Nagios® XI™

High Availability

Prepared by – Vaibhav Daud

Table of Contents

Purpose

High Availability –

This is to ensure that the application is running without any business impact

Note:-

- Given steps are according to our requirement and setup, so might be we have not moved some components i.e. mrtg, but those can be move like this
- Use quorum disk only when you have fence device else leave the step and configure HA without quorum disk.

Prerequisites

- One physical system with OS & Oracle Virtual Box
- Two Linux VM with CentOS 6.5 for Nagios XI 2014 2.7
- One Linux VM for Openfiler (Virtual storage)
- Network connectivity with dns and internet.
- One Virtual IP and shared partitions
- Cluster packages – rgmanager, ricci, luci, gfs2-utils, iscsi

Base Plan

Create 3 VM's using a single physical system and virtualization application.

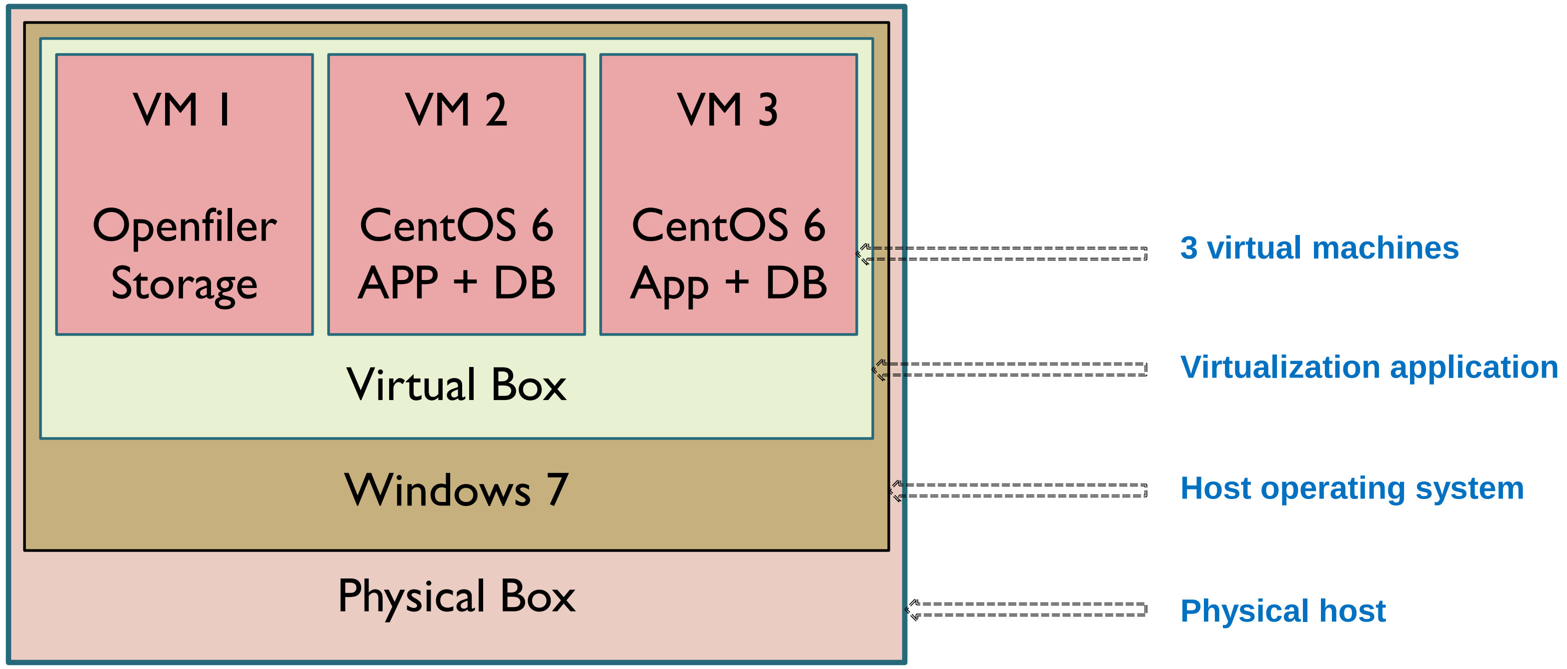

Base Architecture

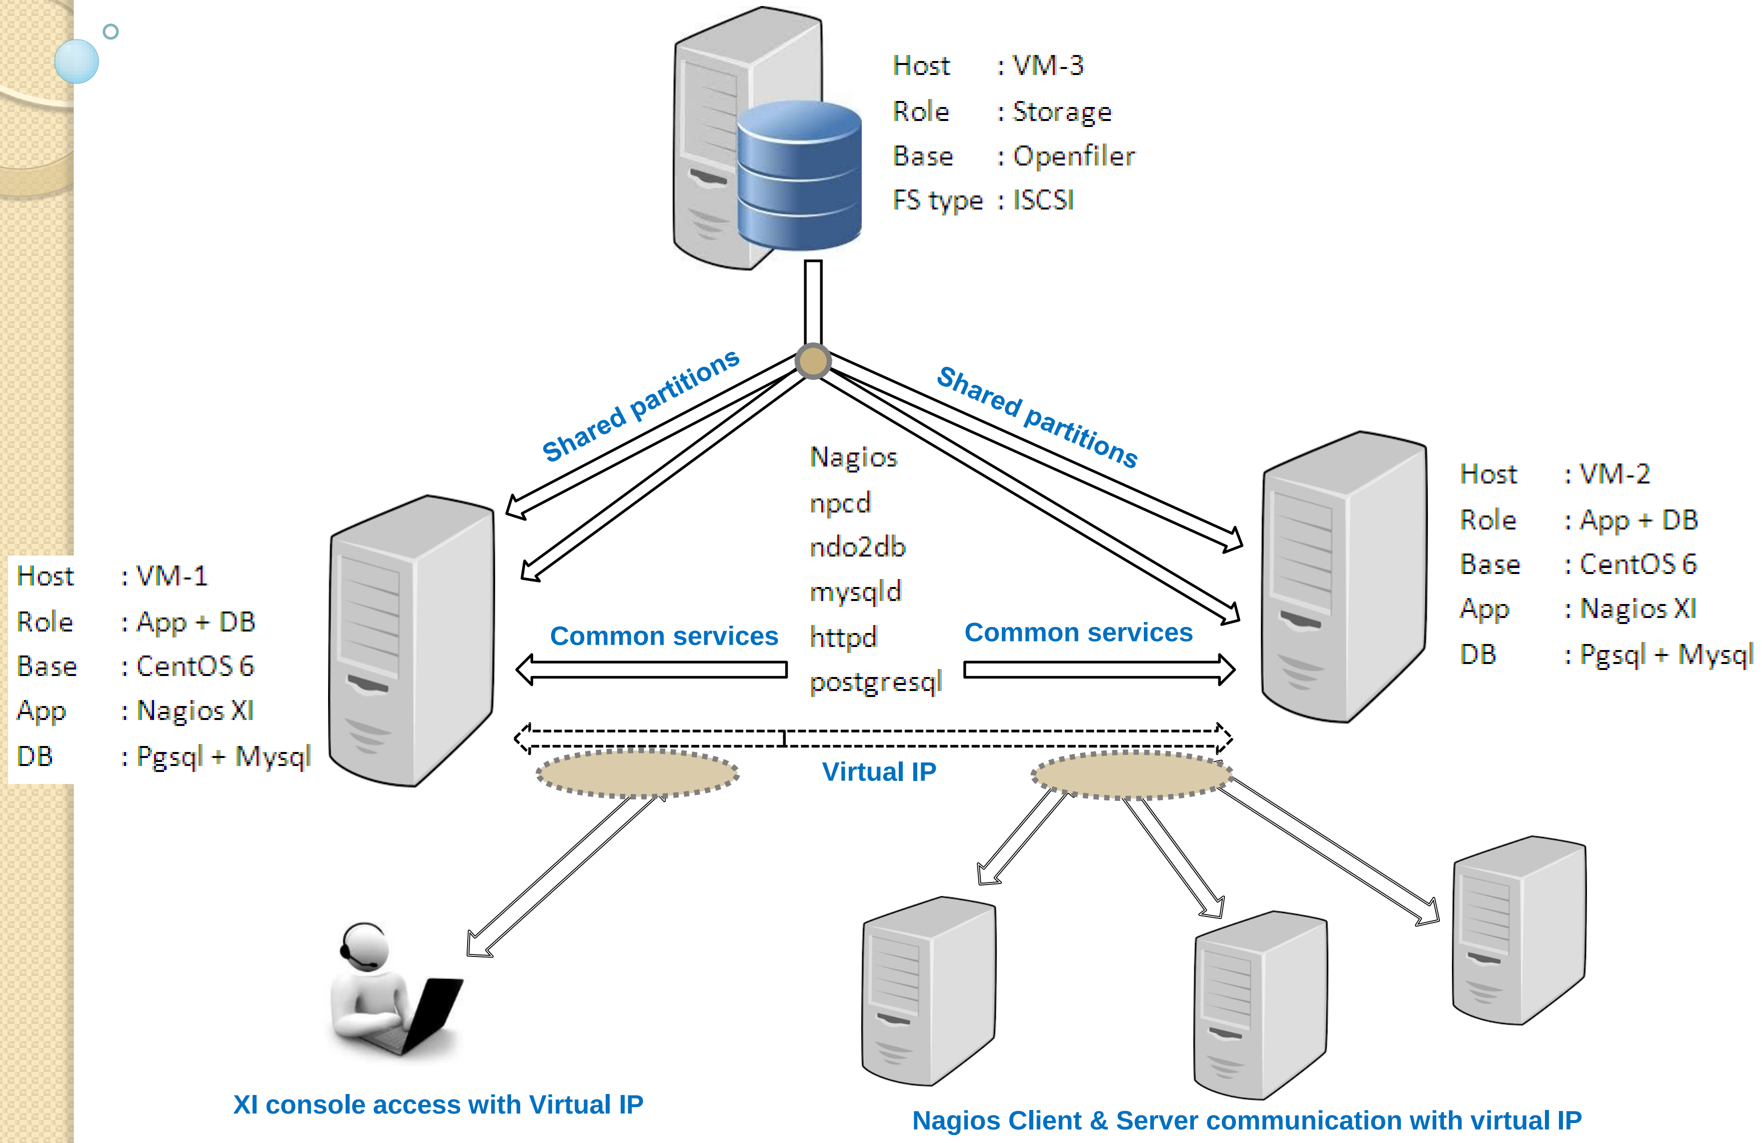

Overview

Here we are using one physical box with Intel core i5 processor, 12 gb ram, 500 gb hdd and Windows 7 OS. On which we have installed Virtual box for creating three virtual machines.

- VM-1 using as primary application server where Nagios XI, Mysql, Pgsql and its sub components are installed
- VM-2 using as secondary application server where Nagios XI, Mysql, Pgsql and its sub components are installed
