsol2-priv vmsol2-priv.mydom.com
-----
```

Set the UDP parameters on both vmsol1/vmsol2

To check current settings for `udp_rcv_hiwat` and `udp_xmit_hiwat`, enter the following Commands as root user on both nodes.

```
# ndd /dev/udp udp_xmit_hiwat
# ndd /dev/udp udp_rcv_hiwat
```

On Solaris 10, to set the values of these parameters to **65536 bytes in current memory**,

Enter the following commands for setting until next boot up:

```
# ndd -set /dev/udp udp_xmit_hiwat 65536
# ndd -set /dev/udp udp_rcv_hiwat 65536
```

Make values persistent across reboots.

As root user on **both nodes**.

```
# cat>/etc/rc2.d/S99ndd<<EOF
ndd -set /dev/udp udp_xmit_hiwat 65536
ndd -set /dev/udp udp_rcv_hiwat 65536
EOF
```

12.4 Create Users, Groups and Directories for Grid & RDBMS home.

Run below on **bothnodes** as **root** user:

----Cut&Paste as root on vmsol1 & vmsol2-----

```
groupadd -g 1000 oinstall
groupadd -g 1200 asmadmin
groupadd -g 1201 asmdba
groupadd -g 1202 asmoper
groupadd -g 1300 dba
groupadd -g 1301 oper
```

```
useradd -u 1100 -g oinstall -G asmadmin,asmdba,asmoper -d /export/home/grid \
-m -s /usr/bin/bash -c "Grid Infrastructure Owner" grid
passwd grid
```

```
useradd -u 1101 -g oinstall -G dba,oper,asmdba -d /export/home/oracle -m -s \
/usr/bin/bash -c "Oracle Software Owner" oracle
passwd oracle
```

```
mkdir -p /u01/app/grid
mkdir -p /u01/app/11.2.0/grid
chown -R grid:oinstall /u01
mkdir -p /u01/app/oracle
chown oracle:oinstall /u01/app/oracle
chmod -R 775 /u01
```

----Cut&Paste as root on vmsol1 & vmsol2-----

Check for nobody user (id nobody), it will exists as part of the solaris install

12.5 Setup .profiles for Users (oracle & grid)

Login to vmsol1/vmsol2 as root and switch to **grid** user as below and add the entires.

***Here comment the ORACLE_SID as per node.*

eg:- Comment ORACLE_SID=+ASM2 on vmsol1

```

root# su - grid
-----.profile-----
ORACLE_SID=+ASM1; export ORACLE_SID
ORACLE_SID=+ASM2; export ORACLE_SID

JAVA_HOME=/usr/local/java; export JAVA_HOME
ORACLE_BASE=/u01/app/grid; export ORACLE_BASE
ORACLE_HOME=/u01/app/11.2.0/grid; export ORACLE_HOME
ORACLE_TERM=xterm; export ORACLE_TERM
NLS_DATE_FORMAT="DD-MON-YYYY HH24:MI:SS"; export NLS_DATE_FORMAT
TNS_ADMIN=$ORACLE_HOME/network/admin; export TNS_ADMIN
ORA_NLS11=$ORACLE_HOME/nls/data; export ORA_NLS11
PATH=.:${JAVA_HOME}/bin:${PATH}:${HOME}/bin:$ORACLE_HOME/bin
PATH=${PATH}:/usr/bin:/bin:/usr/bin/X11:/usr/local/bin
PATH=${PATH}:/u01/app/common/oracle/bin
export PATH
LD_LIBRARY_PATH=$ORACLE_HOME/lib
LD_LIBRARY_PATH=${LD_LIBRARY_PATH}:${ORACLE_HOME/oracm/lib}
LD_LIBRARY_PATH=${LD_LIBRARY_PATH}:/lib:/usr/lib:/usr/local/lib
export LD_LIBRARY_PATH
CLASSPATH=$ORACLE_HOME/JRE
CLASSPATH=${CLASSPATH}:${ORACLE_HOME/jlib}
CLASSPATH=${CLASSPATH}:${ORACLE_HOME/rdbms/jlib}
CLASSPATH=${CLASSPATH}:${ORACLE_HOME/network/jlib}
export CLASSPATH
THREADS_FLAG=native; export THREADS_FLAG
export TEMP=/tmp
export TMPDIR=/tmp
umask 022
set -o vi
-----.profile-----

```

Login to vmsol1/vmsol2 as root and switch to **oracle** user as below and add the entires.

***Here comment the ORACLE_SID as per node.*

eg:- Comment ORACLE_SID=+racdb2 on vmsol1

```
root# su - oracle
```



```

----.profile----
# Set the below ORACLE_SID as per node
ORACLE_SID=racdb1; export ORACLE_SID
ORACLE_SID=racdb2; export ORACLE_SID

ORACLE_UNQNAME=racdb; export ORACLE_UNQNAME
JAVA_HOME=/usr/local/java; export JAVA_HOME
ORACLE_BASE=/u01/app/oracle; export ORACLE_BASE
ORACLE_HOME=$ORACLE_BASE/product/11.2.0/dbhome_1; export ORACLE_HOME
ORACLE_PATH=/u01/app/common/oracle/sql; export ORACLE_PATH
ORACLE_TERM=xterm; export ORACLE_TERM
NLS_DATE_FORMAT="DD-MON-YYYY HH24:MI:SS"; export NLS_DATE_FORMAT
TNS_ADMIN=$ORACLE_HOME/network/admin; export TNS_ADMIN
ORA_NLS11=$ORACLE_HOME/nls/data; export ORA_NLS11
PATH=.:{JAVA_HOME}/bin:{PATH}:$HOME/bin:$ORACLE_HOME/bin
PATH=${PATH}:/usr/bin:/bin:/usr/bin/X11:/usr/local/bin
PATH=${PATH}:/u01/app/common/oracle/bin
export PATH
LD_LIBRARY_PATH=$ORACLE_HOME/lib
LD_LIBRARY_PATH=${LD_LIBRARY_PATH}:$ORACLE_HOME/oracm/lib
LD_LIBRARY_PATH=${LD_LIBRARY_PATH}:/lib:/usr/lib:/usr/local/lib
export LD_LIBRARY_PATH
CLASSPATH=$ORACLE_HOME/JRE
CLASSPATH=${CLASSPATH}:$ORACLE_HOME/jlib
CLASSPATH=${CLASSPATH}:$ORACLE_HOME/rdbms/jlib
CLASSPATH=${CLASSPATH}:$ORACLE_HOME/network/jlib
export CLASSPATH
THREADS_FLAG=native; export THREADS_FLAG
export TEMP=/tmp
export TMPDIR=/tmp
umask 022
set -o vi
----.profile----

```

12.6 Setup Kernel parameters

***Below steps needs to be done on both vmsol1 & vmsol2*

Setup /etc/system

Oracle Clusterware does not set project information when starting processes, some /etc/system processes that are deprecated but not removed must still be set for Oracle Clusterware. Oracle recommends that you set both resource control and /etc/system parameters. Operating system parameters not replaced by resource controls continue to affect performance and security on Solaris 10 systems.

Append the below parameters to /etc/system of both vmsol1 & vmsol2.

---Cut & Paste in /etc/system on vmsol1 & vmsol2---

```
set noexec_user_stack = 1
set semsys:seminfo_semmni = 100
set semsys:seminfo_semmns = 1024
set semsys:seminfo_semmsl = 256
set semsys:seminfo_semvmx = 32767
set shmsys:shminfo_shmmax = 4294967295
set shmsys:shminfo_shmmni = 100
set max_nprocs = 20000
set rlim_fd_max = 65536
set rlim_fd_cur = 1024
```

---Cut & Paste in /etc/system on vmsol1 & vmsol2---

***Below steps needs to be done on both vmsol1 & vmsol2*

Create Project for Grid user and set Kernel parameters

Login as root user (on vmsol1/2)

root# *projadd grid* <= This will create a project by name grid

Login into grid user,

```
root# su - grid
-bash-3.00$ id -p
uid=1100(grid) gid=1000(oinstall) projid=3(default)
```

Here grid is assigned "**default**" to change this to grid add a line shown below in **"/etc/user_attr"** as root user on **vmsol1/2**.

```
root# cat>>/etc/user_attr<<EOF
grid:::project=grid
EOF
```

Relogin into grid user

```
--log--
# su - grid
Sun Microsystems Inc. SunOS 5.10 Generic January 2005
-bash-3.00$ id -p
uid=1100(grid) gid=1000(oinstall) projid=100(grid)
-bash-3.00$
--log--
```

Here now grid user is assigned **project ID=100** and name **grid**.

Below are the Kernel parameters required to be set for the "grid" project.

Parameter	Replaced by Resource Control	Recommended Values
noexec_user_stack	NA	1
semsys:seminfo_semmni	project.max-sem-ids	100
semsys:seminfo_semmns	NA	1024
semsys:seminfo_semmns	process.max-sem-nsems	256
semsys:seminfo_semmns	NA	32767
shmsys:shminfo_shmmax	project.max-shm-memory	4294967295
shmsys:shminfo_shmmni	project.max-shm-ids	100

***The only value which needs to be modified here is of "project.max-shm-memory" in our install. Remaining values are either equal or more than above mentioned.*

Determine the current value of shared memory

```
grid$ prctl -n project.max-shm-memory -i project grid
--log--
-bash-3.00$ id -p
uid=1100(grid) gid=1000(oinstall) projid=100(grid)
-bash-3.00$ prctl -n project.max-shm-memory -i project grid
project: 100: grid
NAME PRIVILEGE VALUE FLAG ACTION RECIPIENT
project.max-shm-memory
  privileged 748MB - deny -
  system 16.0EB max deny -
-bash-3.00$
--log--
```

Permanently set the value in Grid project

```
root# projmod -sK "project.max-shm-memory=(privileged,4G,deny)" grid
```

Verify that the resource control is active and reflects the values:

