

Creating shared storage on SUSE Linux Enterprise Server 10 using Xen and OCFS2

Arun Singh
arun.singh@novell.com
Novell Inc.

1. Introduction

This paper is to describe the process of creating shared storage on a standalone system using SUSE Linux Enterprise Server 10 (SLES10) i.e. OCFS2 and Xen Virtualization. Cool part is there is no real expensive shared storage is used. Information provided here works with real shared storage as well.

2. Steps

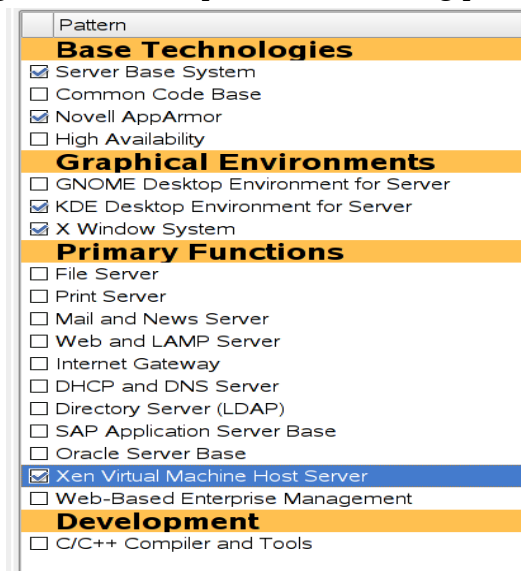
1. Disk Partition Layout

Creating partition for shared storage is important. Please plan ahead to avoid re-install of SLES10. Here is sample disk layout from my server, to create shared storage i.e. /dev/sda3.

/dev/sda3	10.0 GB	Linux native	/shared	K	3708	5013
/dev/sda4	36.1 GB	Extended			5014	9728
/dev/sda5	2.0 GB	Linux swap	swap	K	5014	5275
/dev/sda6	34.1 GB	Linux native	/	K	5276	9728

2. Install SLES10 with Xen

Follow the standard SUSE Linux Enterprise Server install. Select “**Xen Virtual Machine Host Server**” pattern. Also search and select “**ocfs2-tools & ocfs2console**” to install OCFS2 packages. Here is sample screen, showing package selection.



3. Boot SLES10 Server with XEN kernel

Once SLES10 with Xen is installed successfully, reboot the server and select XEN boot option.



Verify that server is booted with xen kernel.

```
sles10:~ # uname -a
Linux sles10 2.6.16.21-0.8-xen #1 SMP Mon Jul 3 18:25:39 UTC 2006 i686 i686 i386 GNU/Linux
sles10:~ #
sles10:~ # rpm -qa |grep ocfs2
ocfs2-tools-1.2.1-5.6
ocfs2console-1.2.1-5.6
sles10:~ #
```

4. Upgrade XEN tools on SLES10 Server (Optional)

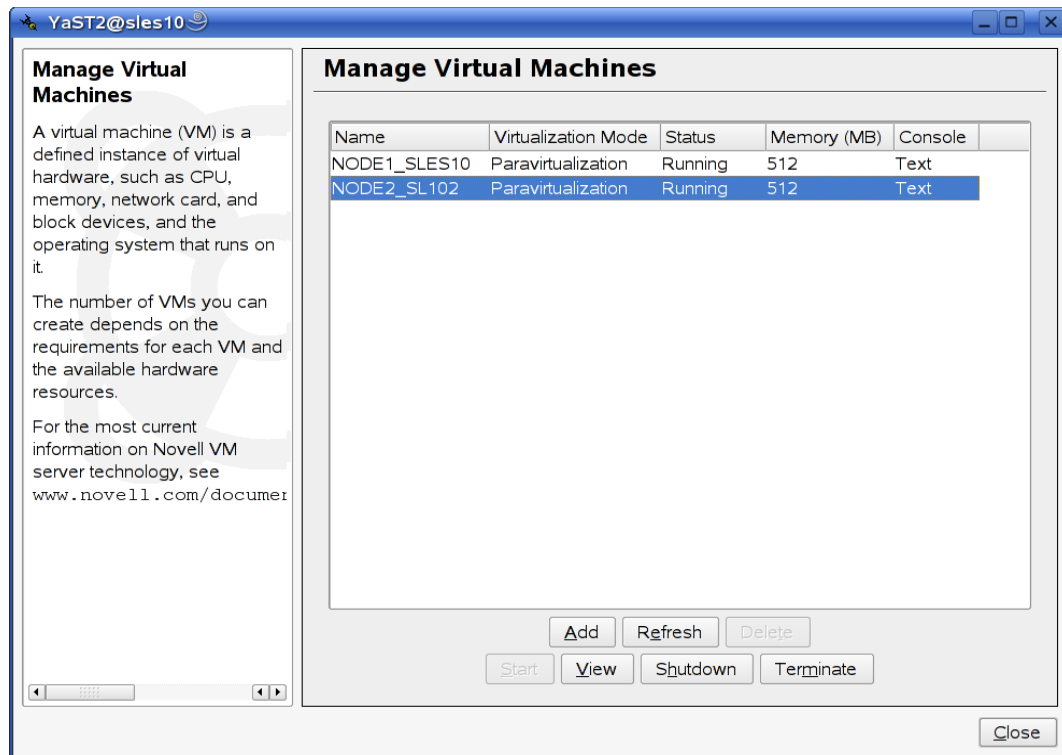
This is an optional step but an important one to upgrade xen tools to newer available version. Visit Novell support web site (<http://support.novell.com/linux/psdb/>) to download newer version of xen tools. You will find installing xen guest very slow without these updated xen tools. Here is screen shot showing upgrade procedure.

```
sles10:/xen_update # ls -al
total 5265
drwxr-xr-x  2 root root   368 Dec  3 14:04 .
drwxr-xr-x 25 root root   600 Dec  3 14:01 ..
-rw-r--r--  1 root root 4230253 Dec  3 14:02 xen-3.0.2_09763-0.8.i586.rpm
-rw-r--r--  1 root root 137476 Dec  3 14:02 xen-devel-3.0.2_09763-0.8.i586.rpm
-rw-r--r--  1 root root  85683 Dec  3 14:02 xen-libs-3.0.2_09763-0.8.i586.rpm
-rw-r--r--  1 root root 502110 Dec  3 14:02 xen-tools-3.0.2_09763-0.8.i586.rpm
-rw-r--r--  1 root root 306892 Dec  3 14:03 xen-tools-ioemu-3.0.2_09763-0.8.i586.rpm
-rw-r--r--  1 root root 115027 Dec  3 14:03 yast2-vm-2.13.62-4.2.i586.rpm
sles10:/xen_update # rpm -Fvh xen-* yast2-vm-2.13.62-4.2.i586.rpm
Preparing...
 1:xen
 2:xen-libs
 3:xen-tools
Updating etc/sysconfig/xendomains...
 4:xen-tools-ioemu
 5:yast2-vm
sles10:/xen_update #
```