- VM-3 using as storage where openfiler is installed, created two volumes and mapped to both XI servers as shared partitions.

After completion of OS & XI installation on both server, we installed the redhat cluster packages on both server, created a cluster, add both node into cluster, mapped shared partition to both server, created fs failover and IP failover. Then we moved required XI configuration files and db's from primary server to shared partition. Then created a script which will check status of each node and according to instruction will stop and start required services.

Configuration Steps

- Installed Virtual-box on base host and created CentOS VM as Nagios Primary

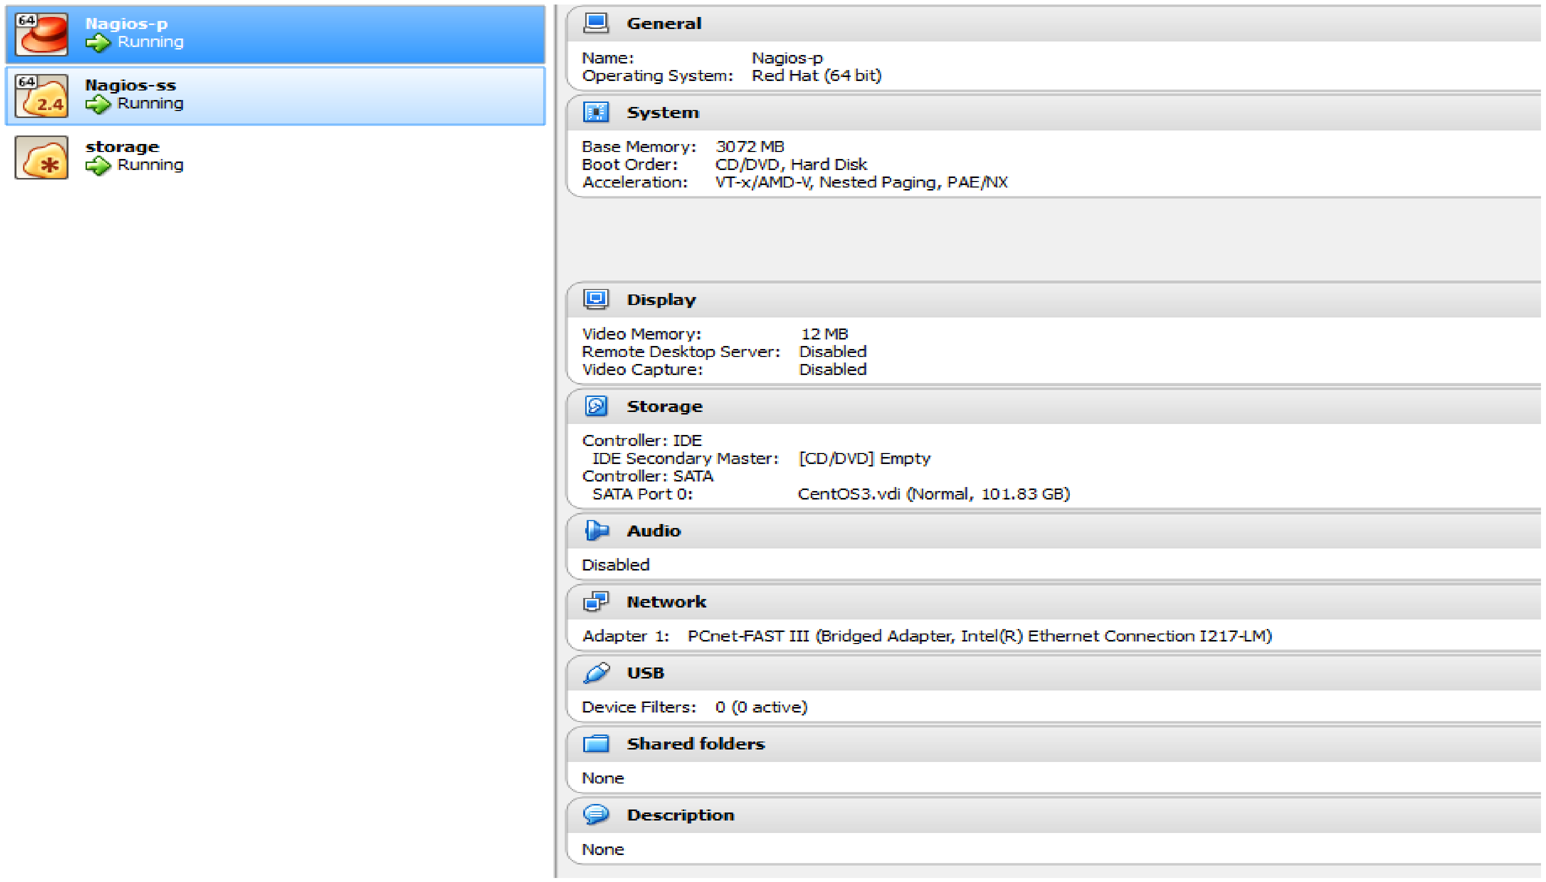

Nagios-p
Running
Nagios-ss
Running
storage
Running

General

Name: Nagios-p
Operating System: Red Hat (64 bit)

System

Base Memory: 3072 MB
Boot Order: CD/DVD, Hard Disk
Acceleration: VT-x/AMD-V, Nested Paging, PAE/NX

Display

Video Memory: 12 MB
Remote Desktop Server: Disabled
Video Capture: Disabled

Storage

Controller: IDE
IDE Secondary Master: [CD/DVD] Empty
Controller: SATA
SATA Port 0: CentOS3.vdi (Normal, 101.83 GB)

Audio

Disabled

Network

Adapter 1: PCnet-FAST III (Bridged Adapter, Intel(R) Ethernet Connection I217-LM)

USB

Device Filters: 0 (0 active)