```
cat /etc/project
--log--
# cat /etc/project
system:0:::
user.root:1:::
noproject:2:::
default:3:::
group.staff:10:::
grid:100:::project.max-shm-memory=(privileged,4294967296,deny)
--log--
```

****Below steps needs to be done on both vmsol1 & vmsol2**

Create Project for Oracle user and set Kernel parameters

Login as root user on vmsol1/2

```
root# projadd oracle <= This will create a project by name oracle
root# su - oracle
-bash-3.00$ id -p
uid=1101(oracle) gid=1000(oinstall) projid=3(default)
```

Here oracle is assigned "**default**" to change this to oracle add a line shown below in "**/etc/user_attr**" as root user in vmsol1/2

```
root# cat >> /etc/user_attr << EOF
oracle:::project=oracle
EOF
```

Relogin into oracle

```
--log--
-bash-3.00$ id -p
uid=1101(oracle) gid=1000(oinstall) projid=101(oracle)
-bash-3.00$
--log--
```

Here now oracle user is assigned **project ID=101** and name **oracle**.

Below are the Kernel parameters required to be set for the "**oracle**" project.

Parameter	Replaced by Resource Control	Recommended Values
noexec_user_stack	NA	1
semsys:seminfo_semmni	project.max-sem-ids	100
semsys:seminfo_semmns	NA	1024
semsys:seminfo_semmsl	process.max-sem-nsems	256
semsys:seminfo_semvmx	NA	32767
shmsys:shminfo_shmmax	project.max-shm-memory	4294967295
shmsys:shminfo_shmmni	project.max-shm-ids	100

The only value which needs to be modified here is of "**project.max-shm-memory**" in our install. Remaining values are either equal or more than above mentioned.

Determine the current value of shared memory

```
oracle$ prctl -n project.max-shm-memory -i project oracle
--log--
-bash-3.00$ id -p
uid=1101(oracle) gid=1000(oinstall) projid=101(oracle)
```

```
-bash-3.00$ prctl -n project.max-shm-memory -i project oracle
project: 101: oracle
NAME PRIVILEGE VALUE FLAG ACTION RECIPIENT
project.max-shm-memory
  privileged 748MB - deny -
  system 16.0EB max deny -
-bash-3.00$
--log--
```

Permanently set the value in Oracle project

```
root# projmod -sK "project.max-shm-memory=(privileged,4G,deny)" oracle
```

Verify that the resource control is active and reflects the values:

```
cat /etc/project
--log--
# cat /etc/project
system:0:::
user.root:1:::
noproject:2:::
default:3:::
group.staff:10:::
grid:100:::project.max-shm-memory=(privileged,4294967296,deny)
oracle:101:::project.max-shm-memory=(privileged,4294967296,deny)
#
--log--
```

12.7 Enable ssh daemon & setup ssh trust between vmsol1 & vmsol2

Enable sshd service on both vmsol1/2 as root user.

```
bothnodes# svcadm enable -t network/ssh:default
bothnodes# ps -ef|grep -i ssh
root 514 1 0 19:03:23 ? 0:00 /usr/lib/ssh/sshd
```

For this setup I have used Grid Infrastructure **in built ability to setup trust** between grid users and as a downside li will not be able to run CVU successfully. But while setting up Grid Infrastructure it would internally run CVU before start of Grid Install, so you can very will let us know the fixup script things if parameters necessary for the install is set Incorrectly.

12.8 Take a Snapshot of vmsol1 & vmsol2

Shutdown the vmsol1 & vmsol2 and follow Steps as mentioned in Step 7.5 to take a Snapshot and give each snap a new name and a description.

13. Configure Shared Storage (ZFS ISCSI Targets)

13.1 Enable RPC on all nodes

Execute as root on vmsol1, vmsol1 & vmsol2.

```
allnodes# svccfg -s svc:/network/rpc/bind setprop config/local_only=false
allnodes# svcadm refresh svc:/network/rpc/bind
```

13.2 Disable ISCSI Multipathing on vmsol1 & vmsol2:

As root user edit the file /kernel/drv/iscsi.conf and set `mpxio-disable="yes"`

13.3 Enable ISCSI on vmsol1 & vmsol2

Enable iscsi initiator:

```
bothnodes# svcadm enable svc:/network/iscsi/initiator:default
bothnodes# svcs -a | grep -i iscsi
--log--
# svcs -a | grep -i iscsi
disabled    10:25:33 svc:/system/iscsitgt:default
online      10:25:41 svc:/network/iscsi/initiator:default
--log--
```

13.4 Discover ISCSI Targets and Verify

Execute the below as **root** user on both vmsol1 & vmsol2

```
bothnodes# iscsiadm modify discovery --sendtargets enable
bothnodes# iscsiadm add discovery-address 192.168.1.96
bothnodes# devfsadm -i iscsi
```

Check on vmsol1 & vmsol2 if ISCSI volumes are mounted.

As root user on vmsol1 & vmsol2 execute command "format" to display all the disks mounted on the Solaris server.

```
# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
 0. c1t0d0 <DEFAULT cyl 5218 alt 2 hd 255 sec 63>
   /pci@0,0/pci15ad,1976@10/sd@0,0
 1. c2t600144F04B8A020600000C29F654C300d0 <DEFAULT cyl 1533 alt 2 hd 128 sec 32>
   /scsi_vhci/disk@g600144f04b8a020600000c29f654c300
 2. c2t600144F04B8A020800000C29F654C300d0 <DEFAULT cyl 1303 alt 2 hd 255 sec 63>
   /scsi_vhci/disk@g600144f04b8a020800000c29f654c300
 3. c2t600144F04B8A020900000C29F654C300d0 <DEFAULT cyl 1956 alt 2 hd 255 sec 63>
   /scsi_vhci/disk@g600144f04b8a020900000c29f654c300
Specify disk (enter its number): ^C
```

All ISCSI disks are mounted on vmsol1 & vmsol2.

Here you can see that apart from c1t0d0 on vmsol1/2, you see 3 other disks which are the Iscsi volumes mounted from vmutil server via ISCSI protocol.

Here these 3 ISCSI disks mounted have a long name format like below

"c2t600144F04B8A020600000C29F654C300d0" which is due to ISCSI Multipathing which is currently enabled. After a reboot (Step 13.2 iscsi multipathing is disabled) you would see the disk names changed to shorter naming pattern like "c1t1d0".

We can also Verify ISCSI connections to the vmutil server from vmsol1 & vmsol2.

Execute the below as root user on vmutil.

```
--log--
```

```
# iscsitadm list target -v
```

```
Target: iscsi_target/crs01
  iSCSI Name: iqn.1986-03.com.sun:02:0d7b372c-c3b2-6041-bd24-b2cea2a34761
  Alias: iscsi_target/crs01
  Connections: 2
  Initiator:
    iSCSI Name: iqn.1986-03.com.sun:01:2a9520f2ffff.4b89d1da
    Alias: vmsol2
  Initiator:
    iSCSI Name: iqn.1986-03.com.sun:01:2a9520f2ffff.4b8971d9
    Alias: vmsol1
  ACL list:
  TPGT list:
  LUN information:
    LUN: 0
    GUID: 600144f04b8a020600000c29f654c300
    VID: SUN
    PID: SOLARIS
    Type: dis
    Size: 3.0G
    Backing store: /dev/zvol/rdisk/iscsi_target/crs01
    Status: online
Target: iscsi_target/data01
  iSCSI Name: iqn.1986-03.com.sun:02:b5b07821-899e-4eb3-b8fd-c540f4ab05ff
  Alias: iscsi_target/data01
  Connections: 2
  Initiator:
    iSCSI Name: iqn.1986-03.com.sun:01:2a9520f2ffff.4b89d1da
    Alias: vmsol2
  Initiator:
    iSCSI Name: iqn.1986-03.com.sun:01:2a9520f2ffff.4b8971d9
    Alias: vmsol1
  ACL list:
  TPGT list:
  LUN information:
    LUN: 0
    GUID: 600144f04b8a020800000c29f654c300
    VID: SUN
    PID: SOLARIS
```

```

Type: disk
Size: 10G
Backing store: /dev/zvol/rdisk/iscsi_target/data01
Status: online
Target: iscsi_target/fra01
iSCSI Name: iqn.1986-03.com.sun:02:5b272782-6438-c793-a575-cbf42e8b9503
Alias: iscsi_target/fra01
Connections: 2
Initiator:
  iSCSI Name: iqn.1986-03.com.sun:01:2a9520f2ffff.4b89d1da
  Alias: vmsol2
Initiator:
  iSCSI Name: iqn.1986-03.com.sun:01:2a9520f2ffff.4b8971d9
  Alias: vmsol1
ACL list:
TPGT list:
LUN information:
  LUN: 0
  GUID: 600144f04b8a020900000c29f654c300
  VID: SUN
  PID: SOLARIS
  Type: disk
  Size: 15G
  Backing store: /dev/zvol/rdisk/iscsi_target/fra01
  Status: online
#
---log---
```

13.5 Format the ZFS ISCSI volumes from ONLY on vmsol1

As root user format the ISCSI volumes be used for ASM disks **only from vmsol1** as the partitions once created on vmsol1 will be automatically visible on vmsol2.

Here would "Skip first 3 cylinders" of the "Slice 1" for each disk while partitioning, as ASM writes header information in first few blocks of each disk.

On vmsol1 only

root# format [as shown below for the three disks]

***Follow the text marked in RED for formatting disks.*

```
--log--
```

```
# format
```

```
Searching for disks...done
```

```
AVAILABLE DISK SELECTIONS:
```

0. c1t0d0 <DEFAULT cyl 5218 alt 2 hd 255 sec 63>
/pci@0,0/pci15ad,1976@10/sd@0,0
1. c2t600144F04B8A020600000C29F654C300d0 <DEFAULT cyl 1533 alt 2 hd 128 sec 32>
/scsi_vhci/disk@g600144f04b8a020600000c29f654c300
2. c2t600144F04B8A020800000C29F654C300d0 <DEFAULT cyl 1303 alt 2 hd 255 sec 63>