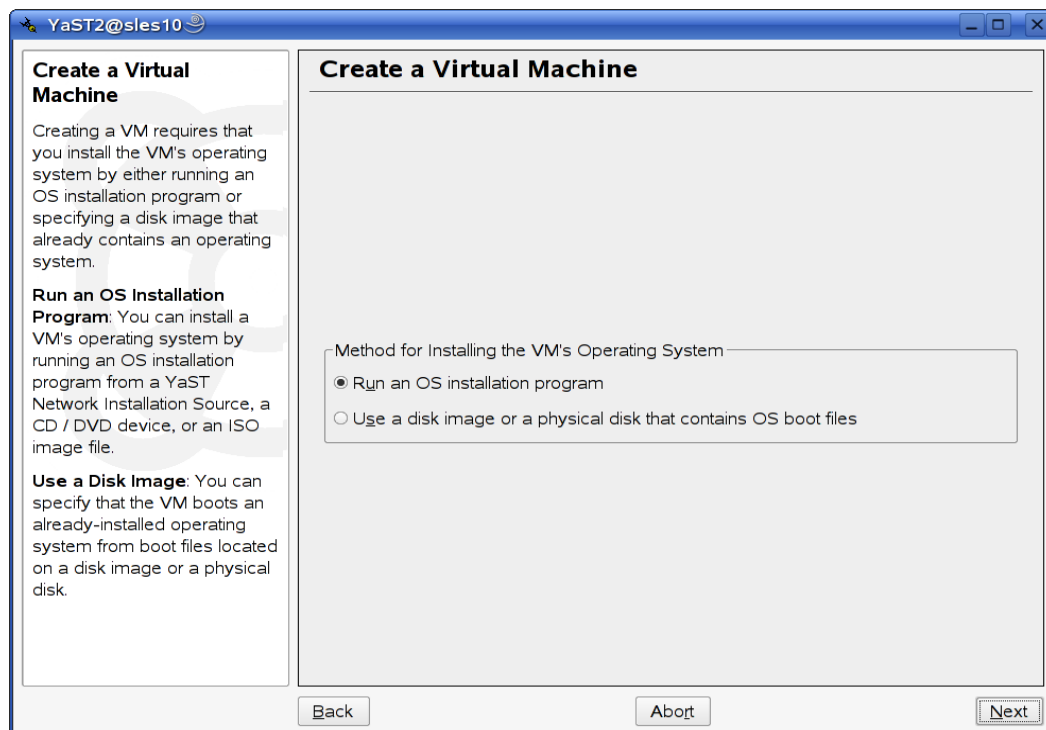
5. Install Xen Guests OS

Installing xen guset OS is very easy in SLES10. There are many documents available explaining installation procedure, but here are some quick steps.

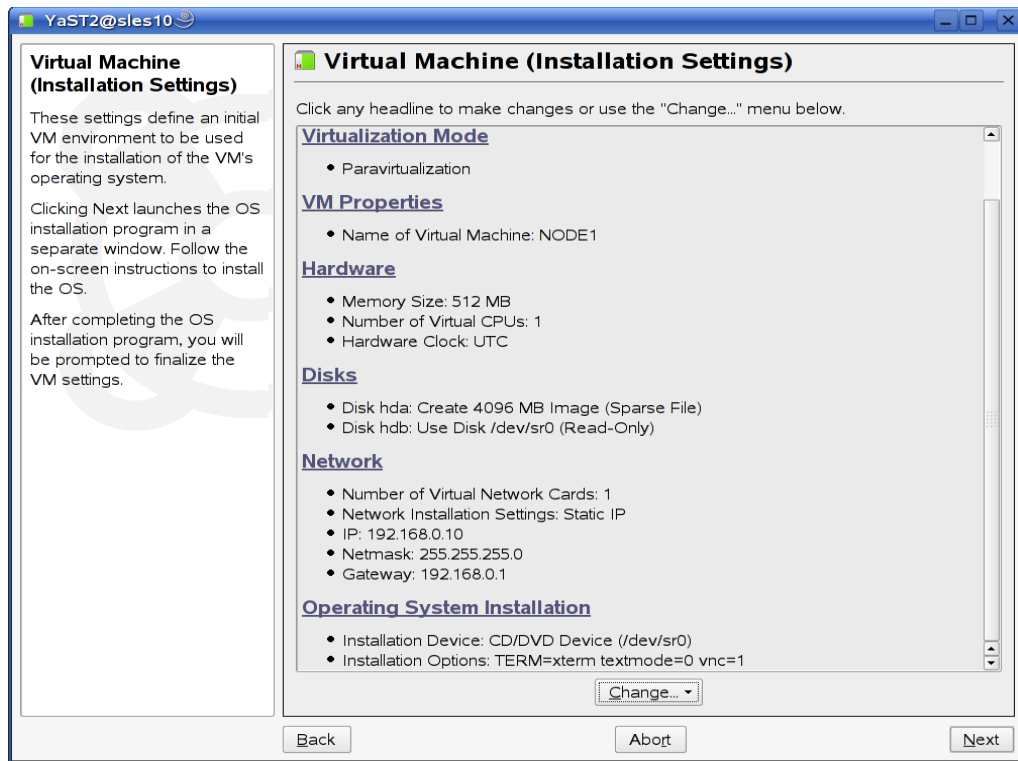
1. Start YaST xen console i.e. /sbin/yast2 xen