Shared folders

None

Description

None

Configuration Steps

- Installed Virtual-box on base host and created CentOS VM as Nagios Secondary

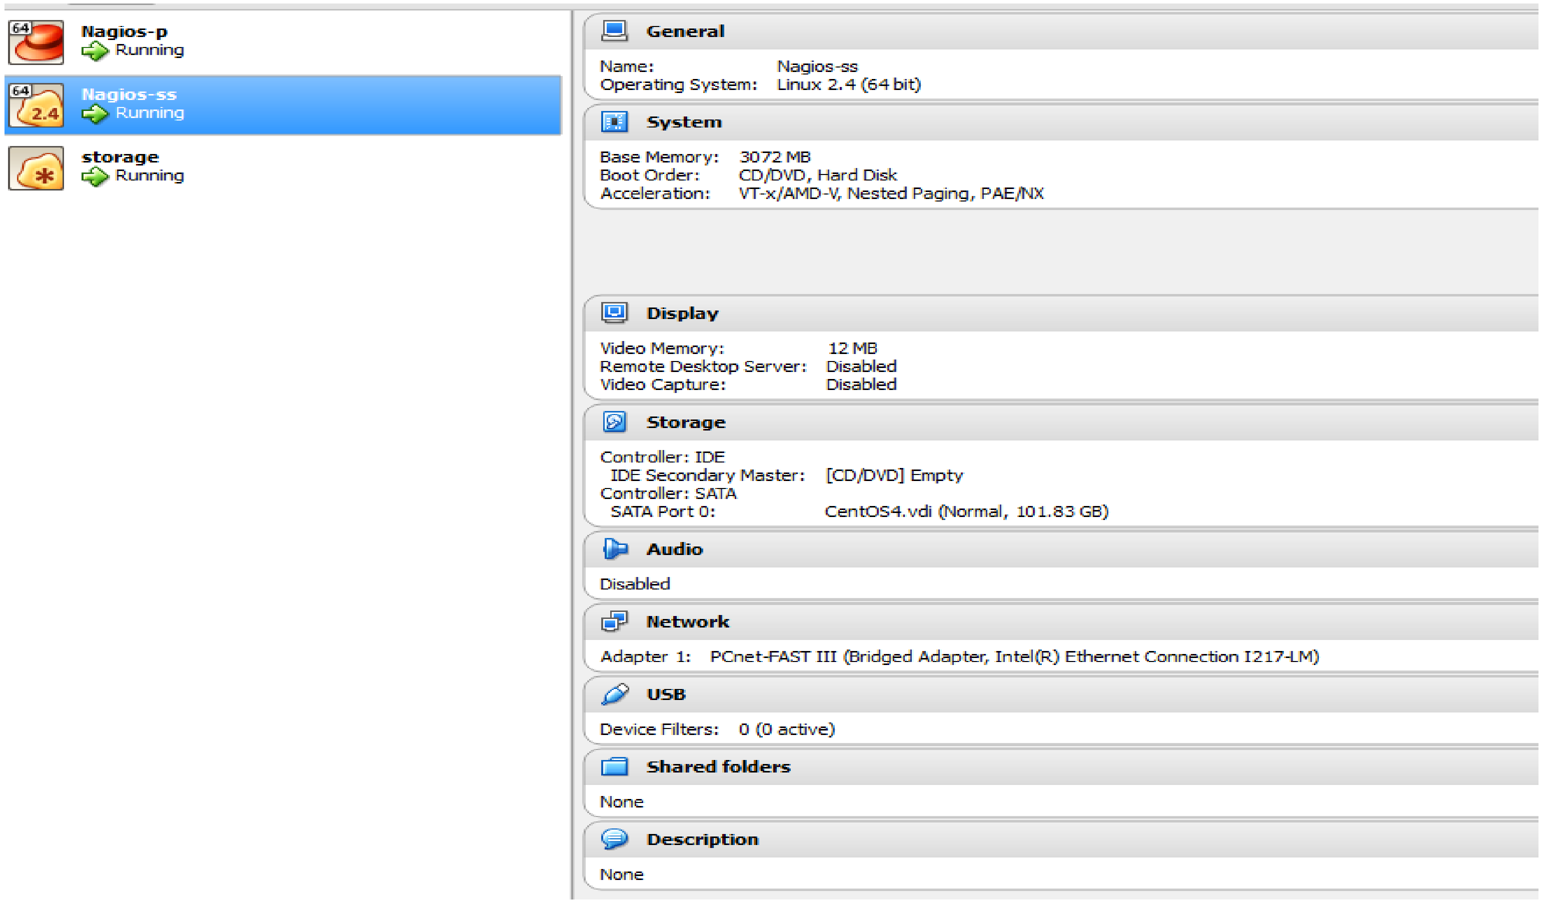

Nagios-p
Running

Nagios-ss
Running

storage
Running

General

Name:	Nagios-ss
Operating System:	Linux 2.4 (64 bit)

System

Base Memory:	3072 MB
Boot Order:	CD/DVD, Hard Disk
Acceleration:	VT-x/AMD-V, Nested Paging, PAE/NX

Display

Video Memory:	12 MB
Remote Desktop Server:	Disabled
Video Capture:	Disabled

Storage

Controller: IDE	
IDE Secondary Master:	[CD/DVD] Empty
Controller: SATA	
SATA Port 0:	CentOS4.vdi (Normal, 101.83 GB)

Audio

Disabled

Network

Adapter 1: PCnet-FAST III (Bridged Adapter, Intel(R) Ethernet Connection I217-LM)

USB

Device Filters: 0 (0 active)

Shared folders

None

Description

None

Configuration Steps

- Installed Virtual-box on base host and created Openfiler VM as Storage

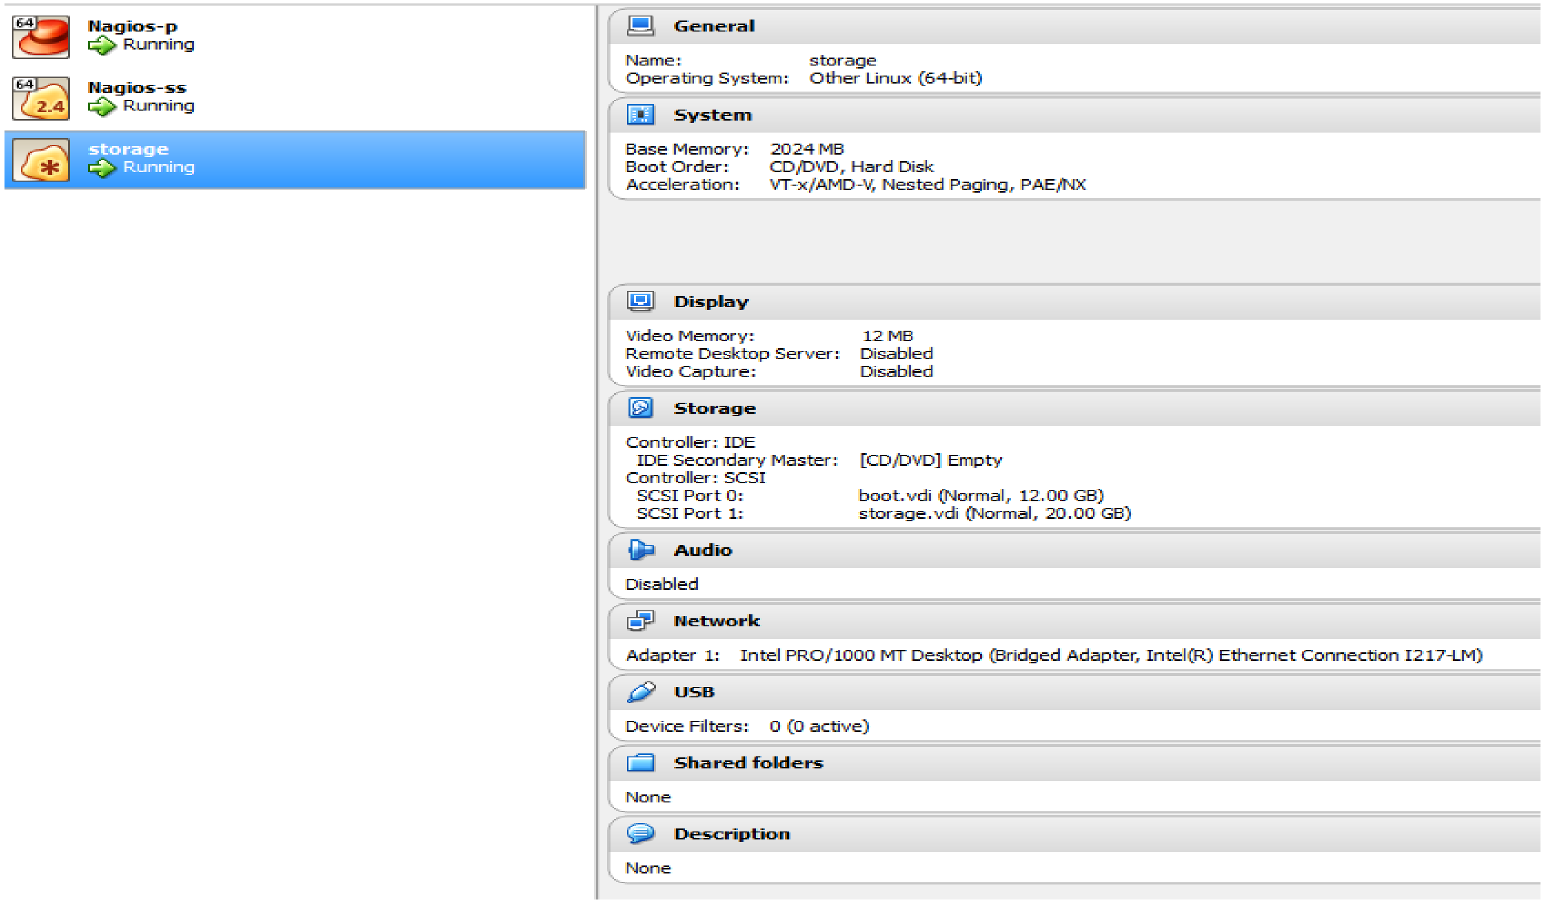

Nagios-p
Running

Nagios-ss
Running

storage
Running

General

Name:	storage
Operating System:	Other Linux (64-bit)

System

Base Memory:	2024 MB
Boot Order:	CD/DVD, Hard Disk
Acceleration:	VT-x/AMD-V, Nested Paging, PAE/NX

Display

Video Memory:	12 MB
Remote Desktop Server:	Disabled
Video Capture:	Disabled

Storage

Controller: IDE	
IDE Secondary Master:	[CD/DVD] Empty
Controller: SCSI	
SCSI Port 0:	boot.vdi (Normal, 12.00 GB)
SCSI Port 1:	storage.vdi (Normal, 20.00 GB)

Audio

Disabled

Network

Adapter 1: Intel PRO/1000 MT Desktop (Bridged Adapter, Intel(R) Ethernet Connection I217-LM)

USB

Device Filters: 0 (0 active)