```

/scsi_vhci/disk@g600144f04b8a020800000c29f654c300
3. c2t600144F04B8A020900000C29F654C300d0 <DEFAULT cyl 1956 alt 2 hd 255 sec 63>
/scsi_vhci/disk@g600144f04b8a020900000c29f654c300

```

Specify disk (enter its number): 1

```

selecting c2t600144F04B8A020600000C29F654C300d0
[disk formatted]

```

FORMAT MENU:

```

disk    - select a disk
type    - select (define) a disk type
partition - select (define) a partition table
current - describe the current disk
format  - format and analyze the disk
fdisk   - run the fdisk program
repair  - repair a defective sector
label   - write label to the disk
analyze - surface analysis
defect  - defect list management
backup  - search for backup labels
verify  - read and display labels
save    - save new disk/partition definitions
inquiry - show vendor, product and revision
volname - set 8-character volume name
!<cmd> - execute <cmd>, then return
quit

```

format> **fdisk**

No fdisk table exists. The default partition for the disk is:

a 100% "SOLARIS System" partition

Type "y" to accept the default partition, otherwise type "n" to edit the partition table.

y

format> **partition**

PARTITION MENU:

```

0  - change `0' partition
1  - change `1' partition
2  - change `2' partition
3  - change `3' partition
4  - change `4' partition
5  - change `5' partition
6  - change `6' partition
7  - change `7' partition
select - select a predefined table
modify - modify a predefined partition table
name   - name the current table
print  - display the current table

```

label - write partition map and label to the disk
 !<cmd> - execute <cmd>, then return
 quit

partition> **print**

Current partition table (original):

Total disk cylinders available: 1532 + 2 (reserved cylinders)

Part	Tag	Flag	Cylinders	Size	Blocks
0	unassigned	wm	0	0	(0/0/0) 0
1	unassigned	wm	0	0	(0/0/0) 0
2	backup	wu	0 - 1531	2.99GB	(1532/0/0) 6275072
3	unassigned	wm	0	0	(0/0/0) 0
4	unassigned	wm	0	0	(0/0/0) 0
5	unassigned	wm	0	0	(0/0/0) 0
6	unassigned	wm	0	0	(0/0/0) 0
7	unassigned	wm	0	0	(0/0/0) 0
8	boot	wu	0 - 0	2.00MB	(1/0/0) 4096
9	unassigned	wm	0	0	(0/0/0) 0

partition> **1**

Part	Tag	Flag	Cylinders	Size	Blocks
1	unassigned	wm	0	0	(0/0/0) 0

Enter partition id tag[unassigned]: **<Press Enter key>**

Enter partition permission flags[w]: **<Press Enter key>**

Enter new starting cyl[0]: **3**

Enter partition size[0b, 0c, 3e, 0.00mb, 0.00gb]: **1527c**

partition> **label**

Ready to label disk, continue? **y**

partition> **print**

Current partition table (unnamed):

Total disk cylinders available: 1532 + 2 (reserved cylinders)

Part	Tag	Flag	Cylinders	Size	Blocks
0	unassigned	wm	0	0	(0/0/0) 0
1	unassigned	wm	3 - 1529	2.98GB	(1527/0/0) 6254592
2	backup	wu	0 - 1531	2.99GB	(1532/0/0) 6275072
3	unassigned	wm	0	0	(0/0/0) 0
4	unassigned	wm	0	0	(0/0/0) 0
5	unassigned	wm	0	0	(0/0/0) 0
6	unassigned	wm	0	0	(0/0/0) 0
7	unassigned	wm	0	0	(0/0/0) 0
8	boot	wu	0 - 0	2.00MB	(1/0/0) 4096
9	unassigned	wm	0	0	(0/0/0) 0

partition> **quit**

format> **disk 2**

```
selecting c2t600144F04B8A020800000C29F654C300d0
[disk formatted]
```

```
format> fdisk
```

```
No fdisk table exists. The default partition for the disk is:
```

```
a 100% "SOLARIS System" partition
```

```
Type "y" to accept the default partition, otherwise type "n" to edit the
partition table.
```

```
y
```

```
format> partition
```

```
PARTITION MENU:
```

```
0 - change `0' partition
1 - change `1' partition
2 - change `2' partition
3 - change `3' partition
4 - change `4' partition
5 - change `5' partition
6 - change `6' partition
7 - change `7' partition
select - select a predefined table
modify - modify a predefined partition table
name - name the current table
print - display the current table
label - write partition map and label to the disk
!<cmd> - execute <cmd>, then return
quit
```

```
partition> print
```

```
Current partition table (original):
```

```
Total disk cylinders available: 1302 + 2 (reserved cylinders)
```

Part	Tag	Flag	Cylinders	Size	Blocks
0	unassigned	wm	0	0	(0/0/0) 0
1	unassigned	wm	0	0	(0/0/0) 0
2	backup	wu	0 - 1301	9.97GB	(1302/0/0) 20916630
3	unassigned	wm	0	0	(0/0/0) 0
4	unassigned	wm	0	0	(0/0/0) 0
5	unassigned	wm	0	0	(0/0/0) 0
6	unassigned	wm	0	0	(0/0/0) 0
7	unassigned	wm	0	0	(0/0/0) 0
8	boot	wu	0 - 0	7.84MB	(1/0/0) 16065
9	unassigned	wm	0	0	(0/0/0) 0

```
partition> 1
```

Part	Tag	Flag	Cylinders	Size	Blocks
1	unassigned	wm	0	0	(0/0/0) 0

Enter partition id tag[unassigned]: **<Press Enter key>**
 Enter partition permission flags[wm]: **<Press Enter key>**
 Enter new starting cyl[0]: **3**
 Enter partition size[0b, 0c, 3e, 0.00mb, 0.00gb]: **1297c**
 partition> **label**
 Ready to label disk, continue? **y**

partition> **print**
 Current partition table (unnamed):
 Total disk cylinders available: 1302 + 2 (reserved cylinders)

Part	Tag	Flag	Cylinders	Size	Blocks
0	unassigned	wm	0	0	(0/0/0) 0
1	unassigned	wm	3 - 1299	9.94GB	(1297/0/0) 20836305
2	backup	wu	0 - 1301	9.97GB	(1302/0/0) 20916630
3	unassigned	wm	0	0	(0/0/0) 0
4	unassigned	wm	0	0	(0/0/0) 0
5	unassigned	wm	0	0	(0/0/0) 0
6	unassigned	wm	0	0	(0/0/0) 0
7	unassigned	wm	0	0	(0/0/0) 0
8	boot	wu	0 - 0	7.84MB	(1/0/0) 16065
9	unassigned	wm	0	0	(0/0/0) 0

partition> **quit**
 format> **disk 3**
 selecting c2t600144F04B8A020900000C29F654C300d0
 [disk formatted]
 format> **fdisk**
 No fdisk table exists. The default partition for the disk is:

a 100% "SOLARIS System" partition

Type "y" to accept the default partition, otherwise type "n" to edit the partition table.

y
 format> **partition**

PARTITION MENU:
 0 - change `0' partition
 1 - change `1' partition
 2 - change `2' partition
 3 - change `3' partition
 4 - change `4' partition
 5 - change `5' partition
 6 - change `6' partition
 7 - change `7' partition
 select - select a predefined table
 modify - modify a predefined partition table

name - name the current table
 print - display the current table
 label - write partition map and label to the disk
 !<cmd> - execute <cmd>, then return
 quit

partition> **print**

Current partition table (original):

Total disk cylinders available: 1955 + 2 (reserved cylinders)

Part	Tag	Flag	Cylinders	Size	Blocks
0	unassigned	wm	0	0	(0/0/0) 0
1	unassigned	wm	0	0	(0/0/0) 0
2	backup	wu	0 - 1954	14.98GB	(1955/0/0) 31407075
3	unassigned	wm	0	0	(0/0/0) 0
4	unassigned	wm	0	0	(0/0/0) 0
5	unassigned	wm	0	0	(0/0/0) 0
6	unassigned	wm	0	0	(0/0/0) 0
7	unassigned	wm	0	0	(0/0/0) 0
8	boot	wu	0 - 0	7.84MB	(1/0/0) 16065
9	unassigned	wm	0	0	(0/0/0) 0

partition> **1**

Part	Tag	Flag	Cylinders	Size	Blocks
1	unassigned	wm	0	0	(0/0/0) 0

Enter partition id tag[unassigned]: **<Press Enter key>**

Enter partition permission flags[wu]: **<Press Enter key>**

Enter new starting cyl[0]: **3**

Enter partition size[0b, 0c, 3e, 0.00mb, 0.00gb]: **1950c**

partition> **label**

Ready to label disk, continue? **y**

partition> **print**

Current partition table (unnamed):

Total disk cylinders available: 1955 + 2 (reserved cylinders)

Part	Tag	Flag	Cylinders	Size	Blocks
0	unassigned	wm	0	0	(0/0/0) 0
1	unassigned	wm	3 - 1952	14.94GB	(1950/0/0) 31326750
2	backup	wu	0 - 1954	14.98GB	(1955/0/0) 31407075
3	unassigned	wm	0	0	(0/0/0) 0
4	unassigned	wm	0	0	(0/0/0) 0
5	unassigned	wm	0	0	(0/0/0) 0
6	unassigned	wm	0	0	(0/0/0) 0
7	unassigned	wm	0	0	(0/0/0) 0
8	boot	wu	0 - 0	7.84MB	(1/0/0) 16065
9	unassigned	wm	0	0	(0/0/0) 0

partition> **quit**

--log--

You can verify if these partitions are accessible on vmsol2 by below steps

Login as root user on vmsol2.

```
root# format
  [choose disk 1]
format> partition
partition> print
```

13.6 Reboot both vmsol1 & vmsol2 for disabling the ISCSI Multipathing(Step 13.2) to take effect.

```
bothnodes# reboot
```

13.7 Identify & Setup permissions for ISCSI shares

Execute below as root on vmsol1 & vmsol2 to see the new naming convention of the ISCSI disks.