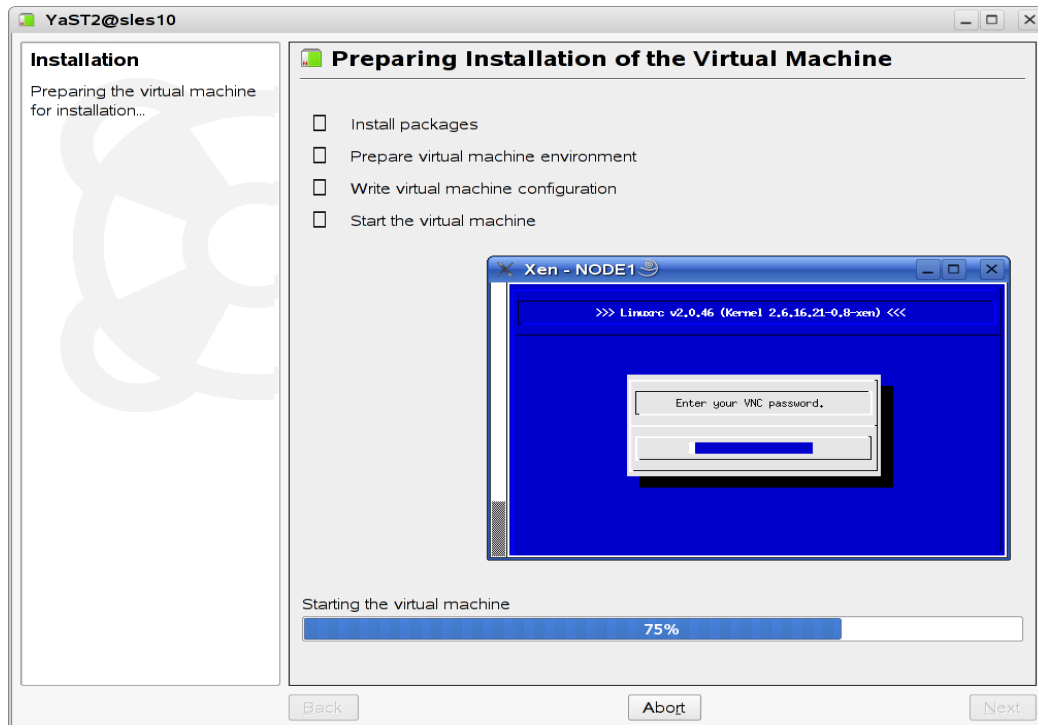
2. Click "Add" to add new virtual machine.



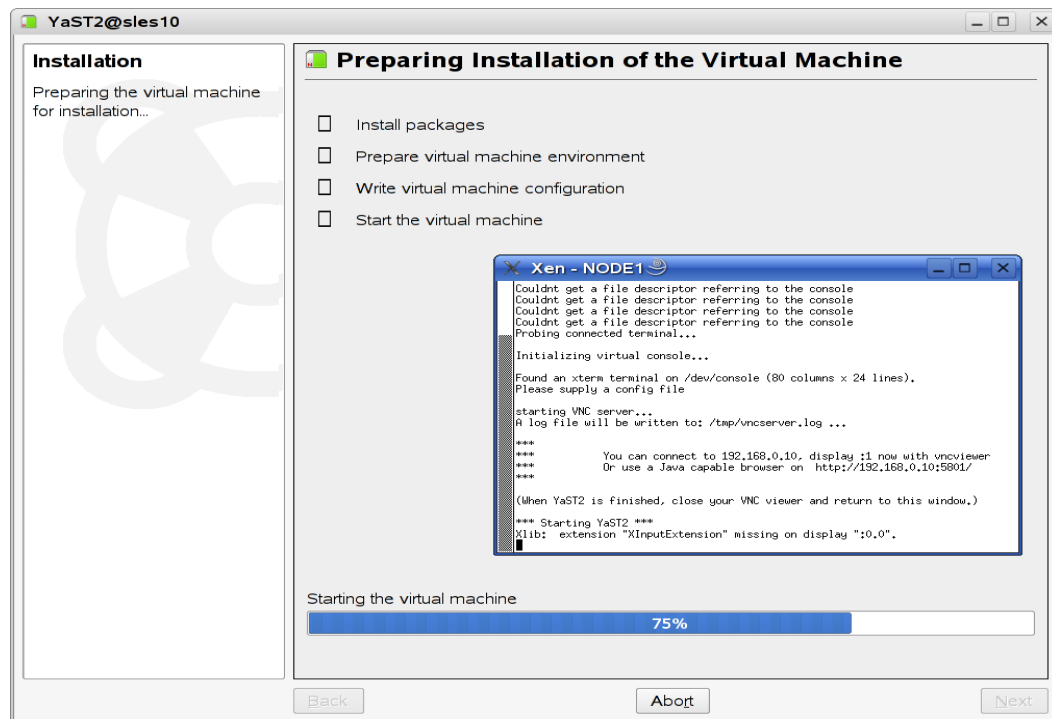
3. Change Settings, especially **Network** and **Operating Sysystem Installation**.



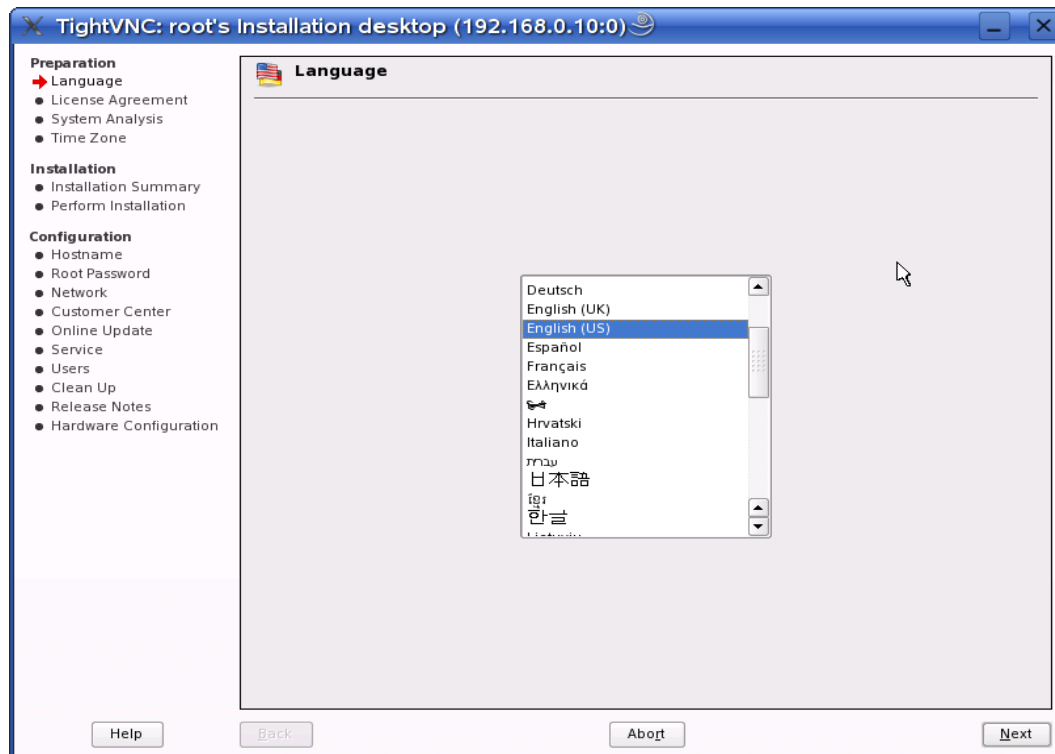
4. Here vnc based installation is selected. You may get different screen for text based install.