Shared folders

None

Description

None

Configuration Steps

- Did the online installation of Nagios XI on both CentOS VM

  We used Nagios XI 2014 2.7 version

```
cd /tmp
wget http://assets.nagios.com/downloads/nagiosxi/xi-latest.tar.gz
tar xzf xi-latest.tar.gz
cd /tmp/nagiosxi
./fullinstall
```

- Did the installation of Redhat cluster packages on both server

```
yum install ricci
yum install rgmanager
yum install luci
yum install gfs2-utils
yum install iscsi
```

Configuration Steps

- Created a cluster as test using luci console and added both node in to cluster

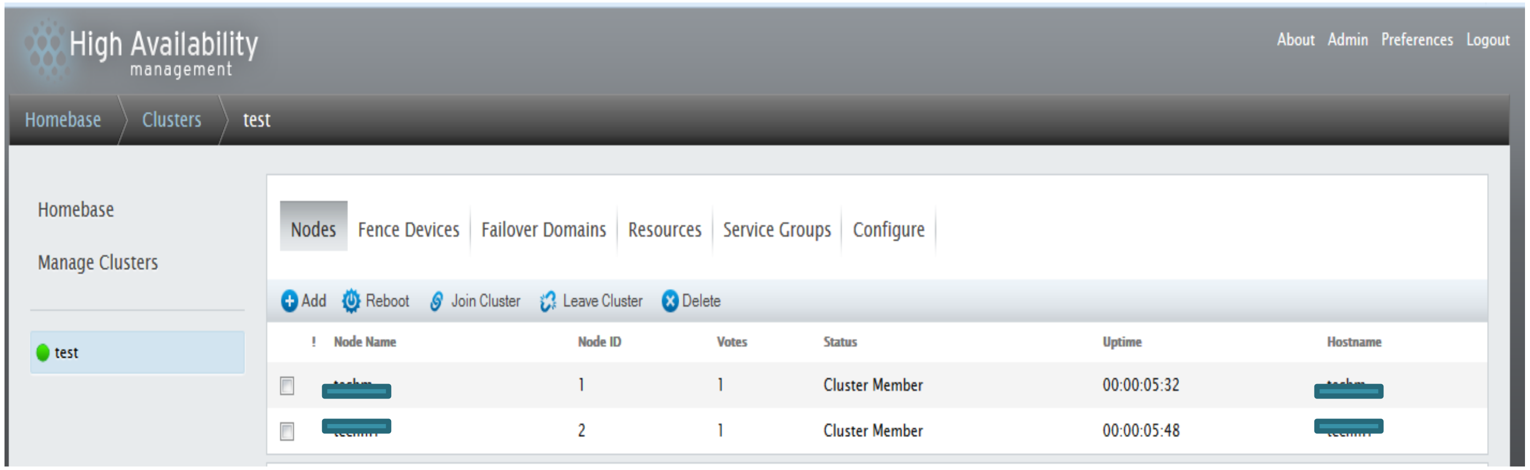

!	Node Name	Node ID	Votes	Status	Uptime	Hostname
	[redacted]	1	1	Cluster Member	00:00:05:32	[redacted]
	[redacted]	2	1	Cluster Member	00:00:05:48	[redacted]

Configuration Steps

- Created two volumes on storage using volume option

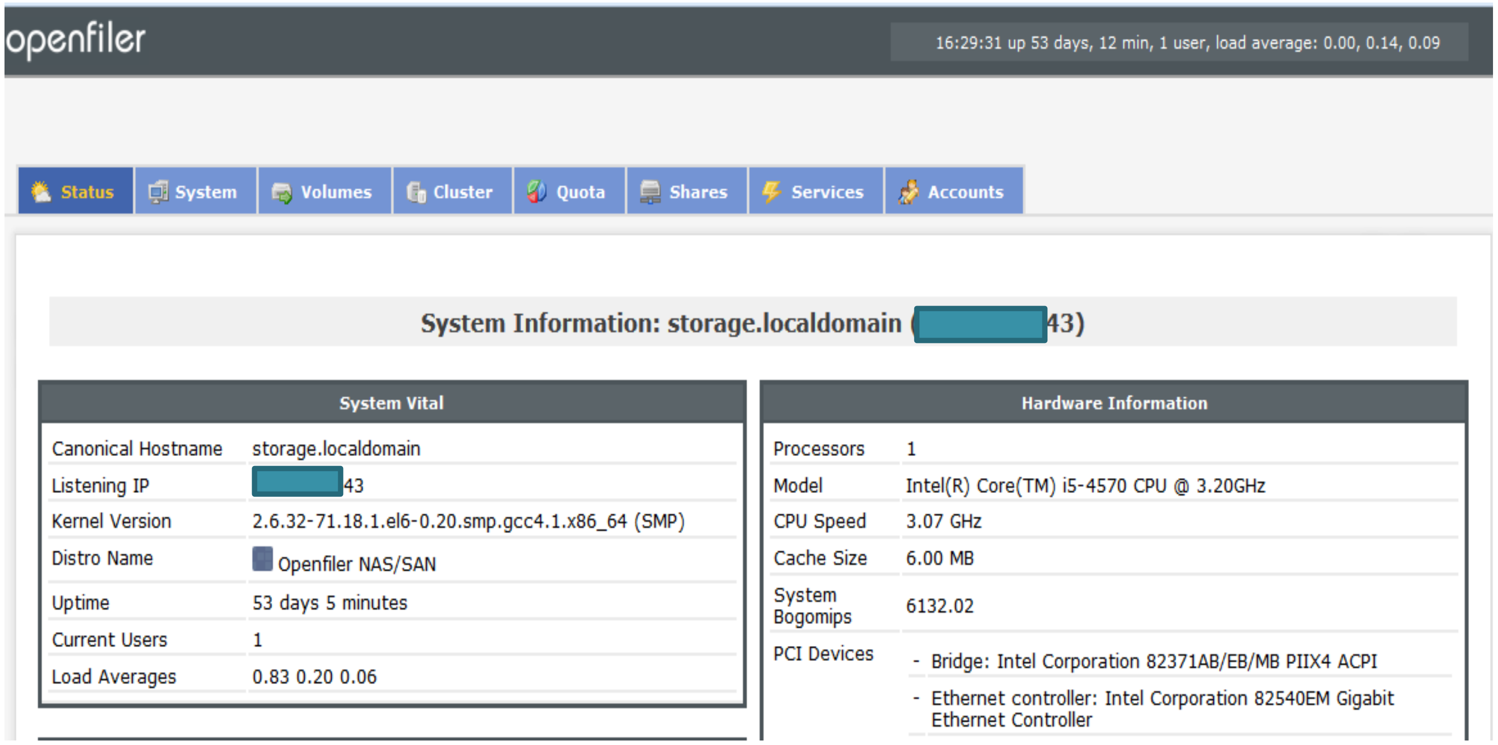

16:29:31 up 53 days, 12 min, 1 user, load average: 0.00, 0.14, 0.09

System Information: storage.localdomain ([redacted]43)

System Vital	
Canonical Hostname	storage.localdomain
Listening IP	[redacted]43
Kernel Version	2.6.32-71.18.1.el6-0.20.smp.gcc4.1.x86_64 (SMP)
Distro Name	Openfiler NAS/SAN
Uptime	53 days 5 minutes
Current Users	1
Load Averages	0.83 0.20 0.06

Hardware Information	
Processors	1
Model	Intel(R) Core(TM) i5-4570 CPU @ 3.20GHz
CPU Speed	3.07 GHz
Cache Size	6.00 MB
System Bogomips	6132.02
PCI Devices	- Bridge: Intel Corporation 82371AB/EB/MB PIIX4 ACPI - Ethernet controller: Intel Corporation 82540EM Gigabit Ethernet Controller