```
--log--
bothnodes# iscsiadm list target -S|grep "OS Device Name:"
```

```
vmsol1# iscsiadm list target -S|grep "OS Device Name:"
OS Device Name: /dev/rdisk/c3t5d0s2 =>crs01
OS Device Name: /dev/rdisk/c3t6d0s2 =>data01
OS Device Name: /dev/rdisk/c3t7d0s2 =>fra01
```

```
vmsol2# iscsiadm list target -S|grep "OS Device Name:"
OS Device Name: /dev/rdisk/c3t5d0s2 =>crs01
OS Device Name: /dev/rdisk/c3t6d0s2 =>data01
OS Device Name: /dev/rdisk/c3t7d0s2 =>fra01
```

```
--log--
```

Here, on both nodes.

```
c3t5d0s2 => crs01
c3t6d0s2 => data01
c3t7d0s2 => fra01
```

We had formatted **Slice 1** of the above mentioned disks. So changing the permissions of **Slice 1 (c3t5d0s1, c3t6d0s1 & c3t7d0s1)** to group "grid:asmadmin" on both vmsol1 & vmsol2.

As root on vmsol1 & vmsol2

```
bothnodes# chown grid:asmadmin /dev/rdisk/c3t5d0s1 /dev/rdisk/c3t6d0s1 /dev/rdisk/c3t7d0s1
bothnodes# chmod 660 /dev/rdisk/c3t5d0s1 /dev/rdisk/c3t6d0s1 /dev/rdisk/c3t7d0s1
```

Check Permissions on vmsol1 & vmsol2

```
--log--
# cd /devices/iscsi
# ls -ltr |grep grid
crw-rw---- 1 grid  asmadmin  32, 385 Feb 28 11:54 disk@0000iqn.1986-03.com.sun%3A02%3Ab5b07821-
899e-4eb3-b8fd-c540f4ab05ff0001,0:b,raw
crw-rw---- 1 grid  asmadmin  32, 449 Feb 28 11:54 disk@0000iqn.1986-03.com.sun%3A02%3A5b272782-
6438-c793-a575-cbf42e8b95030001,0:b,raw
crw-rw---- 1 grid  asmadmin  32, 321 Feb 28 11:54 disk@0000iqn.1986-03.com.sun%3A02%3A0d7b372c-
c3b2-6041-bd24-b2cea2a347610001,0:b,raw
#
--log--
```

13.8 Take a Snapshot of vmsol1,vmsol2 & vmutil.

Shutdown the vmsol1, vmsol2 & vmutil and follow Steps as mentioned in Step 7.5 to take a Snapshot and give each snap a new name and description.

14. Install Oracle 11gR2 Grid Infrastructure

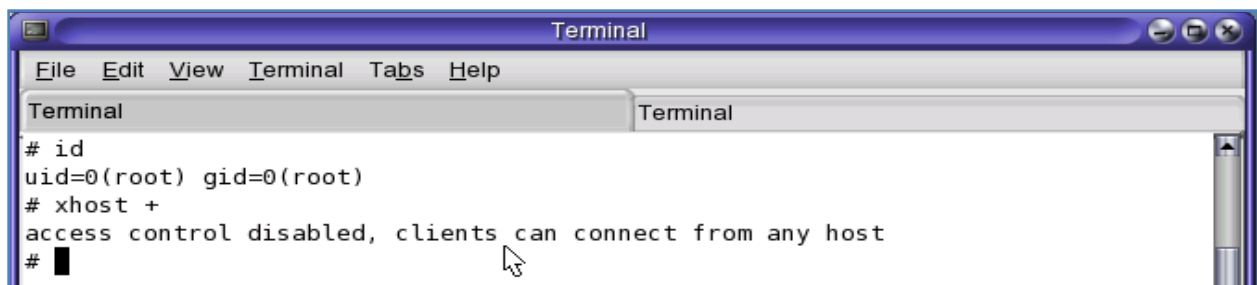
14.1 Stage the binaries for Oracle 11gR2 Grid Infrastructure in the shared folder (/hgfs/shared). So that we can access the binaries from Host machine on to vmsol1 & vmsol2 via /hgfs filesystem.

14.2 Login as root user in vmsol1 and open a "Terminal" with 2 tab's(2 sessions).

Execute "xhost +" on first tab and then in second tab switch to grid user from root user and set the DISPLAY.

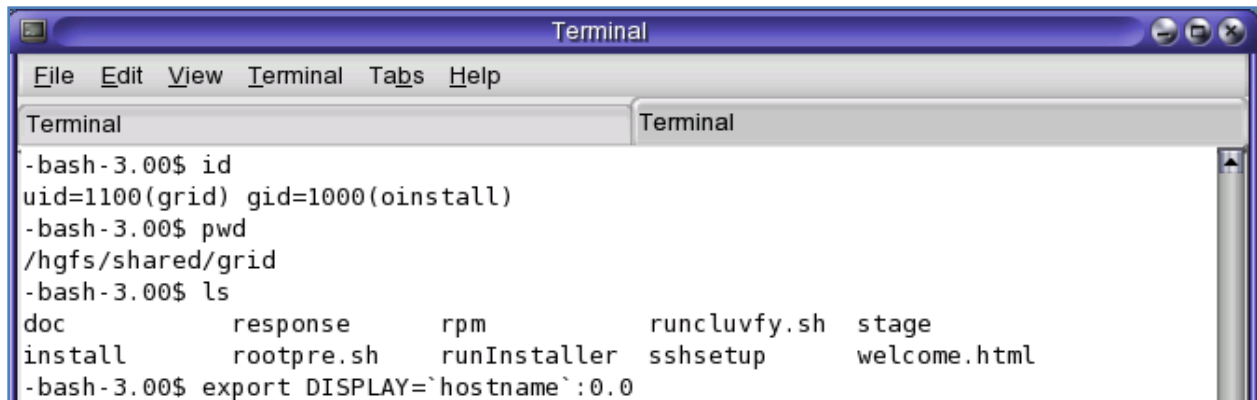
---In First tab---

```
# id
uid=0(root) gid=0(root)
# xhost +
access control disabled, clients can connect from any host
---In First tab---
```



---In second tab---

```
su - grid
cd /hgfs/shared/grid
-bash-3.00$ export DISPLAY=`hostname`:0.0
---In second tab---
```



```

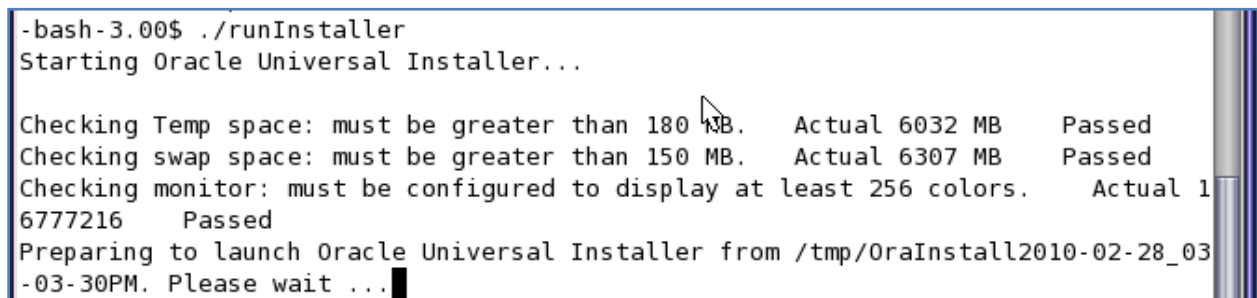
Terminal
File Edit View Terminal Tabs Help
Terminal
Terminal
-bash-3.00$ id
uid=1100(grid) gid=1000(oinstall)
-bash-3.00$ pwd
/hgfs/shared/grid
-bash-3.00$ ls
doc                response          rpm               runcluvfy.sh     stage
install           rootpre.sh       runInstaller     sshsetup         welcome.html
-bash-3.00$ export DISPLAY=`hostname`:0.0

```

14.3 Kick off Grid Infrastructure Installation

In the second tab shown in Step 14.2 (you are logged in as grid user) and *vmutil server must be running*. Check if you are able to *resolve all hosts* (Step 10.5).

As "grid" user execute below in **second tab**,
./runInstaller &



```

-bash-3.00$ ./runInstaller
Starting Oracle Universal Installer...