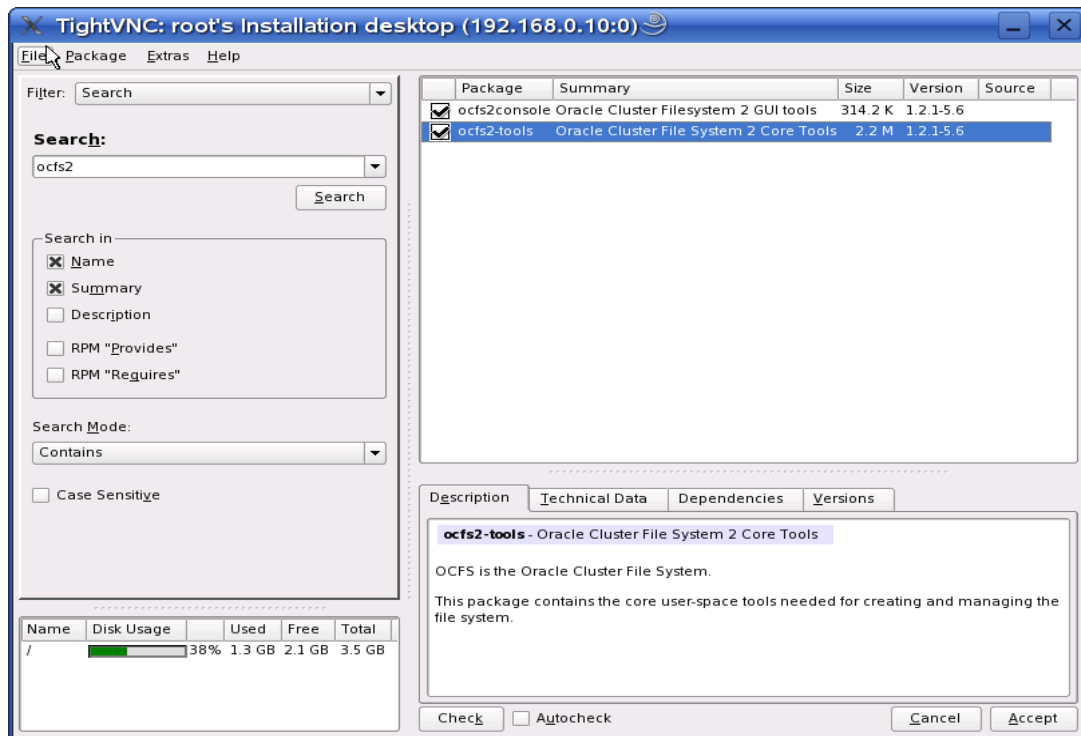
5. Follow the install instruction presnted i.e. open vncviewer and continue installation.




6. Here on-ward it's standard install.



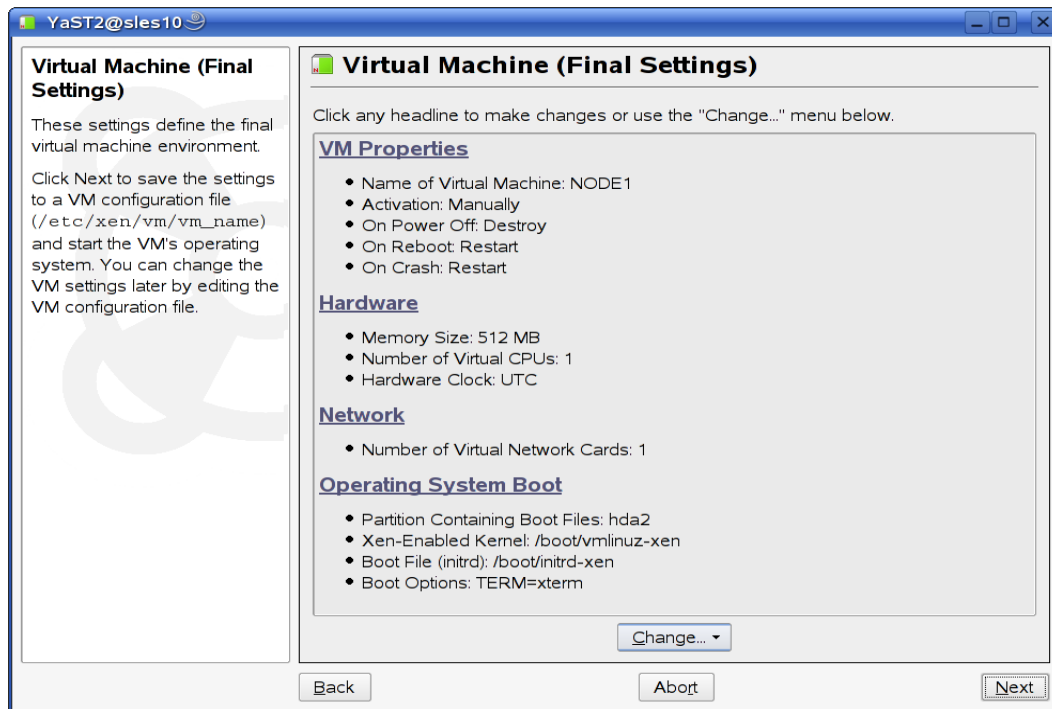
7. Select default package selection. You can search and select additional ocfs2 packages.



Installation continues ...



Now guest OS Installation is complete. Change any boot settings needed.



Once virtual guest OS is installed, you can stop/start from same YaST console. If you like, you can use command line options i.e. **xm**.