Configuration Steps

- Created two volumes on storage

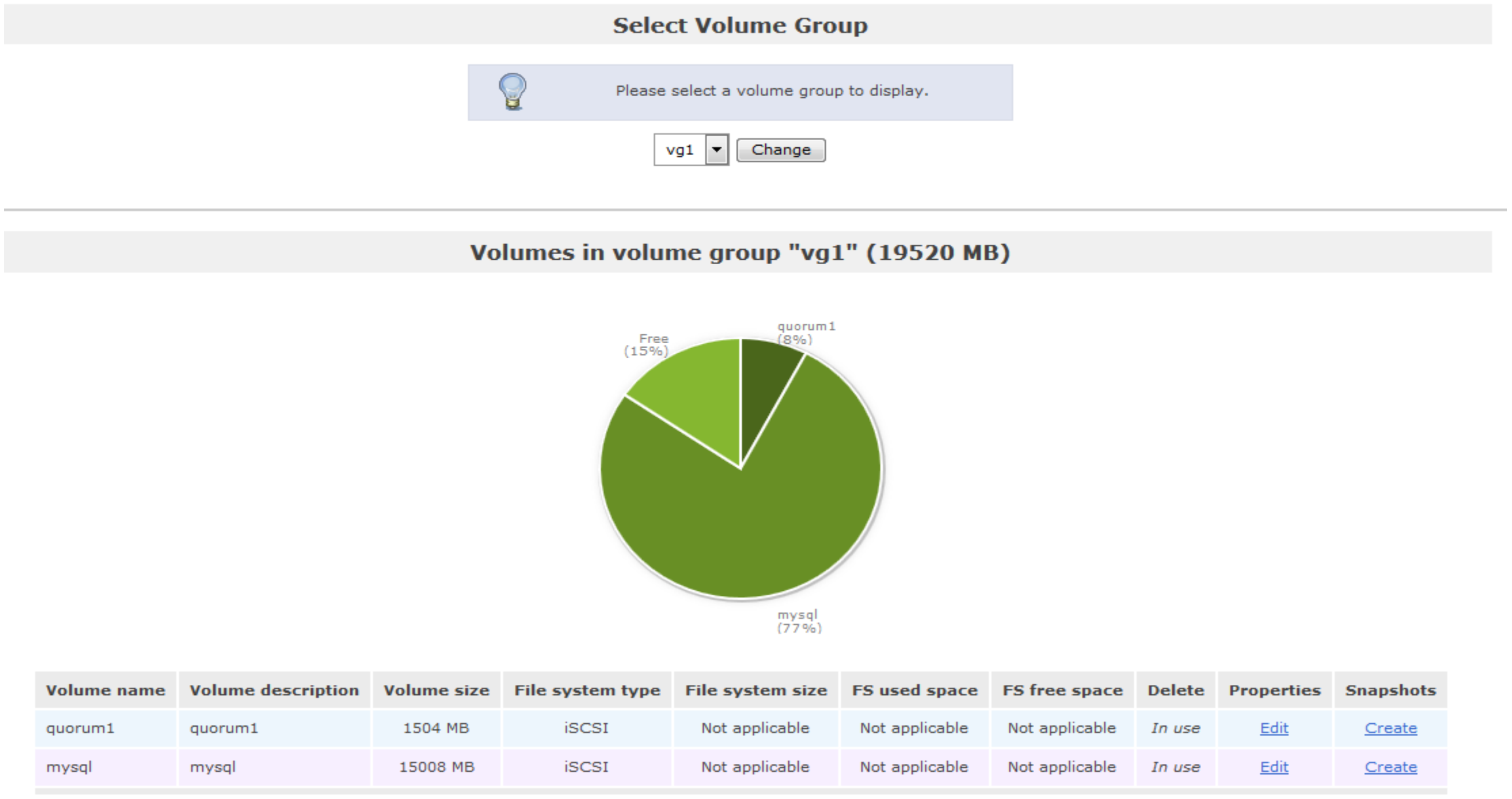

Select Volume Group

Please select a volume group to display.

Volumes in volume group "vg1" (19520 MB)

Volume name	Volume description	Volume size	File system type	File system size	FS used space	FS free space	Delete	Properties	Snapshots
quorum1	quorum1	1504 MB	iSCSI	Not applicable	Not applicable	Not applicable	*In use*	Edit	Create
mysql	mysql	15008 MB	iSCSI	Not applicable	Not applicable	Not applicable	*In use*	Edit	Create

Configuration Steps

- Mapped volumes on both server as sdb & sdc

iscsiadm -m discovery -t sendtargets –p 1[redacted]43

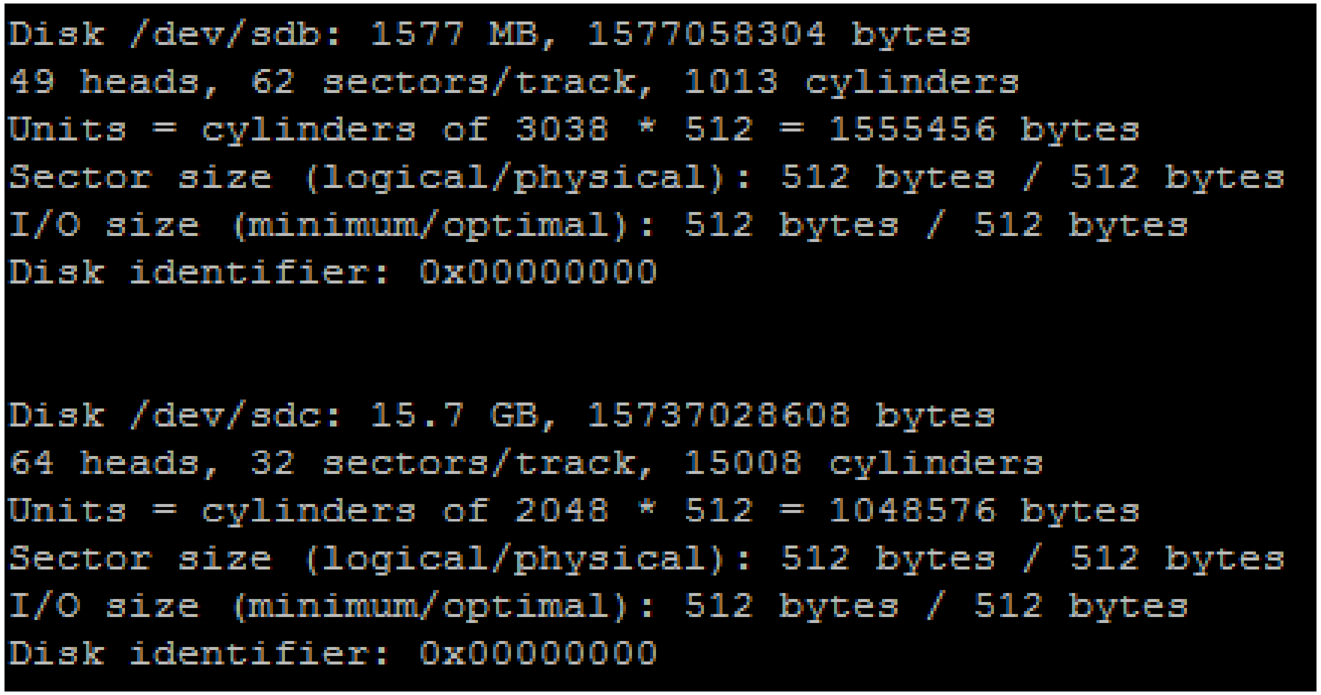

```
Disk /dev/sdb: 1577 MB, 1577058304 bytes
49 heads, 62 sectors/track, 1013 cylinders
Units = cylinders of 3038 * 512 = 1555456 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000000


Disk /dev/sdc: 15.7 GB, 15737028608 bytes
64 heads, 32 sectors/track, 15008 cylinders
Units = cylinders of 2048 * 512 = 1048576 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000000
```

Configuration Steps

- Created quorum disk using the shared partition sdb

mkqdisk -c /dev/sdb -l quorum1

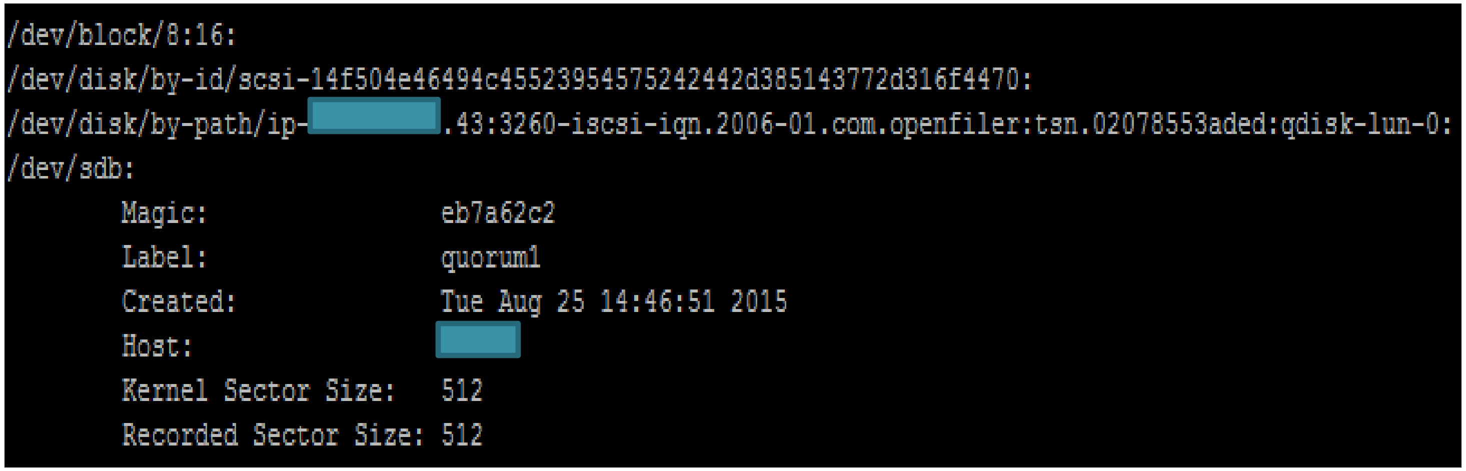

```
/dev/block/8:16:
/dev/disk/by-id/scsi-14f504e46494c45523954575242442d385143772d316f4470:
/dev/disk/by-path/ip-[redacted].43:3260-iscsi-iqn.2006-01.com.openfiler:tsn.02078553aded:qdisk-lun-0:
/dev/sdb:
        Magic:                  eb7a62c2
        Label:                  quorum1
        Created:                Tue Aug 25 14:46:51 2015
        Host:                   [redacted]
        Kernel Sector Size:     512
        Recorded Sector Size:   512
```