Checking Temp space: must be greater than 180 MB.   Actual 6032 MB   Passed
Checking swap space: must be greater than 150 MB.   Actual 6307 MB   Passed
Checking monitor: must be configured to display at least 256 colors.   Actual 16777216   Passed
Preparing to launch Oracle Universal Installer from /tmp/OraInstall2010-02-28_03-03-30PM. Please wait ...

```

This will start Oracle OUI (Oracle 11g Grid Infrastructure's first screen in installation)

Screen Name	Choices/Options
Select Installation Option	Choose "Install and configure Grid Infrastructure for a cluster" Click "Next"
Select Installation Type	Choose "Advanced Installation" Click "Next"
Select Product Languages	Choose "English" [Default] Click "Next"

Grid Plug and Play Information

Enter the below as shown in screen shoot.

*Cluster Name: **vmcluster***

*SCAN Name : **vmcluster-scan***

We are not using GNS [leave it unchecked]

Single Client Access Name (SCAN) allows clients to use one name as a whole. Client connect requests to the SCAN name can be

Cluster Name:

SCAN Name:

SCAN Port:

Configure GNS

GNS Sub Domain:

For example: grid.example.com

GNS VIP Address:

Click "Next"

Cluster Node Information

You will already see vmsol1 & vmsol1-vip in screen

Hostname	Virtual IP Name
vmsol1	vmsol1-vip

Click on the "**Add**" button as it pops up another window [Add Cluster Node Information]. Enter the below in the Text Boxes.

*Hostname : **vmsol2***

*Virtual IP name : **vmsol2-vip***

Click "OK"

Now you will see vmsol2 also registered similar to vmsol1.

Hostname	Virtual IP Name
vmsol1	vmsol1-vip
vmsol2	vmsol2-vip

Now Click on "**SSH Connectivity**" button, it will unhide an area below the button where OS Username: **grid** [non-editable] and a Text box with Label OS Password appears. Enter the **grid OS user password** in it.

Click "**Setup**" Button.[this will setup trust between grid user on vmsol1 & vmsol2].

Click "Next"

This would recheck ssh connectivity between grid user on both nodes and pops up a confirmation window on success. Click "OK".

Click "Next"

Specify Network Interface Usage

aggr1 : Private [For private Interconnect]
e1000g0 : Public

Interface Name	Subnet	Interface Type
aggr1	172.168.0.0	Private
e1000g0	192.168.1.0	Public

Click "Next"

Storage Option Information

Choose "**Automatic Storage Management**"
 Click "Next"

Create ASM Disk Group

Choose as shown in below Image.

Disk group Name : **CRS**
 Redundancy : **External**
 Choose radio : **Canidate Disks**
 Check box : **/dev/rdisk/c3t5d0s1**

Select Disk Group Characteristics and select disks

Disk Group Name

Redundancy High Normal External

Add Disks

Candidate Disks All Disks

<input type="checkbox"/>	Disk Path	Size (in MB)
<input checked="" type="checkbox"/>	/dev/rdisk/c3t5d0s1	3054
<input type="checkbox"/>	/dev/rdisk/c3t6d0s1	10173
<input type="checkbox"/>	/dev/rdisk/c3t7d0s1	15296

Click "Next"

ASM Password

Choose **"Use same passwords for these accounts"**

Fill in the same passwords in both text boxes.

Click "Next"

Privileged Operating System groups

Leave default selected values.

ASM Database Administrator (OSDBA) Group

ASM Instance Administration Operator (OSOPER) Group

ASM Instance Administrator (OSASM) Group

Click "Next"

Specify Installation Location

Oracle Base : **/u01/app/grid**

Software Location : **/u01/app/11.2.0/grid**

Click "Next"

Create Inventory

Inventory Directory : **/u01/app/orainventory**

Click "Next"

Now CVU does the verification and will fail. If any pre-requisites

are missing. In our case all are met.

Summary

Click "Finish", to start the Grid Infrastructure Installation.

Progress

100%

Saving Cluster Inventory

Status

✓	Install Grid Infrastructure for a Cluster	Succeeded
✓	• Prepare	Succeeded
✓	• Copy files	Succeeded
✓	• Link binaries	Succeeded
✓	• Setup files	Succeeded
✓	• Perform remote operations	Succeeded
	Execute Root Scripts for Install Grid Infrastructure for a Cluster	Pending

Execute scripts as root on vmsol1 & vmsol2

Scripts to be executed:

Number	Script Location	Nodes
1	/u01/app/oraInventory/orainstRoot.sh	vmsol1,vmsol2
2	/u01/app/11.2.0/grid/root.sh	vmsol1,vmsol2

Execute "orainstRoot.sh" on vmsol1 and then on vmsol2 as "root" user.

```
# /u01/app/oraInventory/orainstRoot.sh
Changing permissions of /u01/app/oraInventory.
Adding read,write permissions for group.
Removing read,write,execute permissions for world.

Changing groupname of /u01/app/oraInventory to oinstall
The execution of the script is complete.
# █
```

Execute "root.sh" on vmsol1 and then on vmsol2 as "root" user.

On vmsol1:

```
# /u01/app/11.2.0/grid/root.sh
Running Oracle 11g root.sh script...
```

The following environment variables are set as:

```
ORACLE_OWNER= grid
```

```
ORACLE_HOME= /u01/app/11.2.0/grid
```

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

```
Creating /usr/local/bin directory...
```

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

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

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

```
Creating /var/opt/oracle/oratab file...
```

```
Entries will be added to the /var/opt/oracle/oratab file as needed by  
Database Configuration Assistant when a database is created
```

```
Finished running generic part of root.sh script.
```

```
Now product-specific root actions will be performed.
```

```
2010-02-28 15:28:51: Parsing the host name
```

```
2010-02-28 15:28:51: Checking for super user privileges
```

```
2010-02-28 15:28:51: User has super user privileges
```

```
Using configuration parameter file:
```

```
/u01/app/11.2.0/grid/crs/install/crsconfig_params
```

```
Creating trace directory
```

```
LOCAL ADD MODE
```

```
Creating OCR keys for user 'root', privgrp 'root'..
```

```
Operation successful.
```

```
root wallet
```

```
root wallet cert
```

```
root cert export
```

```
peer wallet
```

```
profile reader wallet
```

```
pa wallet
```

```
peer wallet keys
```

```
pa wallet keys
```

```
peer cert request
```

```
pa cert request
```

```
peer cert
```

```
pa cert
```

```
peer root cert TP
```

```
profile reader root cert TP
```

```
pa root cert TP
```

```
peer pa cert TP
```

```
pa peer cert TP
```

```
profile reader pa cert TP
```

```
profile reader peer cert TP
```

```
peer user cert
```

```
pa user cert
```

```
Adding daemon to inittab
```

```
CRS-4123: Oracle High Availability Services has been started.
```

```
ohasd is starting
```

```
CRS-2672: Attempting to start 'ora.gipcd' on 'vmsol1'
```

```
CRS-2672: Attempting to start 'ora.mdnsd' on 'vmsol1'
```

```
CRS-2676: Start of 'ora.gipcd' on 'vmsol1' succeeded
```

```
CRS-2676: Start of 'ora.mdnsd' on 'vmsol1' succeeded
```

```
CRS-2672: Attempting to start 'ora.gpnpd' on 'vmsol1'
```

```

CRS-2676: Start of 'ora.gpnpd' on 'vmsol1' succeeded
CRS-2672: Attempting to start 'ora.cssdmonitor' on 'vmsol1'
CRS-2676: Start of 'ora.cssdmonitor' on 'vmsol1' succeeded
CRS-2672: Attempting to start 'ora.cssd' on 'vmsol1'
CRS-2672: Attempting to start 'ora.diskmon' on 'vmsol1'
CRS-2676: Start of 'ora.diskmon' on 'vmsol1' succeeded
CRS-2676: Start of 'ora.cssd' on 'vmsol1' succeeded
CRS-2672: Attempting to start 'ora.ctssd' on 'vmsol1'
CRS-2676: Start of 'ora.ctssd' on 'vmsol1' succeeded

```

ASM created and started successfully.

DiskGroup CRS created successfully.

```

clscfg: -install mode specified
Successfully accumulated necessary OCR keys.
Creating OCR keys for user 'root', privgrp 'root'..
Operation successful.
CRS-2672: Attempting to start 'ora.crsd' on 'vmsol1'
CRS-2676: Start of 'ora.crsd' on 'vmsol1' succeeded
CRS-4256: Updating the profile
Successful addition of voting disk 78ab93c0751a4f29bf38ac1e72a4f62e.
Successfully replaced voting disk group with +CRS.
CRS-4256: Updating the profile
CRS-4266: Voting file(s) successfully replaced

```

```

## STATE   File Universal Id         File Name Disk group
--  -

```

```

1. ONLINE 78ab93c0751a4f29bf38ac1e72a4f62e (/dev/rdisk/c3t5d0s1)

```

[CRS]

```

Located 1 voting disk(s).
CRS-2673: Attempting to stop 'ora.crsd' on 'vmsol1'
CRS-2677: Stop of 'ora.crsd' on 'vmsol1' succeeded
CRS-2679: Attempting to clean 'ora.crsd' on 'vmsol1'
CRS-2681: Clean of 'ora.crsd' on 'vmsol1' succeeded
CRS-2673: Attempting to stop 'ora.asm' on 'vmsol1'
CRS-2677: Stop of 'ora.asm' on 'vmsol1' succeeded
CRS-2673: Attempting to stop 'ora.ctssd' on 'vmsol1'
CRS-2677: Stop of 'ora.ctssd' on 'vmsol1' succeeded
CRS-2673: Attempting to stop 'ora.cssdmonitor' on 'vmsol1'
CRS-2677: Stop of 'ora.cssdmonitor' on 'vmsol1' succeeded
CRS-2673: Attempting to stop 'ora.cssd' on 'vmsol1'
CRS-2677: Stop of 'ora.cssd' on 'vmsol1' succeeded
CRS-2673: Attempting to stop 'ora.gpnpd' on 'vmsol1'
CRS-2677: Stop of 'ora.gpnpd' on 'vmsol1' succeeded
CRS-2673: Attempting to stop 'ora.gipcd' on 'vmsol1'
CRS-2677: Stop of 'ora.gipcd' on 'vmsol1' succeeded
CRS-2673: Attempting to stop 'ora.mdnsd' on 'vmsol1'
CRS-2677: Stop of 'ora.mdnsd' on 'vmsol1' succeeded
CRS-2672: Attempting to start 'ora.mdnsd' on 'vmsol1'
CRS-2676: Start of 'ora.mdnsd' on 'vmsol1' succeeded
CRS-2672: Attempting to start 'ora.gipcd' on 'vmsol1'
CRS-2676: Start of 'ora.gipcd' on 'vmsol1' succeeded

```

```

CRS-2672: Attempting to start 'ora.gpnpd' on 'vmsol1'
CRS-2676: Start of 'ora.gpnpd' on 'vmsol1' succeeded
CRS-2672: Attempting to start 'ora.cssdmonitor' on 'vmsol1'
CRS-2676: Start of 'ora.cssdmonitor' on 'vmsol1' succeeded
CRS-2672: Attempting to start 'ora.cssd' on 'vmsol1'
CRS-2672: Attempting to start 'ora.diskmon' on 'vmsol1'
CRS-2676: Start of 'ora.diskmon' on 'vmsol1' succeeded
CRS-2676: Start of 'ora.cssd' on 'vmsol1' succeeded
CRS-2672: Attempting to start 'ora.ctssd' on 'vmsol1'
CRS-2676: Start of 'ora.ctssd' on 'vmsol1' succeeded
CRS-2672: Attempting to start 'ora.asm' on 'vmsol1'
CRS-2676: Start of 'ora.asm' on 'vmsol1' succeeded
CRS-2672: Attempting to start 'ora.crsd' on 'vmsol1'
CRS-2676: Start of 'ora.crsd' on 'vmsol1' succeeded
CRS-2672: Attempting to start 'ora.evmd' on 'vmsol1'
CRS-2676: Start of 'ora.evmd' on 'vmsol1' succeeded
CRS-2672: Attempting to start 'ora.asm' on 'vmsol1'
CRS-2676: Start of 'ora.asm' on 'vmsol1' succeeded
CRS-2672: Attempting to start 'ora.CRS.dg' on 'vmsol1'
CRS-2676: Start of 'ora.CRS.dg' on 'vmsol1' succeeded

```