6. Configure OCFS2 on Xen Host Server

Run **o2cb** command to configure and enable cluster.

```
Shell - Konsole
Session Edit View Bookmarks Settings Help
sles10:/ # /etc/init.d/o2cb configure
Configuring the O2CB driver.

This will configure the on-boot properties of the O2CB driver.
The following questions will determine whether the driver is loaded on
boot. The current values will be shown in brackets ('[]'). Hitting
<ENTER> without typing an answer will keep that current value. Ctrl-C
will abort.

Load O2CB driver on boot (y/n) [y]:
Cluster to start on boot (Enter "none" to clear) [ocfs2]:
Use user-space driven heartbeat? (y/n) [n]:
Writing O2CB configuration: OK
Loading module "configfs": OK
Mounting configfs filesystem at /sys/kernel/config: OK
Loading module "ocfs2_nodemanager": OK
Loading module "ocfs2_dlm": OK
Loading module "ocfs2_dlmfs": OK
Mounting ocfs2_dlmfs filesystem at /dlm: OK
Starting cluster ocfs2: OK
sles10:/ # /etc/init.d/o2cb enable
Writing O2CB configuration: OK
Cluster ocfs2 already online
sles10:/ #
```

7. Format shared volume partition

In order to use partition as OCFS2 cluster, you need to format partition using ocfs2 tool i.e. mkfs.ocfs2. This tool is added to your system when ocfs2-tools package is installed. Alternatively, you can use ocfs2console, GUI application to format partition. Here is screen shot showing the formatting of partition /dev/sda3.


First check o2cb clutrsre is online.

```
sles10:~ # chkconfig ocfs2
ocfs2 on
sles10:~ # chkconfig o2cb
o2cb on
sles10:~ # /etc/init.d/o2cb status
Module "configfs": Loaded
Filesystem "configfs": Mounted
Module "ocfs2_nodemanager": Loaded
Module "ocfs2_dlm": Loaded
Module "ocfs2_dlmfs": Loaded
Filesystem "ocfs2_dlmfs": Mounted
Checking cluster ocfs2: Online
Checking heartbeat: Not active
sles10:~ # █
```

Then either use ocfs2console or mkfs.ocfs2 to format partition.

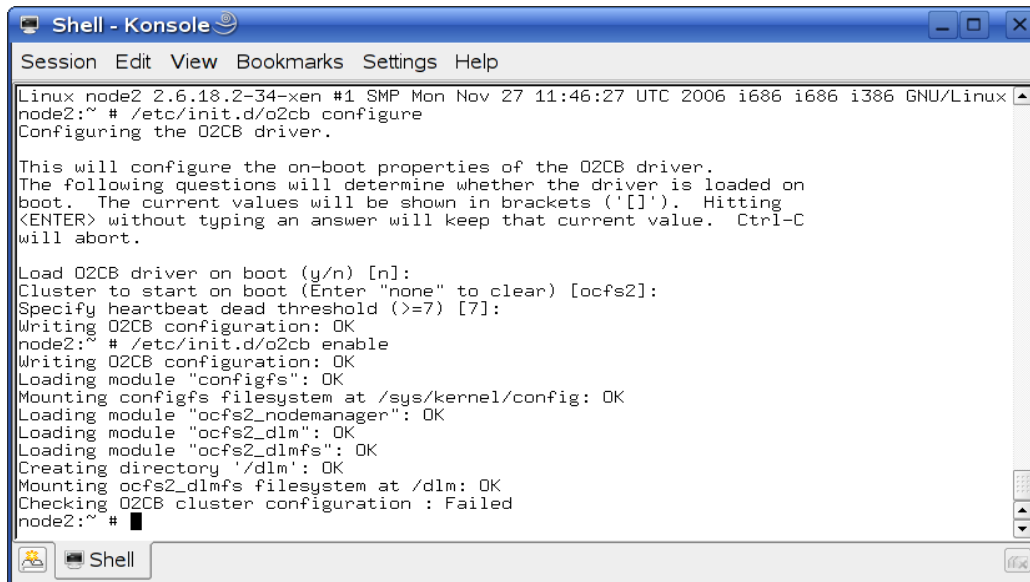
```
sles10:/ # mkfs.ocfs2 -L xen_cluster /dev/sda3
mkfs.ocfs2 1.2.1
Overwriting existing ocfs2 partition.
Proceed (y/N): y
Filesystem label=xen_cluster
Block size=4096 (bits=12)
Cluster size=4096 (bits=12)
Volume size=10742214656 (2622611 clusters) (2622611 blocks)
82 cluster groups (tail covers 9875 clusters, rest cover 32256 clusters)
Journal size=67137536
Initial number of node slots: 4
Creating bitmaps: done
Initializing superblock: done
Writing system files: done
Writing superblock: done
Formatting Journals: done
Writing lost+found: done
mkfs.ocfs2 successful
sles10:/ #
```

Here is ocfs2console view after formatting.



8. Configure OCFS2 Cluster on Xen Guest OS's

You need to configure OCFS2 on each participating xen guest to allow it to mount shared storage. Install `ocfs2-tools` and `ocfs2console`, If you haven't done during xen guest OS installation. Here is screen shot of OCFS2 configuration of Xen Guest OS.