Configuration Steps

- Formatted the sdc partition with gfs2 partition and created a directory name as common on both server

mkfs.gfs2 -p lock_dlm -t test:GFS -j 2 /dev/sdc

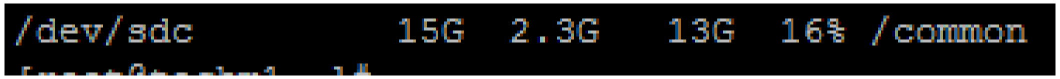

```
/dev/sdc              15G  2.3G   13G  16% /common
```

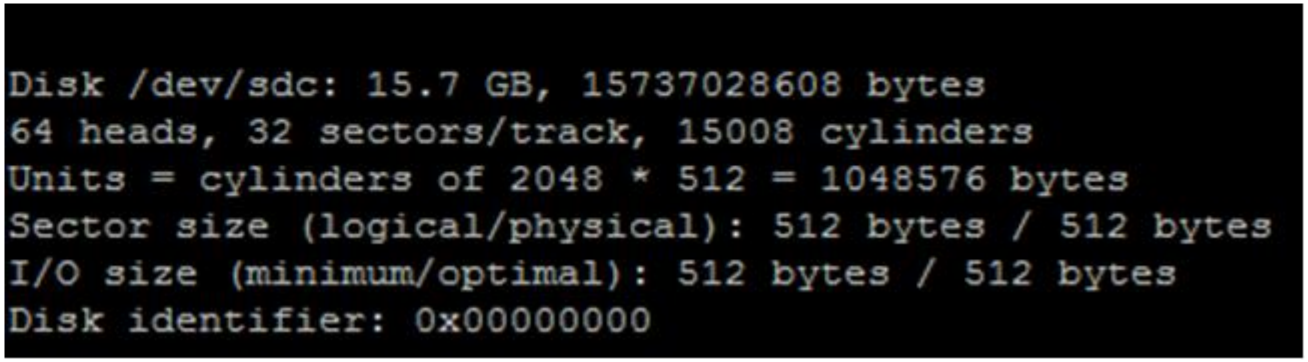

```
Disk /dev/sdc: 15.7 GB, 15737028608 bytes
64 heads, 32 sectors/track, 15008 cylinders
Units = cylinders of 2048 * 512 = 1048576 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000000
```

Configuration Steps

- Created a failover domain and added both device in to

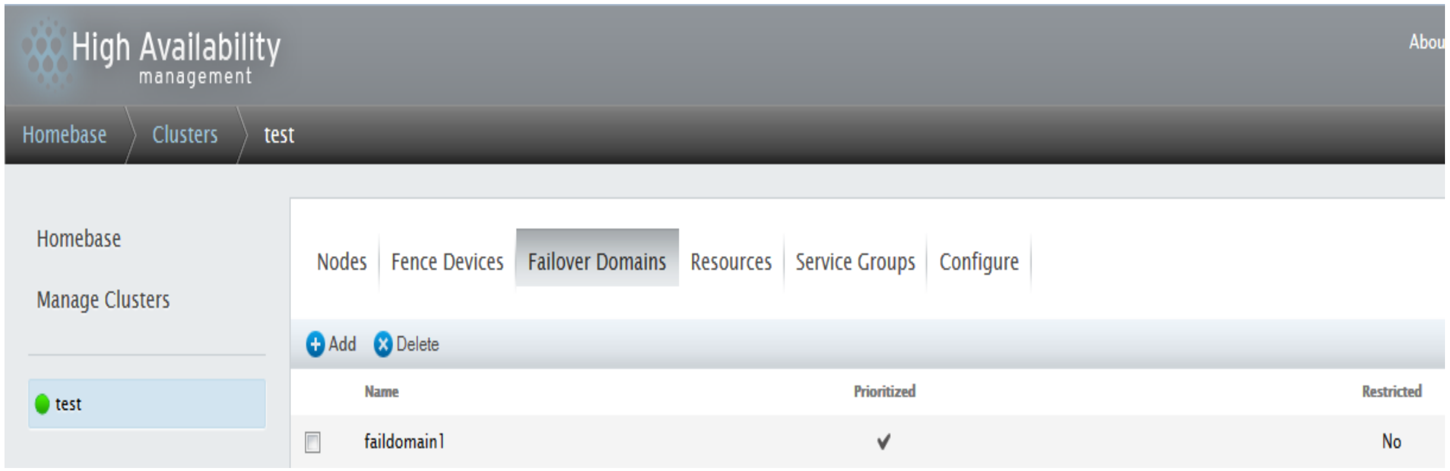

	Name	Prioritized	Restricted
☐	faildomain1	✓	No

Configuration Steps

- Created one resource as gfs for file system failover using common as mount point

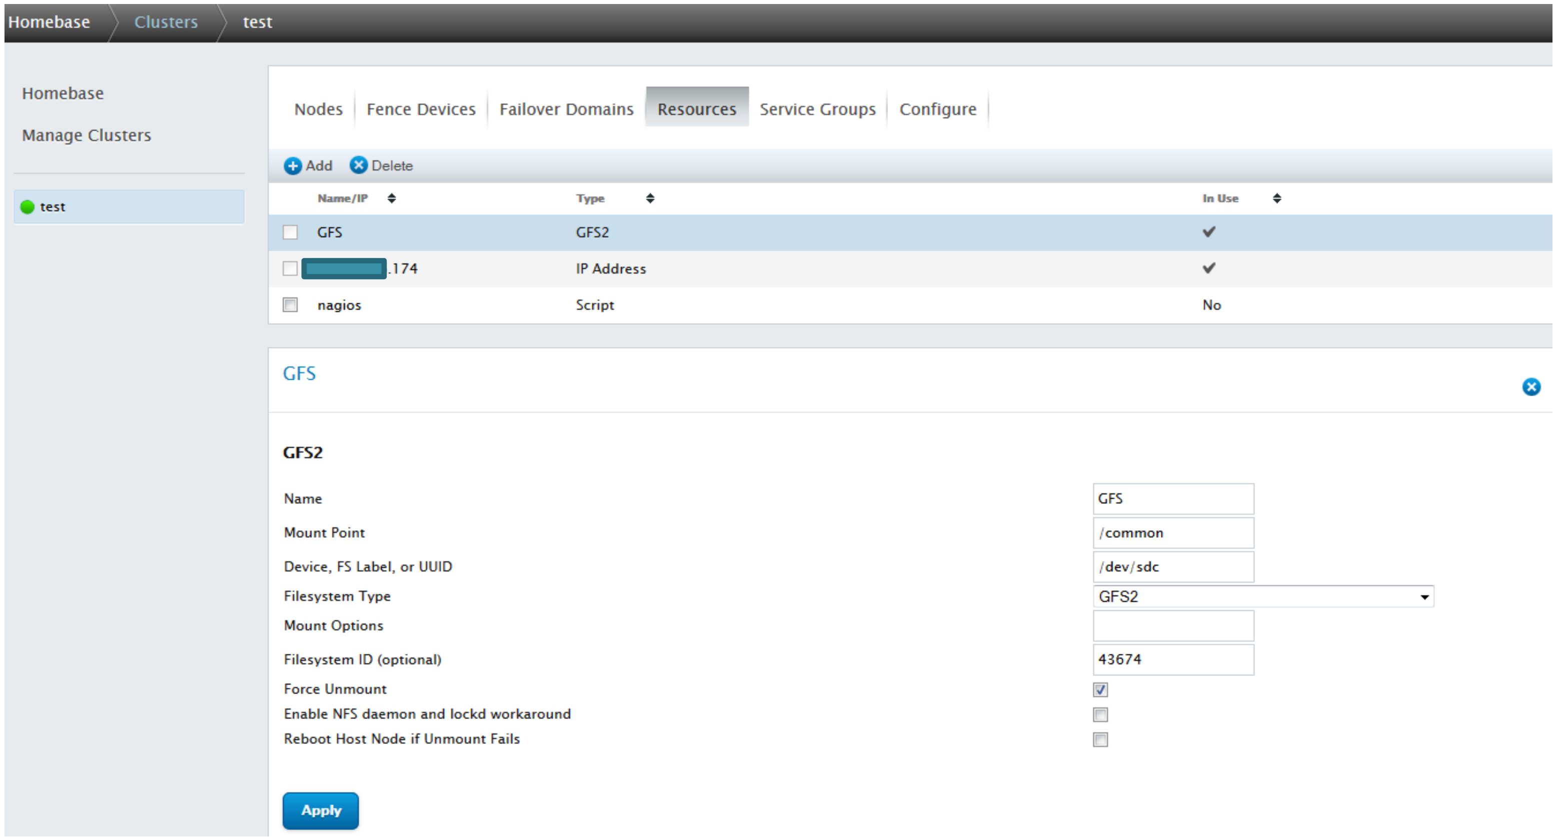

	Name/IP	Type	In Use
☐	GFS	GFS2	✓
☐	[redacted].174	IP Address	✓
☐	nagios	Script	No

GFS

GFS2

Name	GFS
Mount Point	/common
Device, FS Label, or UUID	/dev/sdc
Filesystem Type	GFS2
Mount Options	
Filesystem ID (optional)	43674
Force Unmount	☑
Enable NFS daemon and lockd workaround	☐
Reboot Host Node if Unmount Fails	☐