```

vmsol1 2010/02/28 15:36:27
/u01/app/11.2.0/grid/cdata/vmsol1/backup_20100228_153627.olr
Configure Oracle Grid Infrastructure for a Cluster ... succeeded
Updating inventory properties for clusterware
Starting Oracle Universal Installer...

```

```

Checking swap space: must be greater than 500 MB. Actual 5867 MB
Passed
The inventory pointer is located at /var/opt/oracle/orainst.loc
The inventory is located at /u01/app/orainventory
'UpdateNodeList' was successful.
#

```

All the CRS services are up and running on vmsol1. Now root.sh needs to run as root on vmsol2.

On vmsol2:

```

# /u01/app/11.2.0/grid/root.sh
Running Oracle 11g root.sh script...

```

```

The following environment variables are set as:
ORACLE_OWNER= grid
ORACLE_HOME= /u01/app/11.2.0/grid

```

```

Enter the full pathname of the local bin directory: [/usr/local/bin]:
Creating /usr/local/bin directory...
Copying dbhome to /usr/local/bin ...
Copying oraenv to /usr/local/bin ...
Copying coraenv to /usr/local/bin ...

```

```

Creating /var/opt/oracle/oratab file...

```

```

Entries will be added to the /var/opt/oracle/oratab file as needed by
Database Configuration Assistant when a database is created
Finished running generic part of root.sh script.
Now product-specific root actions will be performed.
2010-02-28 15:39:53: Parsing the host name
2010-02-28 15:39:53: Checking for super user privileges
2010-02-28 15:39:53: User has super user privileges
Using configuration parameter file:
/u01/app/11.2.0/grid/crs/install/crsconfig_params
Creating trace directory
LOCAL ADD MODE
Creating OCR keys for user 'root', privgrp 'root'..
Operation successful.
Adding daemon to inittab
CRS-4123: Oracle High Availability Services has been started.
ohasd is starting
CRS-4402: The CSS daemon was started in exclusive mode but found an
active CSS daemon on node vmsol1, number 1, and is terminating
An active cluster was found during exclusive startup, restarting to join the
cluster
CRS-2672: Attempting to start 'ora.mdnsd' on 'vmsol2'
CRS-2676: Start of 'ora.mdnsd' on 'vmsol2' succeeded
CRS-2672: Attempting to start 'ora.gipcd' on 'vmsol2'
CRS-2676: Start of 'ora.gipcd' on 'vmsol2' succeeded
CRS-2672: Attempting to start 'ora.gpnpd' on 'vmsol2'
CRS-2676: Start of 'ora.gpnpd' on 'vmsol2' succeeded
CRS-2672: Attempting to start 'ora.cssdmonitor' on 'vmsol2'
CRS-2676: Start of 'ora.cssdmonitor' on 'vmsol2' succeeded
CRS-2672: Attempting to start 'ora.cssd' on 'vmsol2'
CRS-2672: Attempting to start 'ora.diskmon' on 'vmsol2'
CRS-2676: Start of 'ora.diskmon' on 'vmsol2' succeeded
CRS-2676: Start of 'ora.cssd' on 'vmsol2' succeeded
CRS-2672: Attempting to start 'ora.ctssd' on 'vmsol2'
CRS-2676: Start of 'ora.ctssd' on 'vmsol2' succeeded
CRS-2672: Attempting to start 'ora.asm' on 'vmsol2'
CRS-2676: Start of 'ora.asm' on 'vmsol2' succeeded
CRS-2672: Attempting to start 'ora.crsd' on 'vmsol2'
CRS-2676: Start of 'ora.crsd' on 'vmsol2' succeeded
CRS-2672: Attempting to start 'ora.evmd' on 'vmsol2'
CRS-2676: Start of 'ora.evmd' on 'vmsol2' succeeded

vmsol2 2010/02/28 15:43:20
/u01/app/11.2.0/grid/cdata/vmsol2/backup_20100228_154320.olr
Configure Oracle Grid Infrastructure for a Cluster ... succeeded
Updating inventory properties for clusterware
Starting Oracle Universal Installer...

Checking swap space: must be greater than 500 MB. Actual 5886 MB
Passed
The inventory pointer is located at /var/opt/oracle/orainst.loc
The inventory is located at /u01/app/oraInventory
'UpdateNodeList' was successful.

```

```
#
Click "OK".
This completes the Grid Infrastructure Installation
```

Total Time Taken **90 minutes**

14.3 Shutdown vmsol1, vmsol2 & vmutil and take a Snapshot.

Shutdown the Cluster ready services (CRS) on vmsol1 & vmsol2

As root user.

```
bothnodes# /u01/app/11.2.0/grid/bin/crsctl stop crs
```

Shutdown VM's vmsol1, vmsol2 and then vmutil.

Take a snapshot and give it a name (like sync-all-1) and description.

15. Install Oracle 11gR2 RDBMS Binaries

Start the vmutil server.

15.1 Stage the binaries for Oracle 11gR2 Database binaries in the shared folder (/hgfs/shared).

So that we can access the binaries from Host machine on to vmsol1 & vmsol2 via /hgfs filesystem.

15.2 Verify if CRS is running on both nodes.

As root user.

```
bothnodes# /u01/app/11.2.0/grid/bin/crsctl crs status
```

15.3 Login as root user in vmsol1 and open a "Terminal" with 2 tab's (2 sessions).

Execute "xhost +" on first tab and then in second tab switch to oracle user from root user and set the DISPLAY.

---In First tab---

```
# id
uid=0(root) gid=0(root)
# xhost +
access control disabled, clients can connect from any host
```

---In First tab---

---In second tab---

```
su - oracle
cd /hgfs/shared/database
```

```
-bash-3.00$ export DISPLAY=`hostname`:0.0
```

---In second tab---

As "oracle" user execute below in **second tab**,

./runInstaller &

Screen Name	Choice/Options
Configure Security Updates	For this article, I left all blank. Click "Next"
Installation Option	Choose " Install database software only " Click "Next"
Node Selection	Choose " Real Application Clusters database Installation " Check both vmsol1 & vmsol2 Now Click on " SSH Connectivity " button, it will unhide an area below the button where <i>OS Username: oracle</i> [non-editable] and a Text box with Label <i>OS Password</i> appears. Enter the oracle OS user password in it. Click " Setup " Button.[this will setup trust between oracle user on vmsol1 & vmsol2]. Click "Next".
Select Product Languages	Choose "English" [default] Click "Next"
Select Database Edition	Choose " Enterprise Edition(4.55 GB) " Click "Next"
Specify Installation Location	Oracle Base : /u01/app/oracle Software Location : /u01/app/oracle/product/11.2.0/dbhome_1 Click "Next"
Privileged Operating System groups	OSDBA group : dba OSOPER group : oper Click "Next"
Perform Prerequisite check	It will run CVU internally for pre-requisite check for RDBMS install.
Summary	Click "Finish" to start the RDBMS S/W installation.
Execute root.sh on both vmsol1 & vmsol2	At the end of installation, a pop-up window will instruct to execute "root.sh" on vmsol1 & vmsol2. Execute "root.sh" on vmsol1 & vmsol2.

```

# /u01/app/oracle/product/11.2.0/dbhome_1/root.sh
Running Oracle 11g root.sh script...

The following environment variables are set as:
  ORACLE_OWNER= oracle
  ORACLE_HOME=
/u01/app/oracle/product/11.2.0/dbhome_1

Enter the full pathname of the local bin directory:
[/usr/local/bin]:
The file "dbhome" already exists in /usr/local/bin.
Overwrite it? (y/n) [n]: y
  Copying dbhome to /usr/local/bin ...
The file "oraenv" already exists in /usr/local/bin.
Overwrite it? (y/n) [n]: y
  Copying oraenv to /usr/local/bin ...
The file "coraenv" already exists in /usr/local/bin.
Overwrite it? (y/n) [n]: y  Copying coraenv to /usr/local/bin
...

Entries will be added to the /var/opt/oracle/oratab file as
needed by
  Database Configuration Assistant when a database is
created
  Finished running generic part of root.sh script.
  Now product-specific root actions will be performed.
  Finished product-specific root actions.
#
Click "Ok"

```

Finish	Completes the installation. Click "Close" to exit OUI.
--------	---

Total Time Taken	60 mins
-------------------------	----------------

15.4 Shutdown vmsol1, vmsol2 & vmutil and take a snapshot.

Shutdown the cluster ready services (CRS) on vmsol1 & vmsol2

As root user.

```
bothnodes# /u01/app/11.2.0/grid/bin/crsctl stop crs
```

Shutdown vmsol1, vmsol2 and then vmutil.

Take a Snapshot and give it a name (sync-all-2) and a description.

16. Create Oracle 11gR2 Real Application Cluster Database using DBCA.

Start first vmutl (log in) then vmsol1 & vmsol2

16.1 Verify is CRS is running on both nodes.

As root user.