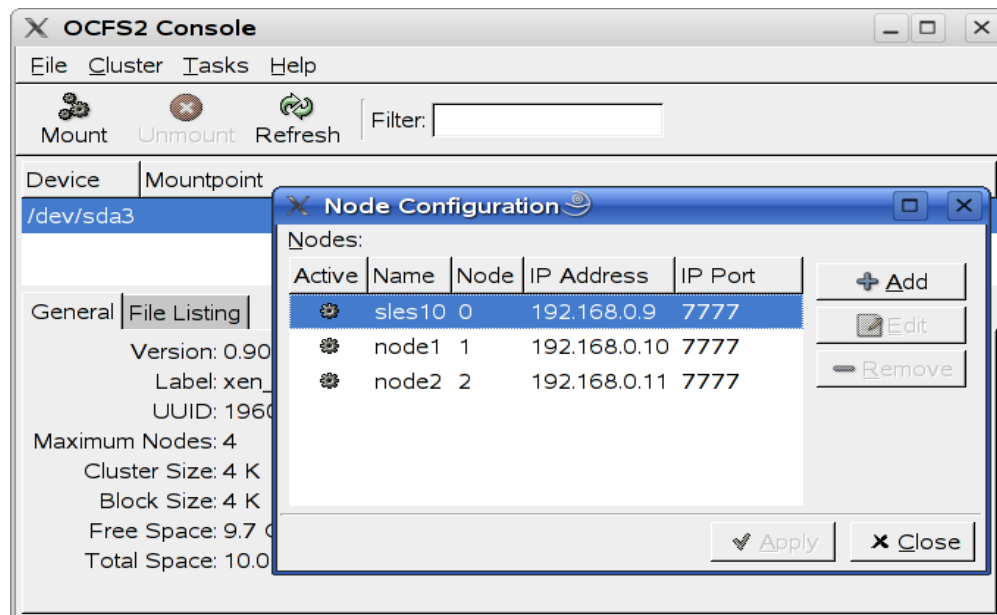
```
Shell - Konsole
Session Edit View Bookmarks Settings Help
Linux node2 2.6.18.2-34-xen #1 SMP Mon Nov 27 11:46:27 UTC 2006 i686 i686 i386 GNU/Linux
node2:~ # /etc/init.d/o2cb configure
Configuring the O2CB driver.

This will configure the on-boot properties of the O2CB driver.
The following questions will determine whether the driver is loaded on
boot. The current values will be shown in brackets ('[]'). Hitting
<ENTER> without typing an answer will keep that current value. Ctrl-C
will abort.

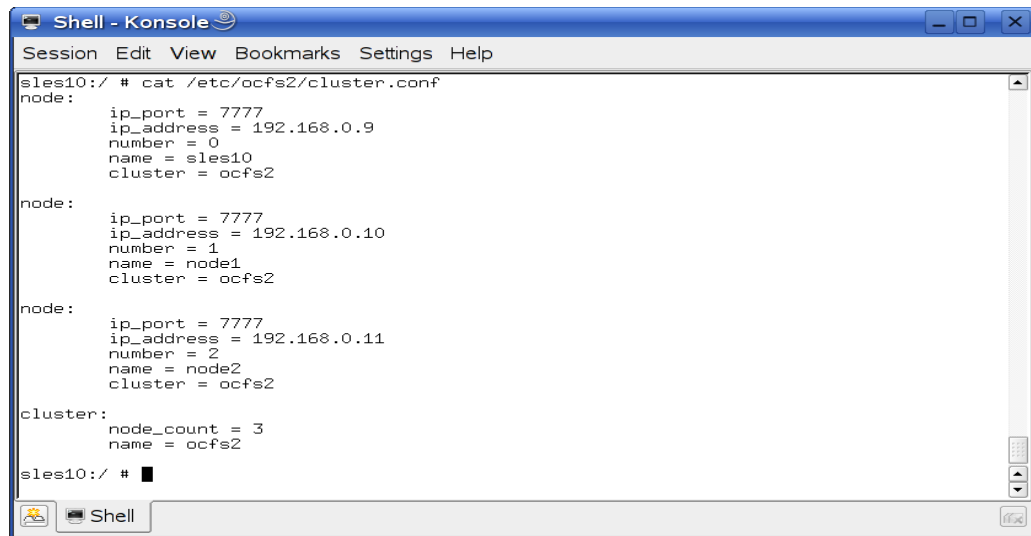
Load O2CB driver on boot (y/n) [n]:
Cluster to start on boot (Enter "none" to clear) [ocfs2]:
Specify heartbeat dead threshold (>=7) [7]:
Writing O2CB configuration: OK
node2:~ # /etc/init.d/o2cb enable
Writing O2CB configuration: OK
Loading module "configfs": OK
Mounting configfs filesystem at /sys/kernel/config: OK
Loading module "ocfs2_nodemanager": OK
Loading module "ocfs2_dlm": OK
Loading module "ocfs2_dlmfs": OK
Creating directory '/dlm': OK
Mounting ocfs2_dlmfs filesystem at /dlm: OK
Checking O2CB cluster configuration : Failed
node2:~ #
```

9. Add OCFS2 nodes to cluster

Once all xen guest are ocfs2 configured and enabled, you need to populate cluster with node information. You can use `ocfs2console` utility to add new node cluster information and propagate to all other nodes in the cluster. Make sure all nodes in the cluster have same `/etc/ocfs2/cluster.conf` information.



Here is /etc/ocfs2/cluster.conf file. Note Static IP address of cluster nodes.



```
Shell - Konsole
Session Edit View Bookmarks Settings Help
sles10:/ # cat /etc/ocfs2/cluster.conf
node:
    ip_port = 7777
    ip_address = 192.168.0.9
    number = 0
    name = sles10
    cluster = ocfs2

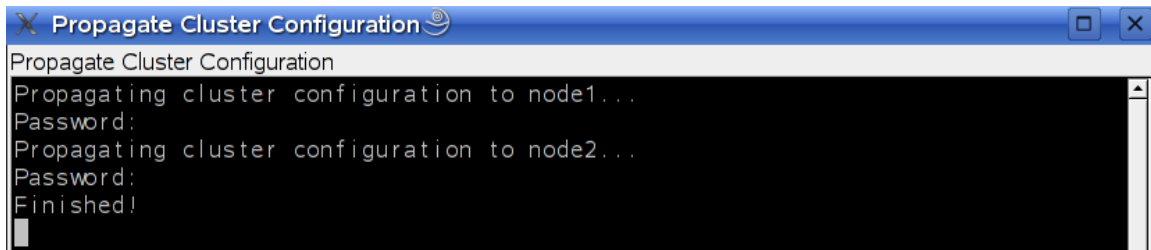
node:
    ip_port = 7777
    ip_address = 192.168.0.10
    number = 1
    name = node1
    cluster = ocfs2

node:
    ip_port = 7777
    ip_address = 192.168.0.11
    number = 2
    name = node2
    cluster = ocfs2

cluster:
    node_count = 3
    name = ocfs2

sles10:/ #
```