Configuration Steps

- Created second resource as VIP, so it will be function able from any cluster node

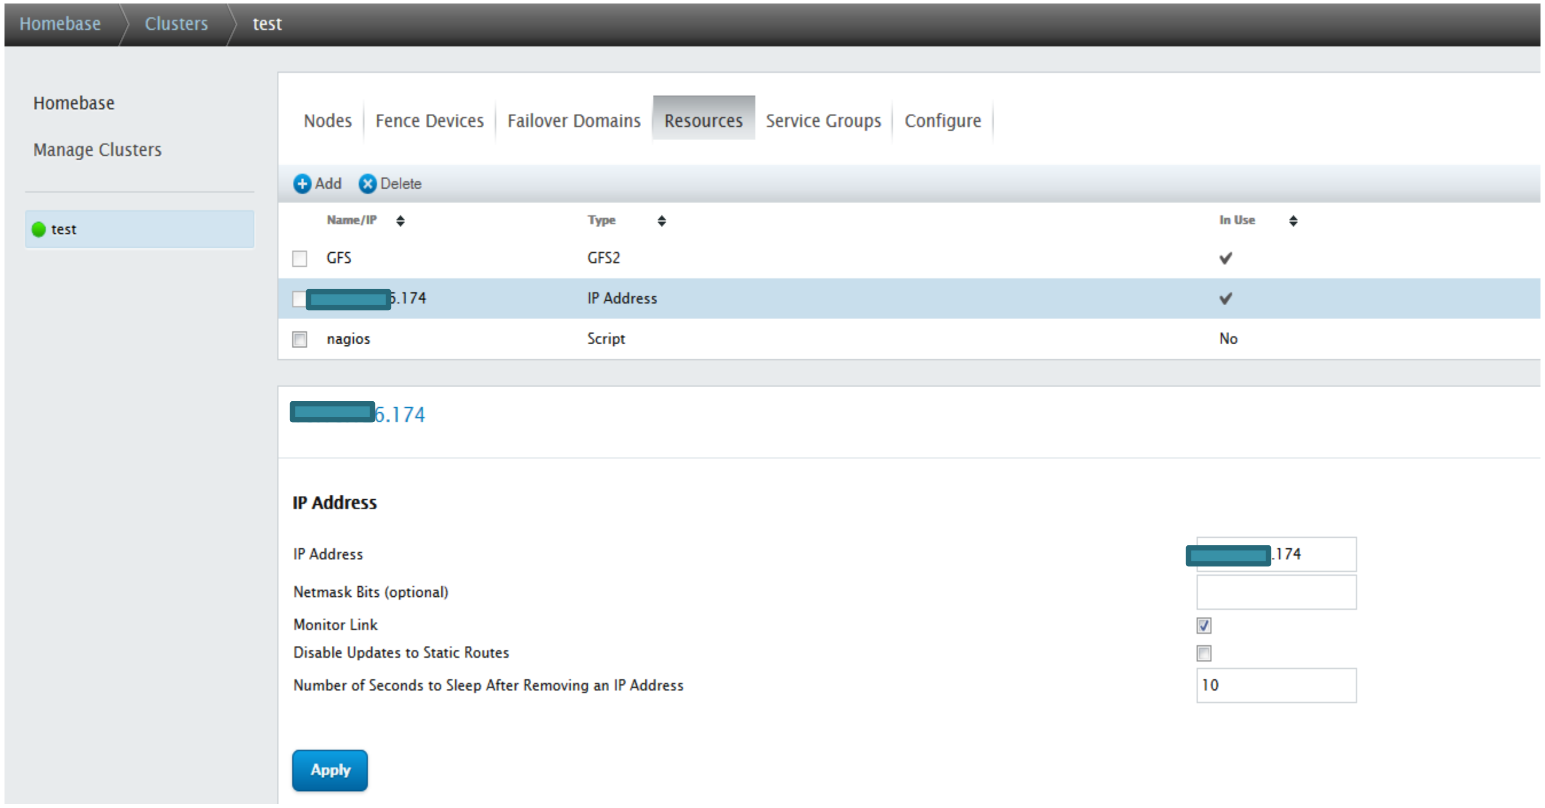

Name/IP	Type	In Use
GFS	GFS2	✓
[redacted].174	IP Address	✓
nagios	Script	No

[redacted].174

IP Address

IP Address	[redacted].174
Netmask Bits (optional)	
Monitor Link	✓
Disable Updates to Static Routes	☐
Number of Seconds to Sleep After Removing an IP Address	10

Configuration Steps

- Created a service group and mapped both resources with it

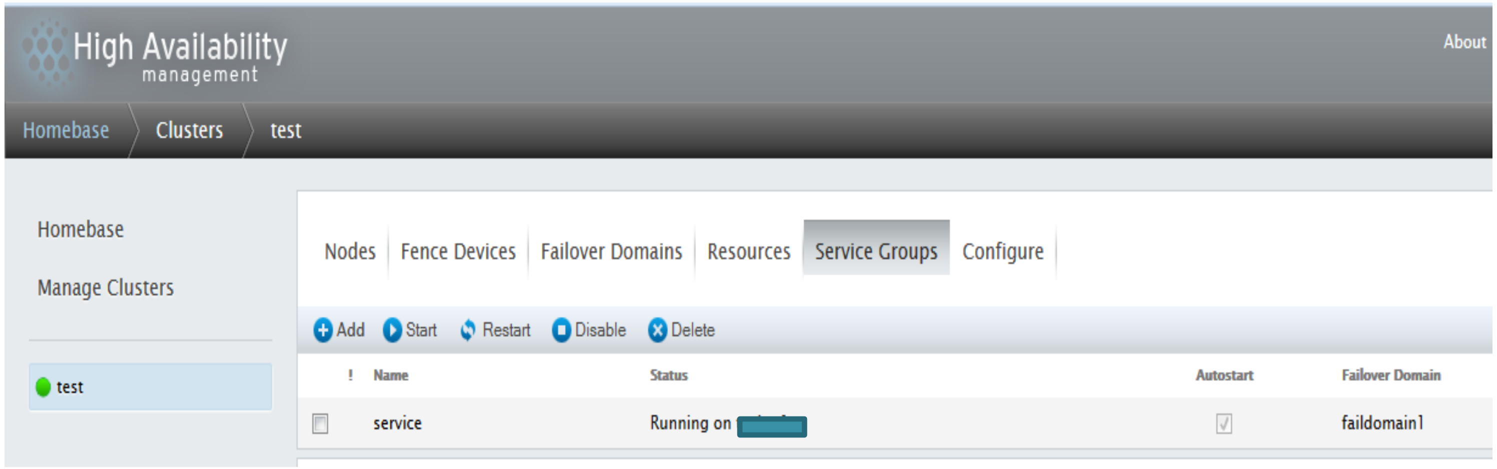

!	Name	Status	Autostart	Failover Domain
	service	Running on [redacted]	✓	faildomain1

Configuration Steps

- Stopped all services on both node and move configuration files and DB from primary server to common partition (/etc/httpd/conf.d, /etc/nagiosql, /usr/local/nagios /usr/local/nagiosxi, /usr/local/nrdp, /var/www/html/nagiosql, /var/lib/mysql, /var/lib/pgsql

```
service nagios stop && chkconfig nagios off
service mysqld stop && chkconfig mysqld off
service postgresql stop && chkconfig postgresql off
service npcd stop && chkconfig npcd off
service ndo2db stop && chkconfig ndo2db off
```

```
mv /etc/httpd/conf.d /common
mv /etc/nagiosql /common
mv /usr/local/nagios /common
mv /usr/local/nagiosxi /common
mv /usr/local/nrdp /common
mv /var/www/html/nagiosql /common/main
mv /var/lib/mysql  /common
mv /var/lib/pgsql /common
```

Configuration Steps

- Removed the default folders from second server. Created sim link for all moved folders from common to their default location on first server then mount common partition on second server and created sim link on second server as well

```
rm –rf /etc/httpd/conf.d && ln –s /common/conf.d /etc/httpd/conf.d
rm –rf /etc/nagiosql && ln –s /common/nagiosql /etc
rm –rf /usr/local/nagios && ln –s /common/nagios /usr/local
rm –rf /usr/local/nagiosxi && ln –s /common/nagiosxi /usr/local
rm –rf /usr/local/nrdp && ln –s /common/nrdp /usr/local
rm –rf /var/www/html/nagiosql && ln-s /common/main/nagiosql /var/www/html
rm –rf /var/lib/mysql && ln –s /common/mysql /var/lib
rm –rf /var/lib/pgsql && ln –s /common/pgsql /var/lib
```

Configuration Steps

- We created a script and put run on both server via cron to start the off services on active node.