```
bothnodes# /u01/app/11.2.0/grid/bin/crsctl crs status
```

16.2 Create ASM disk groups using asmca

Create ASM disk groups for Datafiles (DATA) & Flash recovery area (FRA).

Login as root user in vmsol1 and open a "Terminal" with 2 tab's (2 sessions).

Execute "xhost +" on first tab and then in second tab switch to grid user from root user and set the DISPLAY.

On vmsol1

---In First tab---

```
# id
```

```
uid=0(root) gid=0(root)
```

```
# xhost +
```

```
access control disabled, clients can connect from any host
```

---In First tab---

---In second tab---

```
su - grid
```

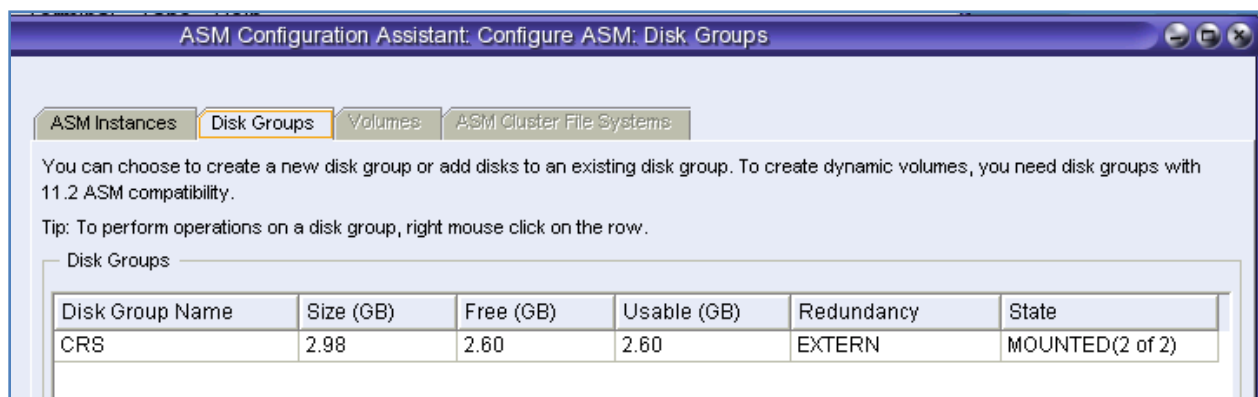
```
-bash-3.00$ export DISPLAY=`hostname`:0.0
```

```
-bash-3.00$ asmca &
```

---In second tab---

This will open window by name "ASM Configuration Assistant Window: Configure ASM:Disk Groups"

Click on the "Disk Groups" tab. You would already see Disk Group "CRS".



Click on the "**Create**" button on the bottom left of the window.
In Window "**Create Disk Group**"

Enter Disk Group Name: **DATA**
Redundancy : **External**
Select Member Disks : **Show Eligible**
Disk Path [Check] : **/dev/rdisk/c3t6d0s1**
Click "OK".

Disk Group Name

Redundancy

Redundancy is achieved by storing multiple copies of the data on different failure groups. Normal redundancy needs 2 failure groups, and high redundancy from atleast three different failure groups.

High Normal External (None)

Select Member Disks

Show Eligible Show All

Quorum failure groups are used to store voting files in extended clusters and do not contain any user data. It requires higher.

<input type="checkbox"/>	Disk Path	Header Status	Disk Name	Size (MB)	Quorum
<input checked="" type="checkbox"/>	/dev/rdisk/c3t6d0s1	CANDIDATE		10173	<input type="checkbox"/>
<input type="checkbox"/>	/dev/rdisk/c3t7d0s1	CANDIDATE		15296	<input type="checkbox"/>

This will take upto 2 mins to create the diskgroup and you will be taken back to the "**ASM Configuration Assistant Window:Configure ASM:Disk Groups**" => **Diskgroups** where now you should see **2 diskgroups CRS & DATA**.

Click on the "**Create**" button on the bottom left of the window.
In Window "**Create Disk Group**"

Enter Disk Group Name: **FRA**
Redundancy : **External**
Select Member Disks : **Show Eligible**
Disk Path [Check] : **/dev/rdisk/c3t7d0s1**
Click "OK".

Disk Group Name

Redundancy

Redundancy is achieved by storing multiple copies of the data on different failure groups. Normal redundancy needs failure groups, and high redundancy from atleast three different failure groups.

High Normal External (None)

Select Member Disks

Show Eligible Show All

Quorum failure groups are used to store voting files in extended clusters and do not contain any user data. It requires higher.

<input type="checkbox"/>	Disk Path	Header Status	Disk Name	Size (MB)	Quorum
<input checked="" type="checkbox"/>	/dew/rdisk/c3t7d0s1	CANDIDATE		15296	<input type="checkbox"/>

This will take upto 2 mins to create the diskgroup and you will be taken back to the "ASM Configuration Assistant Window:Configure ASM:Disk Groups" =>Diskgroups where now you should see **3 diskgroups CRS,DATA & FRA**.

Disk Group Name	Size (GB)	Free (GB)	Usable (GB)	Redundancy	State
DATA	9.93	9.84	9.84	EXTERN	MOUNTED(2 of 2)
FRA	14.94	14.84	14.84	EXTERN	MOUNTED(2 of 2)
CRS	2.98	2.60	2.60	EXTERN	MOUNTED(2 of 2)

Now you have DATA & FRA diskgroups created for creating a RAC database. Click on "Exit" to quit the asmca.

16.3 Create RAC database using dbca.

Execute "xhost +" on first tab and then in second tab switch to oracle user from root user and set the DISPLAY.

On vmsol1

---In First tab---

```
# id
```

```
uid=0(root) gid=0(root)
```

```
# xhost +
```

```
access control disabled, clients can connect from any host
```

---In First tab---

---In second tab---

```
su - oracle
```

```
-bash-3.00$ export DISPLAY=`hostname`:0.0
```

---In second tab---

As "oracle" user execute below in **second tab**,

dbca &

Screen Name	Choice/Options
Database Configuration Assistant: Welcome	Choose " Oracle Real Application Clusters database " Click "Next"
Database Configuration Assistant: Step 1 of 13: Operations	Choose " Create a Database " Click "Next"
Database Configuration Assistant: Step 2 of 13: Database Templates	Choose " Custom Database " Click "Next"
Database Configuration Assistant: Step 3 of 13: Database Identification	Choose Configuration Type as " Admin Managed " Enter Global Database Name: racdb.mydom.com SID Prefix : racdb Click " Select All " to choose vmsol1 & vmsol2 Click "Next"
Database Configuration Assistant: Step 4 of 12: Management Options	In the " Enterprise Manager " tab. Check " Configure Enterprise Manager " [Default] Choose " Configure database control for local management " [Default]
Database Configuration Assistant: Step 5 of 12: Database Credentials	Choose " Use the same Administrative Password For All Accounts " Enter the password twice Click "Next"
Database Configuration Assistant: Step 6 of 12: Database File Locations	Storage Type " Automatic Storage Management(ASM) " Choose " Use Oracle-Managed Files " Enter for Label Database Area : +DATA Click "Next"  This will pop-up " ASM Credentials " Window Enter password and Click " OK "
Database Configuration Assistant: Step 7 of 12: Recovery Configuration	Check the option " Specify Flash Recovery Area " Flash Recovery Area : +FRA Flash Recovery Area Size : 15000 Check the option " Enable Archiving "

Click "Next"

Specify Flash Recovery Area

This is used as the default for all disk based backup and recovery operations required for automatic disk based backup using Enterprise Manager. The database files and recovery files be located on physically different disks to improve backup and performance.

Flash Recovery Area:

Flash Recovery Area Size:

Enable Archiving Edit Archive Mode Parameters

- Database Configuration Assistant: Take all Default values.
Step 8 of 12: Database Content Click "Next"
- Database Configuration Assistant: Choose "**Typical**"
Step 9 of 11: Initialization Keep the option checked for "**Use Automatic Memory Management**"
Parameters Click "Next"
- Database Configuration Assistant: Click "Next"
Step 10 of 11: Database Storage
- Database Configuration Assistant: Check "**Create Database**"
Step 11 of 11: Creation Options Click "**Finish**"
This will pop-up as **Summary** window, Click "**OK**"

** Give it a minute or two, it will start the database creation with % completed.

** When the Installation is 100% it pops-up Another window Click "exit" .

Total Time taken by dbca 270 mins

Now we have a fully functioning RAC Database.

16.4 Shutdown vmsol1, vmsol2 & vmutil and take a snapshot.

Shutdown the cluster ready services (CRS) and RAC Database on vmsol1 & vmsol2

As root user.

```
bothnodes# /u01/app/11.2.0/grid/bin/crsctl stop crs
```


Shutdown vmsol1, vmsol2 and then vmutil.
Take a Snapshot and give it a Name (sync-all-3) and Description.

17. Few commands to administer CRS & RAC Database.

Start vmutil server first, once you are able to login. Start vmsol1 & vmsol2.

Check for crs, asm & rac database availability.

```
grid $ crs_stat -t -v
```

You can now further go-ahead and create users, tablespaces (temp & data) and go on with your experimenting.

- Check health of CRS:

```
grid $ crsctl check crs
```

- Check status of All instances

```
oracle $ srvctl status instance -d racdb
```

- Check ASM status

```
oracle $ srvctl status asm
```

- Check SCAN Status & configuration.