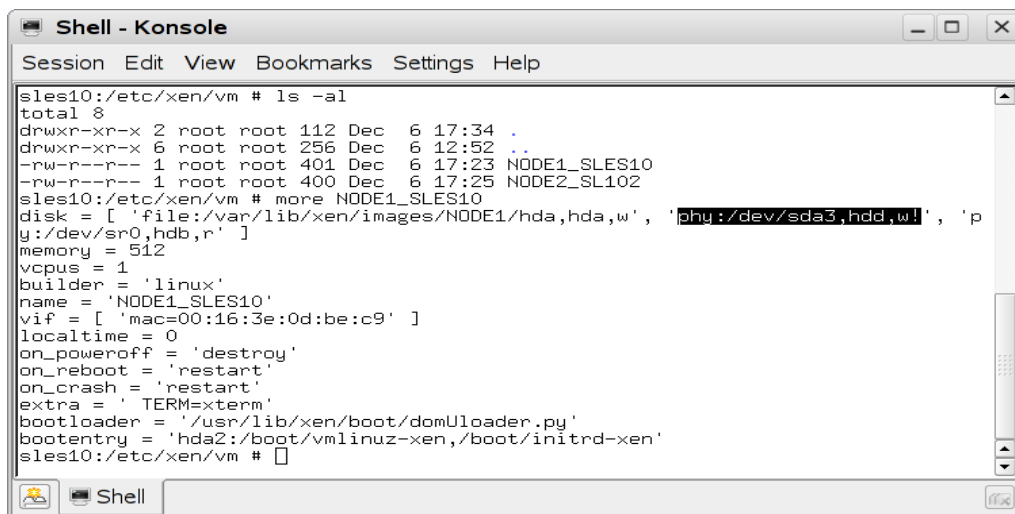
You can copy /etc/ocfs2/cluster.conf manually to all nodes or use **Propagate Clustre Configuration** feature of ocfs2console GUI tool.



```
Propagate Cluster Configuration
Propagating cluster configuration to node1...
Password:
Propagating cluster configuration to node2...
Password:
Finished!
```

10. Add shared partition to Xen Guest OS

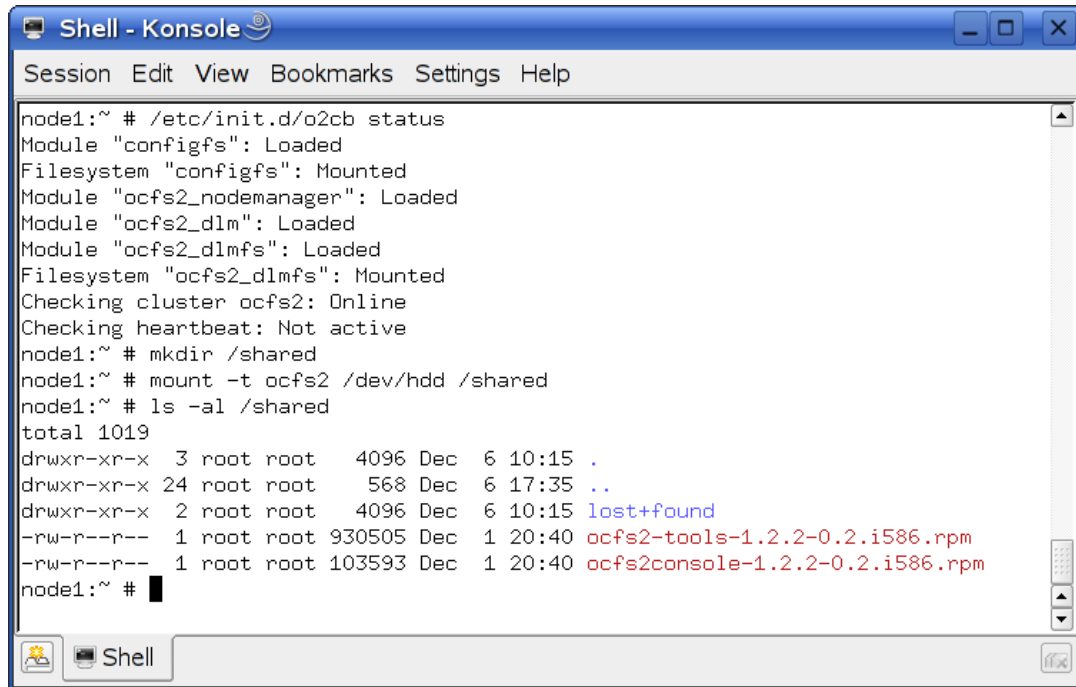
Add shared partition /dev/sda3 detail to xen guest OS setting in /etc/xen/vm/. Here is screen shot showing the detail. Reboot xen guest OS to take change into effect.



```
Shell - Konsole
Session Edit View Bookmarks Settings Help
sles10:/etc/xen/vm # ls -al
total 8
drwxr-xr-x 2 root root 112 Dec  6 17:34 .
drwxr-xr-x 6 root root 256 Dec  6 12:52 ..
-rw-r--r-- 1 root root 401 Dec  6 17:23 NODE1_SLES10
-rw-r--r-- 1 root root 400 Dec  6 17:25 NODE2_SL102
sles10:/etc/xen/vm # more NODE1_SLES10
disk = [ 'file:/var/lib/xen/images/NODE1/hda,hda,w', 'phy:/dev/sda3,hdd,w', 'p
y:/dev/sr0,hdb,r' ]
memory = 512
vcpus = 1
builder = 'linux'
name = 'NODE1_SLES10'
vif = [ 'mac=00:16:3e:0d:be:c9' ]
localtime = 0
on_poweroff = 'destroy'
on_reboot = 'restart'
on_crash = 'restart'
extra = ' TERM=xterm'
bootloader = '/usr/lib/xen/boot/domUloader.py'
bootentry = 'hda2:/boot/vmlinuz-xen,/boot/initrd-xen'
sles10:/etc/xen/vm #
```