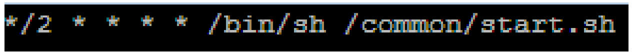

```
*/2 * * * * /bin/sh /common/start.sh
```

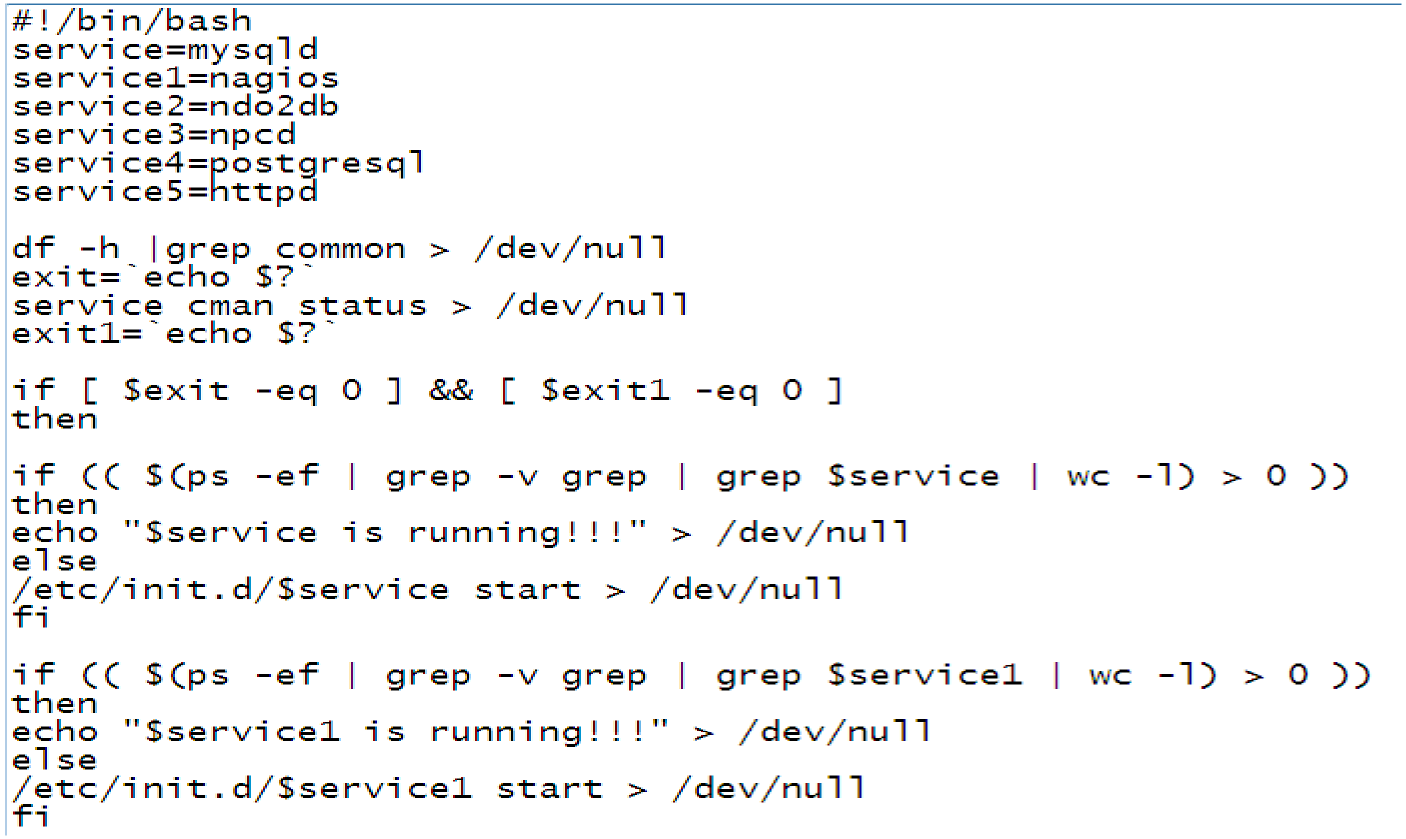

```
#!/bin/bash
service=mysqld
service1=nagios
service2=ndo2db
service3=npcd
service4=postgresql
service5=httpd

df -h |grep common > /dev/null
exit=`echo $?`
service cman status > /dev/null
exit1=`echo $?`

if [ $exit -eq 0 ] && [ $exit1 -eq 0 ]
then

if (( $(ps -ef | grep -v grep | grep $service | wc -l) > 0 ))
then
echo "$service is running!!!" > /dev/null
else
/etc/init.d/$service start > /dev/null
fi

if (( $(ps -ef | grep -v grep | grep $service1 | wc -l) > 0 ))
then
echo "$service1 is running!!!" > /dev/null
else
/etc/init.d/$service1 start > /dev/null
fi
```

Configuration Steps

```
if (( $(ps -ef | grep -v grep | grep $service2 | wc -l) > 0 ))
then
echo "$service2 is running!!!" > /dev/null
else
/etc/init.d/$service2 start > /dev/null
fi
if (( $(ps -ef | grep -v grep | grep $service3 | wc -l) > 0 ))
then
echo "$service3 is running!!!" > /dev/null
else
/etc/init.d/$service3 start > /dev/null
fi

if (( $(ps -ef | grep -v grep | grep $service4 | wc -l) > 0 ))
then
echo "$service4 is running!!!" > /dev/null
else
/etc/init.d/$service4 start > /dev/null
fi

if (( $(ps -ef | grep -v grep | grep $service5 | wc -l) > 0 ))
then
echo "$service5 is running!!!" > /dev/null
else
/etc/init.d/$service5 start > /dev/null
fi

else
echo "FS not available"
fi
```

Configuration Steps

- Test case – Application is successfully work & access whole configuration with VIP if any node down

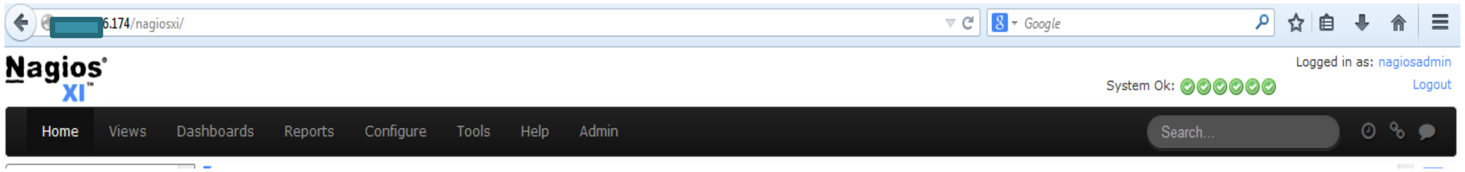

[redacted]6.174/nagiosxi/

- Cluster status on server

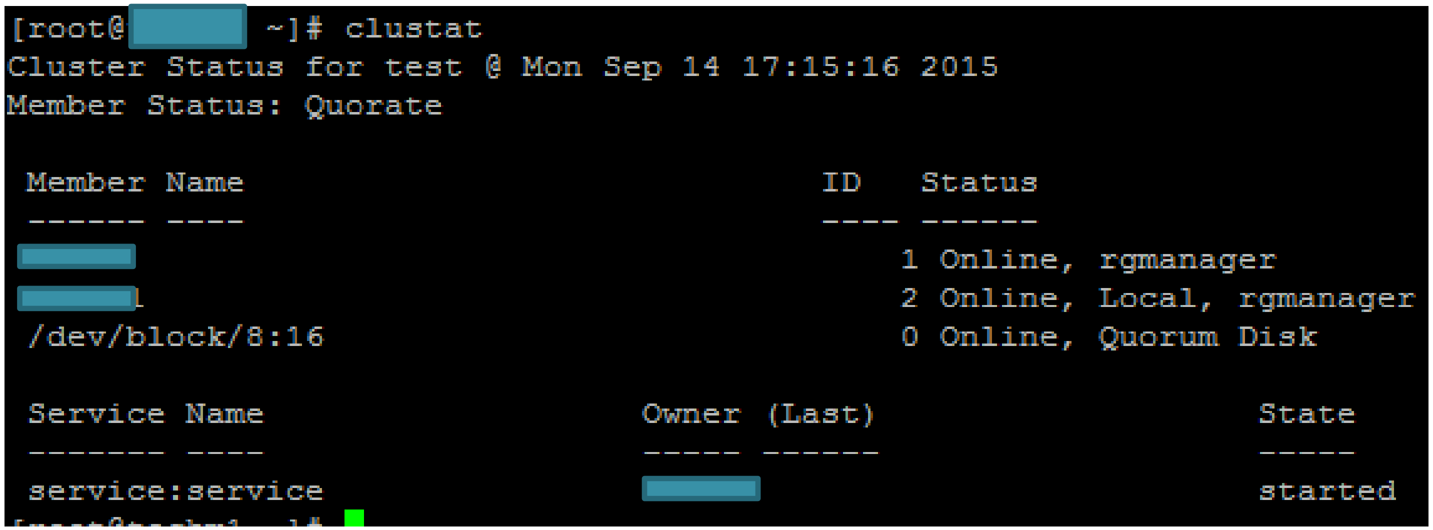

```
[root@[redacted] ~]# clustat
Cluster Status for test @ Mon Sep 14 17:15:16 2015
Member Status: Quorate

 Member Name                                      ID   Status
 ------ ----                                      ---- ------
 [redacted]                                          1 Online, rgmanager
 [redacted]l                                         2 Online, Local, rgmanager
 /dev/block/8:16                                     0 Online, Quorum Disk

 Service Name                     Owner (Last)                       State
 ------- ----                     ----- ------                       -----
 service:service                  [redacted]                         started
```

Thank You