```
oracle $ srvctl status scan
```

```
oracle $ srvctl config scan
```

- Check status of Node Applications

```
oracle $ srvctl status nodeapps
```

- Verify Clock Synchronization demon: CTSS

```
oracle $ cluvfy comp clocksync -verbose
```

- Start & Stop CRS and RAC database

As root user:

```
bothnodes# /u01/app/11.2.0/grid/bin/crsctl stop crs [stops all]
```

```
bothnodes# /u01/app/11.2.0/grid/bin/crsctl start crs [start all]
```

Check srvctl & crsctl help for other options and RAC Admin guide from oracle Documentation

18. Create a Vmware Team for vmutil, vmsol1& vmsol2 [optional]

The VM startup sequence for vmutil, vmsol1 & vmsol2 should be:

1. Startup vmutil and login.

The reason for vmutil to be up before vmsol1 & vmsol2 is, vmutil provides the ISCSI shared storage for CRS & ASM across vmsol1/2. So when the rac nodes (vmsol1/2) boots CRS should be able to find the Shared disks.

2. You can start vmsol1 & vmsol2 parallelly or with a time gap of a minute.

Vmware Workstation provides a feature a creating a Vmware Team, which is a collection of VM's whose boot sequence order can set.

Lets create a Team "Solaris 11gR2 RAC Team" which includes VM's vmutil, vmsol1 & vmsol2.

Go to "Home" tab of Vmware WS

Click "Create a new Team"

=> "New Team Wizard" Click "Next"

=> Give Name: **Solaris 11gR2 RAC Team**

Location: <absolute path for the directory for the Team>

Click "Next"

=> "Add Virtual Machines to the Team"

Use "Add" button to choose in this **sequence** vmutil.vmx, vmsol1.vmx & vmsol2.vmx

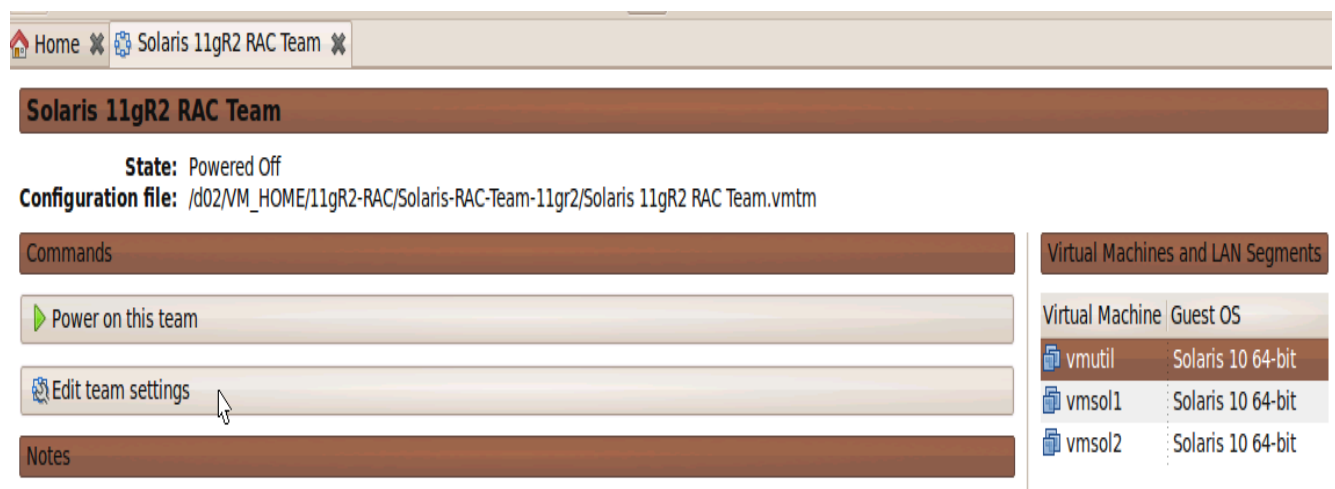
Click "Next"

=> "Add LAN Segments to the Team" [have not added any]

Click "Next"

=> Click "Finish"

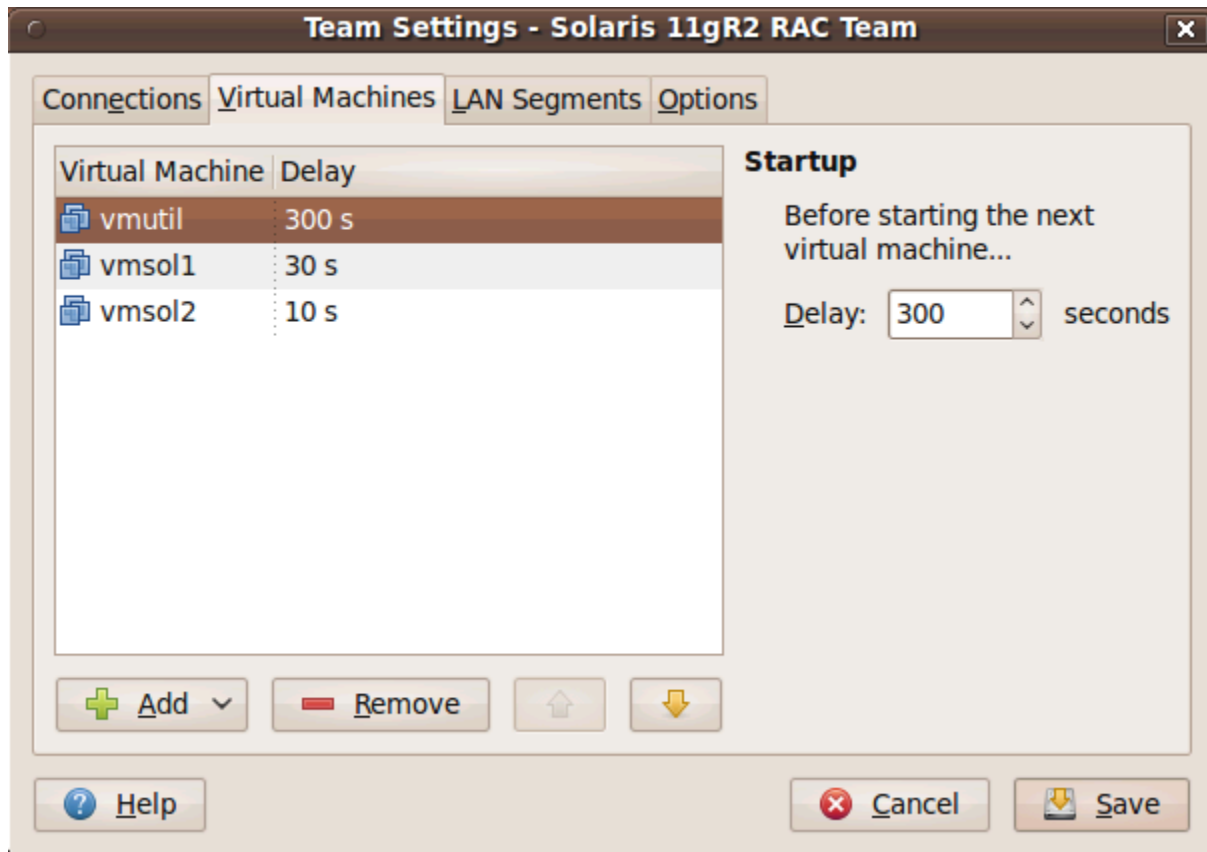
You should see like in below image [you will not see tab's for vmutil, vmsol1 & vmsol2]



Now let's alter the default time (secs) for startup of the 3 VM's to cater our boot sequence. We would first start vmutil, then after 300 secs we will start vmsol1 and after 30 secs vmsol2.

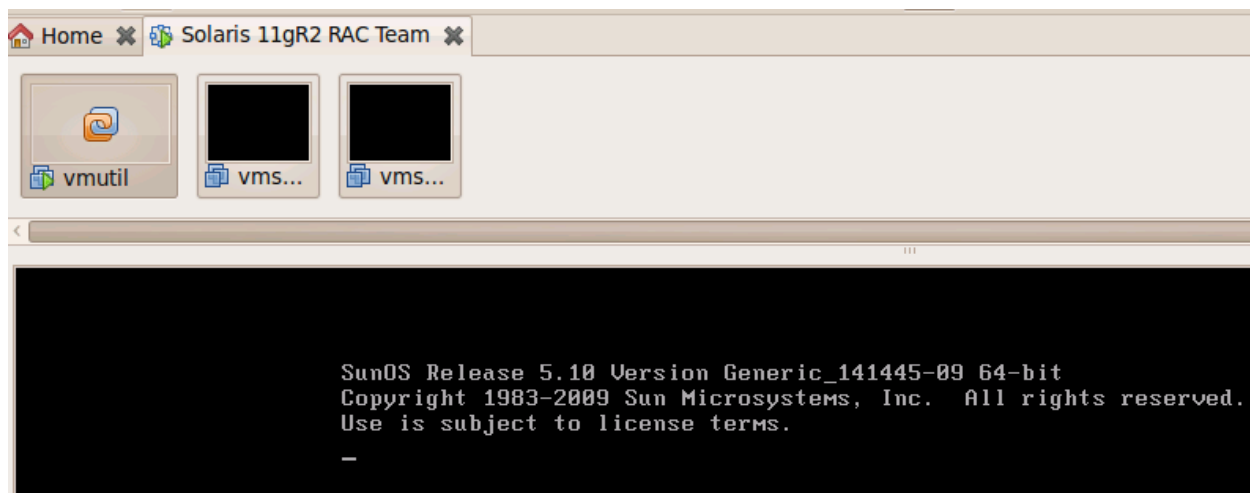
In the “Solaris 11gR2 RAC Team” tab
Click “Edit team Settings”

=> Click on “Virtual Machines” Tab and change the **Delay:** as shown below in image.



Click “Save”

From “Solaris 11gR2 RAC Team” tab, Use “Power on this team” Start the RAC database *racdb*.



19. Issues & Trouble Shooting & Learnings

VM kernel Panic during DBCA:

You may experience kernel panic during rac database creation during dbca (high CPU utilization by java processes) which will result in Automatic VM reboots. So donot allocate all CPU (cores) to a VM as mentioned during the vmutil creation.

Snapshot Mangement & Restore

It is good idea to give a naming convention for Snapshot along with meaningful descriptions of the contents of the snap. So when ever you plan to revert to a snap after step 12 there is dependency of shared storage on all 3 nodes, so ensure whenever you revert all are in sync. So best way to do it is to take a snap after shutting down all the 3 nodes and give it a same snap name (eg: Sync-All-Snap1) so it would be easy to revert. With Vmware WS you can also have child snaps which can be a snap(hot snap) of a running system. So after step 7, try to take snaps of all the three nodes together and when ever you are restoring them restore not just snap on one node but snap of same level on all nodes.

References

- Oracle Documentation.
- Documents on Sun.com [ZFS related]
- Vmware Workstation 7 Documentation

Conclusion

I hope this runbook for 11gR2 RAC on Solaris would help in building your own desktop version of RAC.