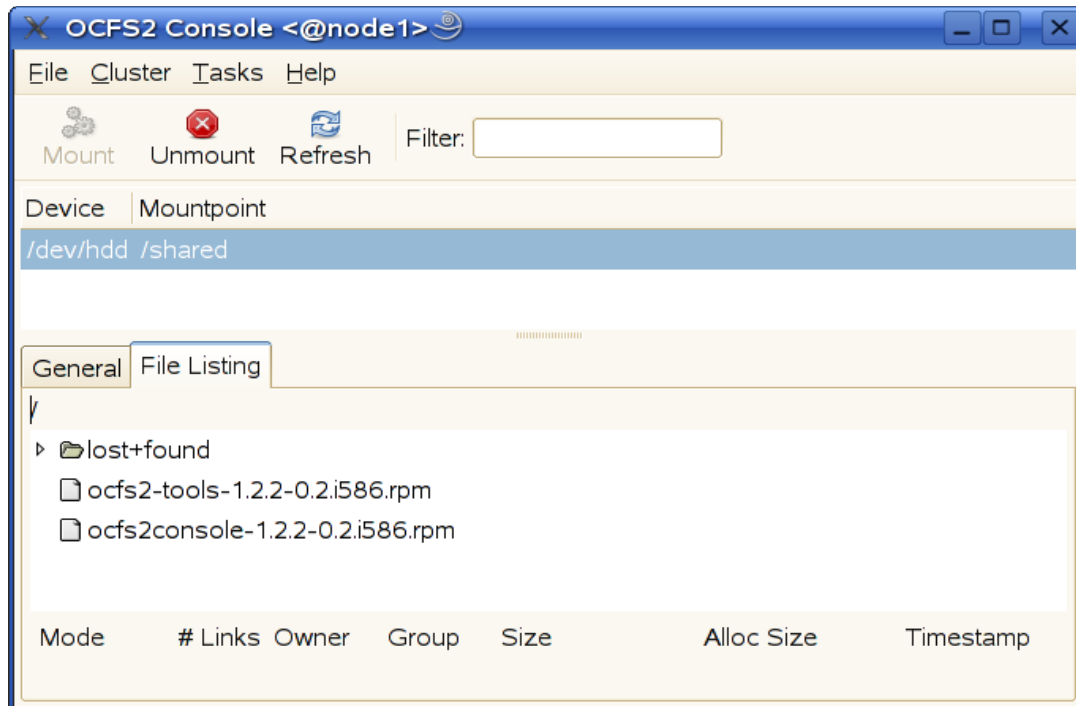
11. Mount shared partition

Now you are ready to mount and use newly created shared partition (/dev/sda3) in Xen guest OS. Here is screen shot showing how to add physical partition on xen guest OS.

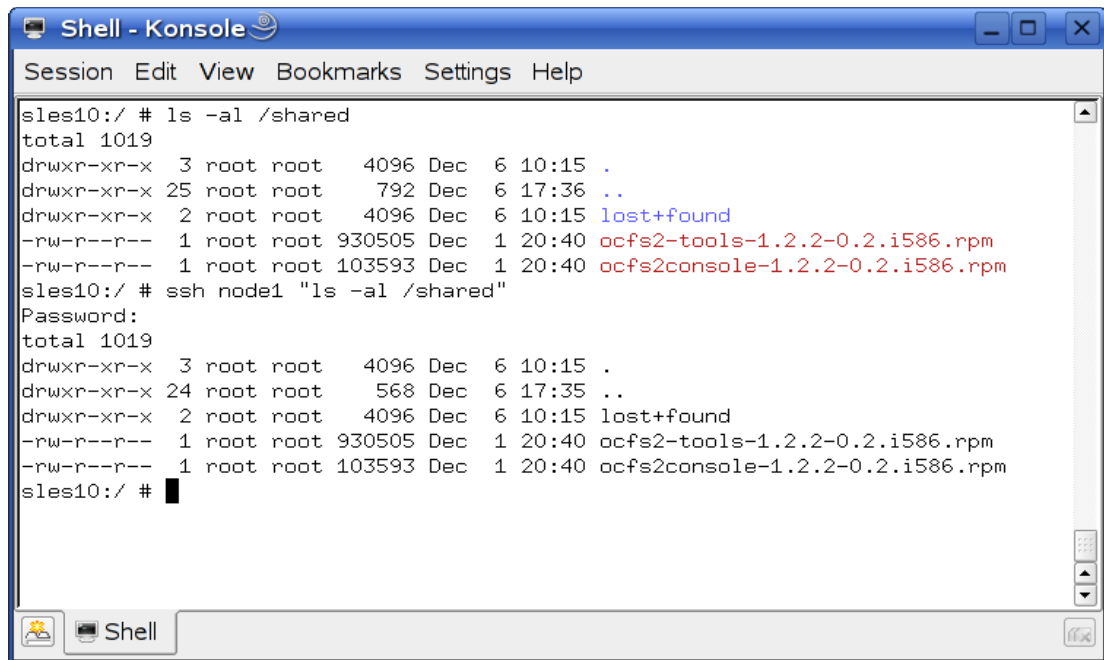


```
node1:~ # /etc/init.d/o2cb status
Module "configfs": Loaded
Filesystem "configfs": Mounted
Module "ocfs2_nodemanager": Loaded
Module "ocfs2_dlm": Loaded
Module "ocfs2_dlmfs": Loaded
Filesystem "ocfs2_dlmfs": Mounted
Checking cluster ocfs2: Online
Checking heartbeat: Not active
node1:~ # mkdir /shared
node1:~ # mount -t ocfs2 /dev/hdd /shared
node1:~ # ls -al /shared
total 1019
drwxr-xr-x  3 root root   4096 Dec  6 10:15 .
drwxr-xr-x 24 root root    568 Dec  6 17:35 ..
drwxr-xr-x  2 root root   4096 Dec  6 10:15 lost+found
-rw-r--r--  1 root root 930505 Dec  1 20:40 ocfs2-tools-1.2.2-0.2.i586.rpm
-rw-r--r--  1 root root 103593 Dec  1 20:40 ocfs2console-1.2.2-0.2.i586.rpm
node1:~ #
```

ocfs2console showing mounted /shared volume on node1.



Here is screen shot showing /shared mounted on Xen host **sles10** and Xen guset **node1**.



```
Shell - Konsole
Session Edit View Bookmarks Settings Help

sles10:/ # ls -al /shared
total 1019
drwxr-xr-x  3 root root   4096 Dec  6 10:15 .
drwxr-xr-x 25 root root    792 Dec  6 17:36 ..
drwxr-xr-x  2 root root   4096 Dec  6 10:15 lost+found
-rw-r--r--  1 root root 930505 Dec  1 20:40 ocfs2-tools-1.2.2-0.2.i586.rpm
-rw-r--r--  1 root root 103593 Dec  1 20:40 ocfs2console-1.2.2-0.2.i586.rpm
sles10:/ # ssh node1 "ls -al /shared"
Password:
total 1019
drwxr-xr-x  3 root root   4096 Dec  6 10:15 .
drwxr-xr-x 24 root root    568 Dec  6 17:35 ..
drwxr-xr-x  2 root root   4096 Dec  6 10:15 lost+found
-rw-r--r--  1 root root 930505 Dec  1 20:40 ocfs2-tools-1.2.2-0.2.i586.rpm
-rw-r--r--  1 root root 103593 Dec  1 20:40 ocfs2console-1.2.2-0.2.i586.rpm
sles10:/ #
```

3. Conclusion

This paper is to help you to understand steps involved in creating shared storage without using expensive shared storage. Using this information you can create shared storage used by all xen guest OS and Host, avoiding copying of files between guest OS's. Hope you will find this paper useful.