

RED HAT
SUMMIT

BOSTON, MA
JUNE 23-26, 2015

RED HAT STORAGE FOR MERE MORTALS

Thomas Cameron, RHCA, RHCSS, RHCDS, RHCVA, RHCX
Chief Architect, Central US

AGENDA

AGENDA

- What is Red Hat Storage?
- History
- Community
- Architecture
- Node Installation
- Installing RHS
- Adding servers to the trusted pool
- Distributed volumes
- Replicated volumes
- Distributed + replicated volumes

AGENDA

- Attaching to Red Hat Storage with NFS

WHAT IS RED HAT STORAGE?

WHAT IS RED HAT STORAGE

- Red Hat Storage is a commercially supported distribution of GlusterFS
- GlusterFS is a scale-out network-attached storage filesystem which uses commodity x86 hardware and inexpensive disks, JBODs, or (rarely) SAN
- It can use ethernet or InfiniBand RDMA for transport

HISTORY

HISTORY

- The company Gluster was founded in 2005 by founder and CTO Anand Babu Periasamy (now an angel investor in Silicon Valley)
- In 2010, Ben Golub (now CEO of Docker) became the CEO
- In 2011, Red Hat acquired Gluster
- In 2012, Red Hat Storage was launched

COMMUNITY

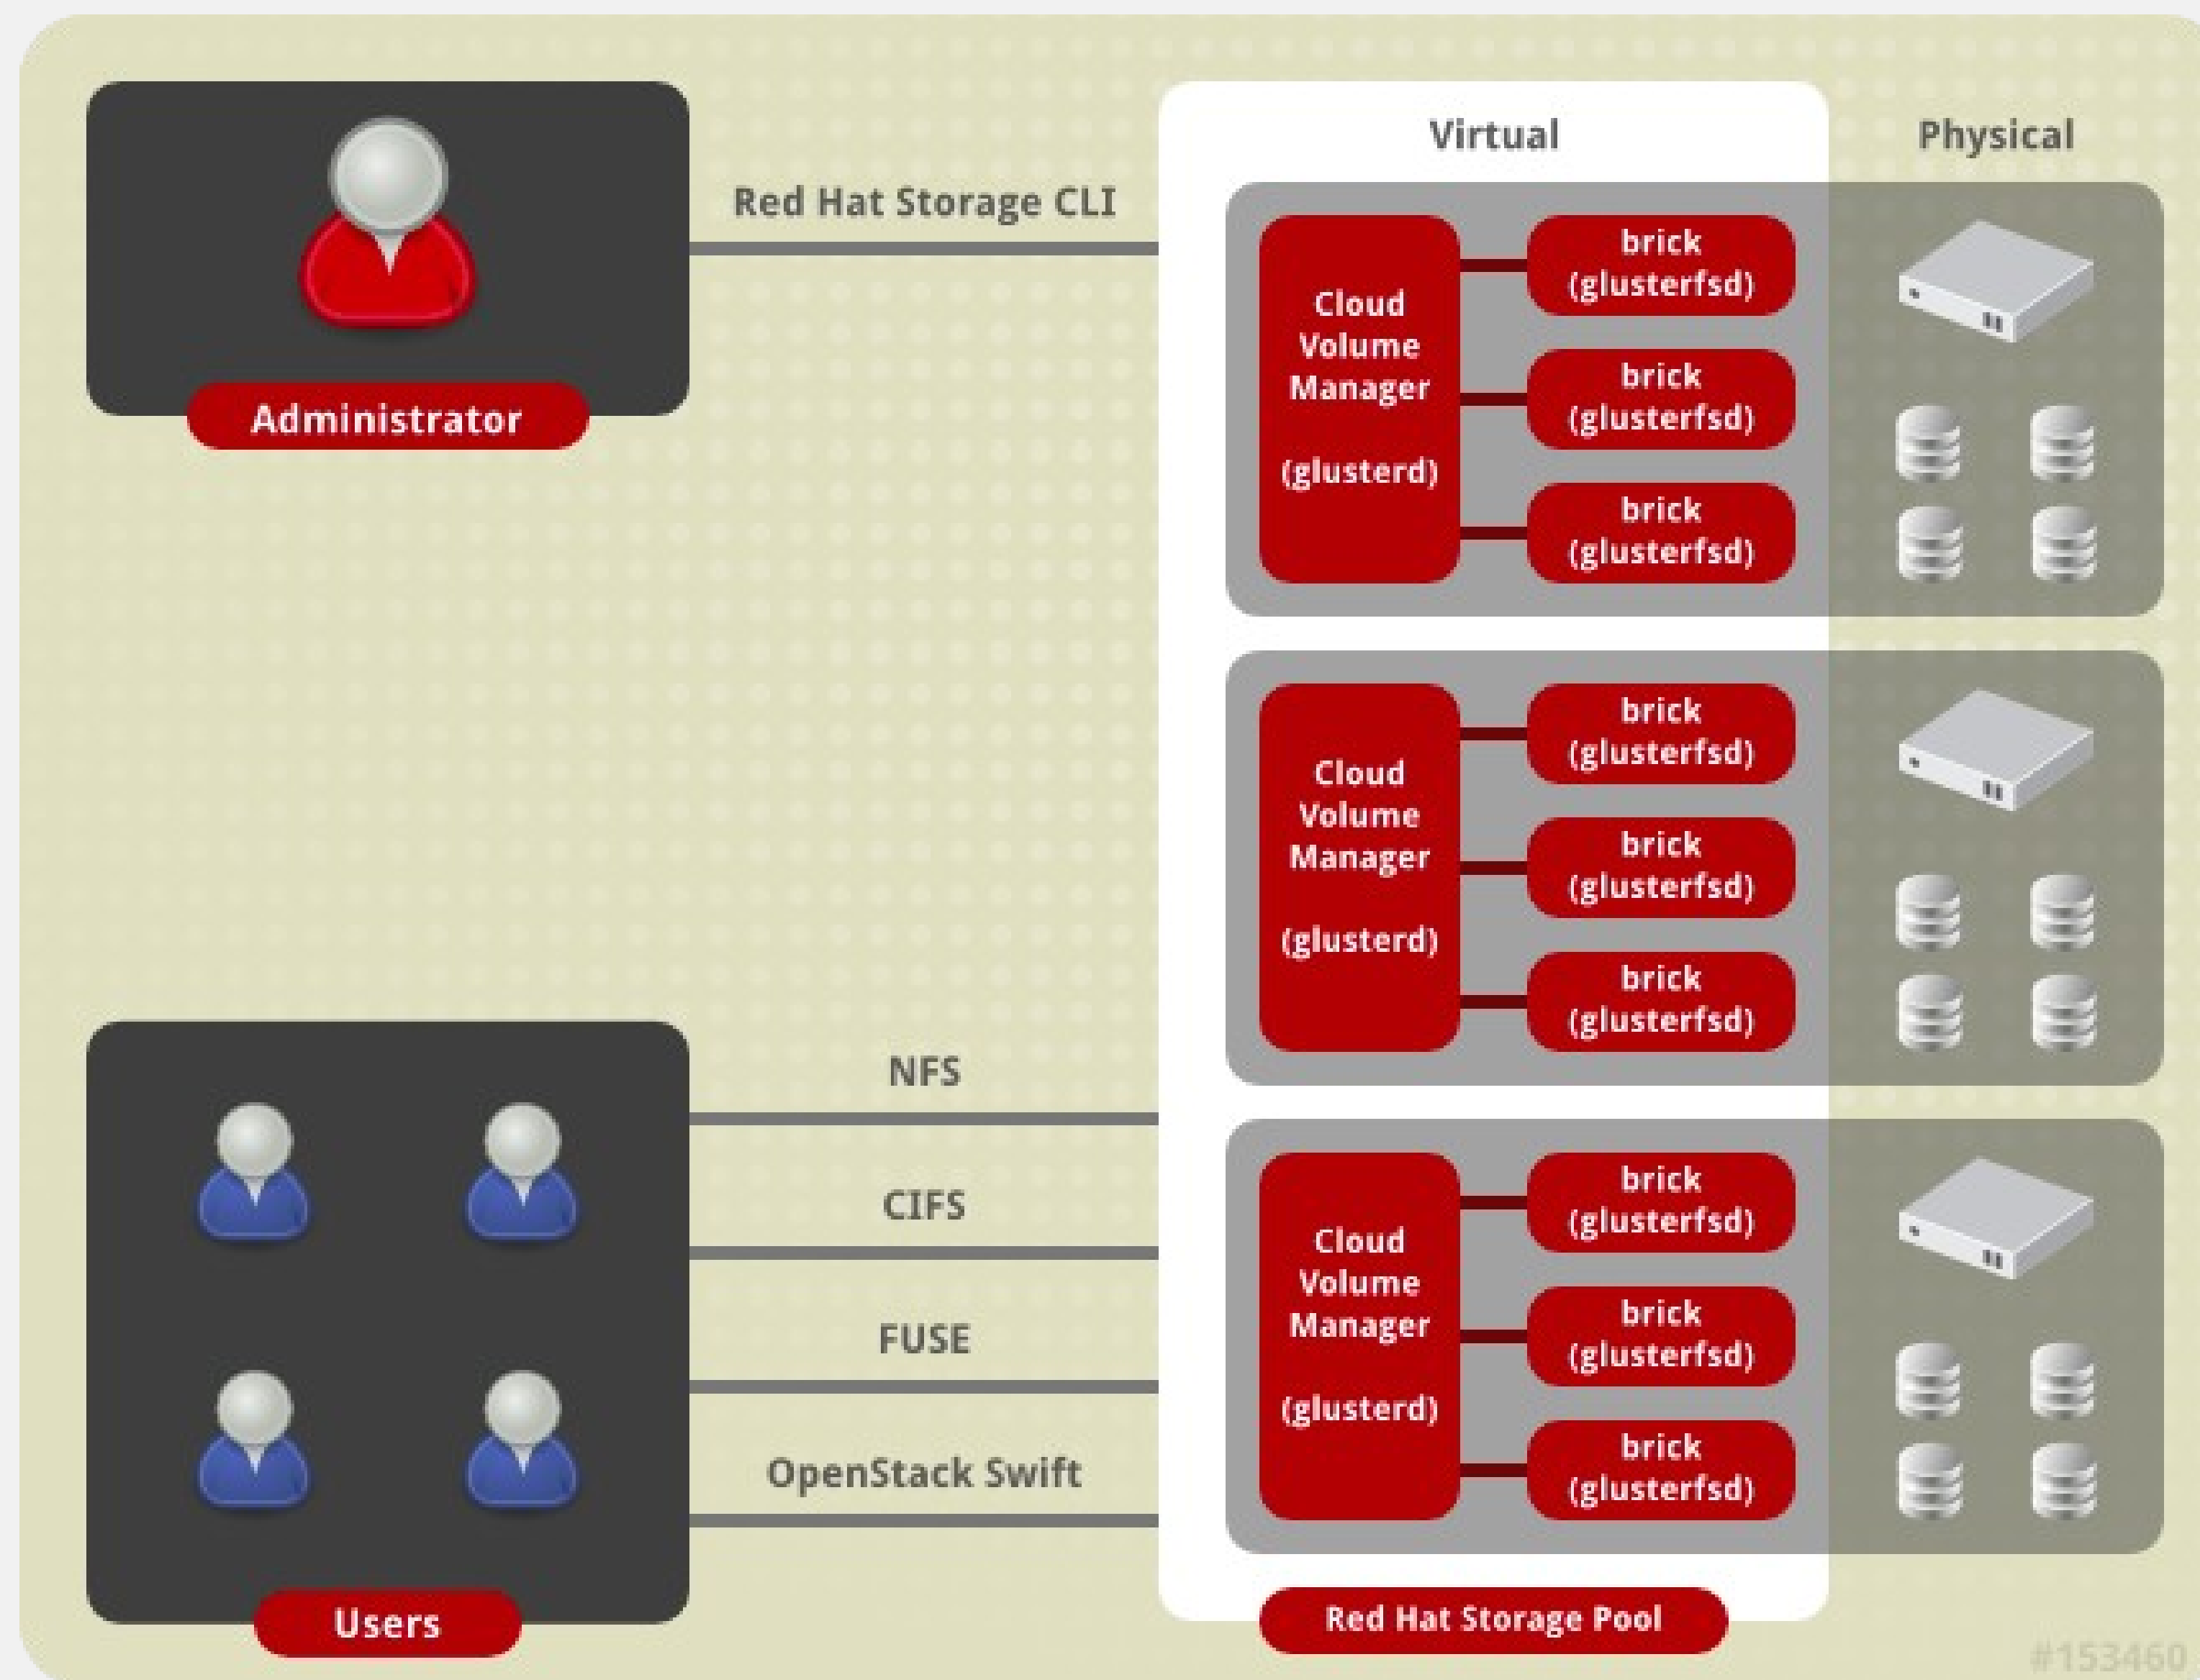
COMMUNITY

- As Fedora is the upstream project from which Red Hat uses technologies to build Red Hat Enterprise Linux, the Gluster Project (www.gluster.org) is the upstream for Red Hat Storage.
- Documentation, IRC channels, development tools and docs, etc. are all there for the community.
- Community governance with Red Hat's sponsorship.

ARCHITECTURE

ARCHITECTURE

- Server component (glusterfsd)
- Client component (glusterfs)
- The client runs in user space (FUSE)
- File location determined by a hashing algorithm - no metadata services
 - No SPOF
- Global namespace



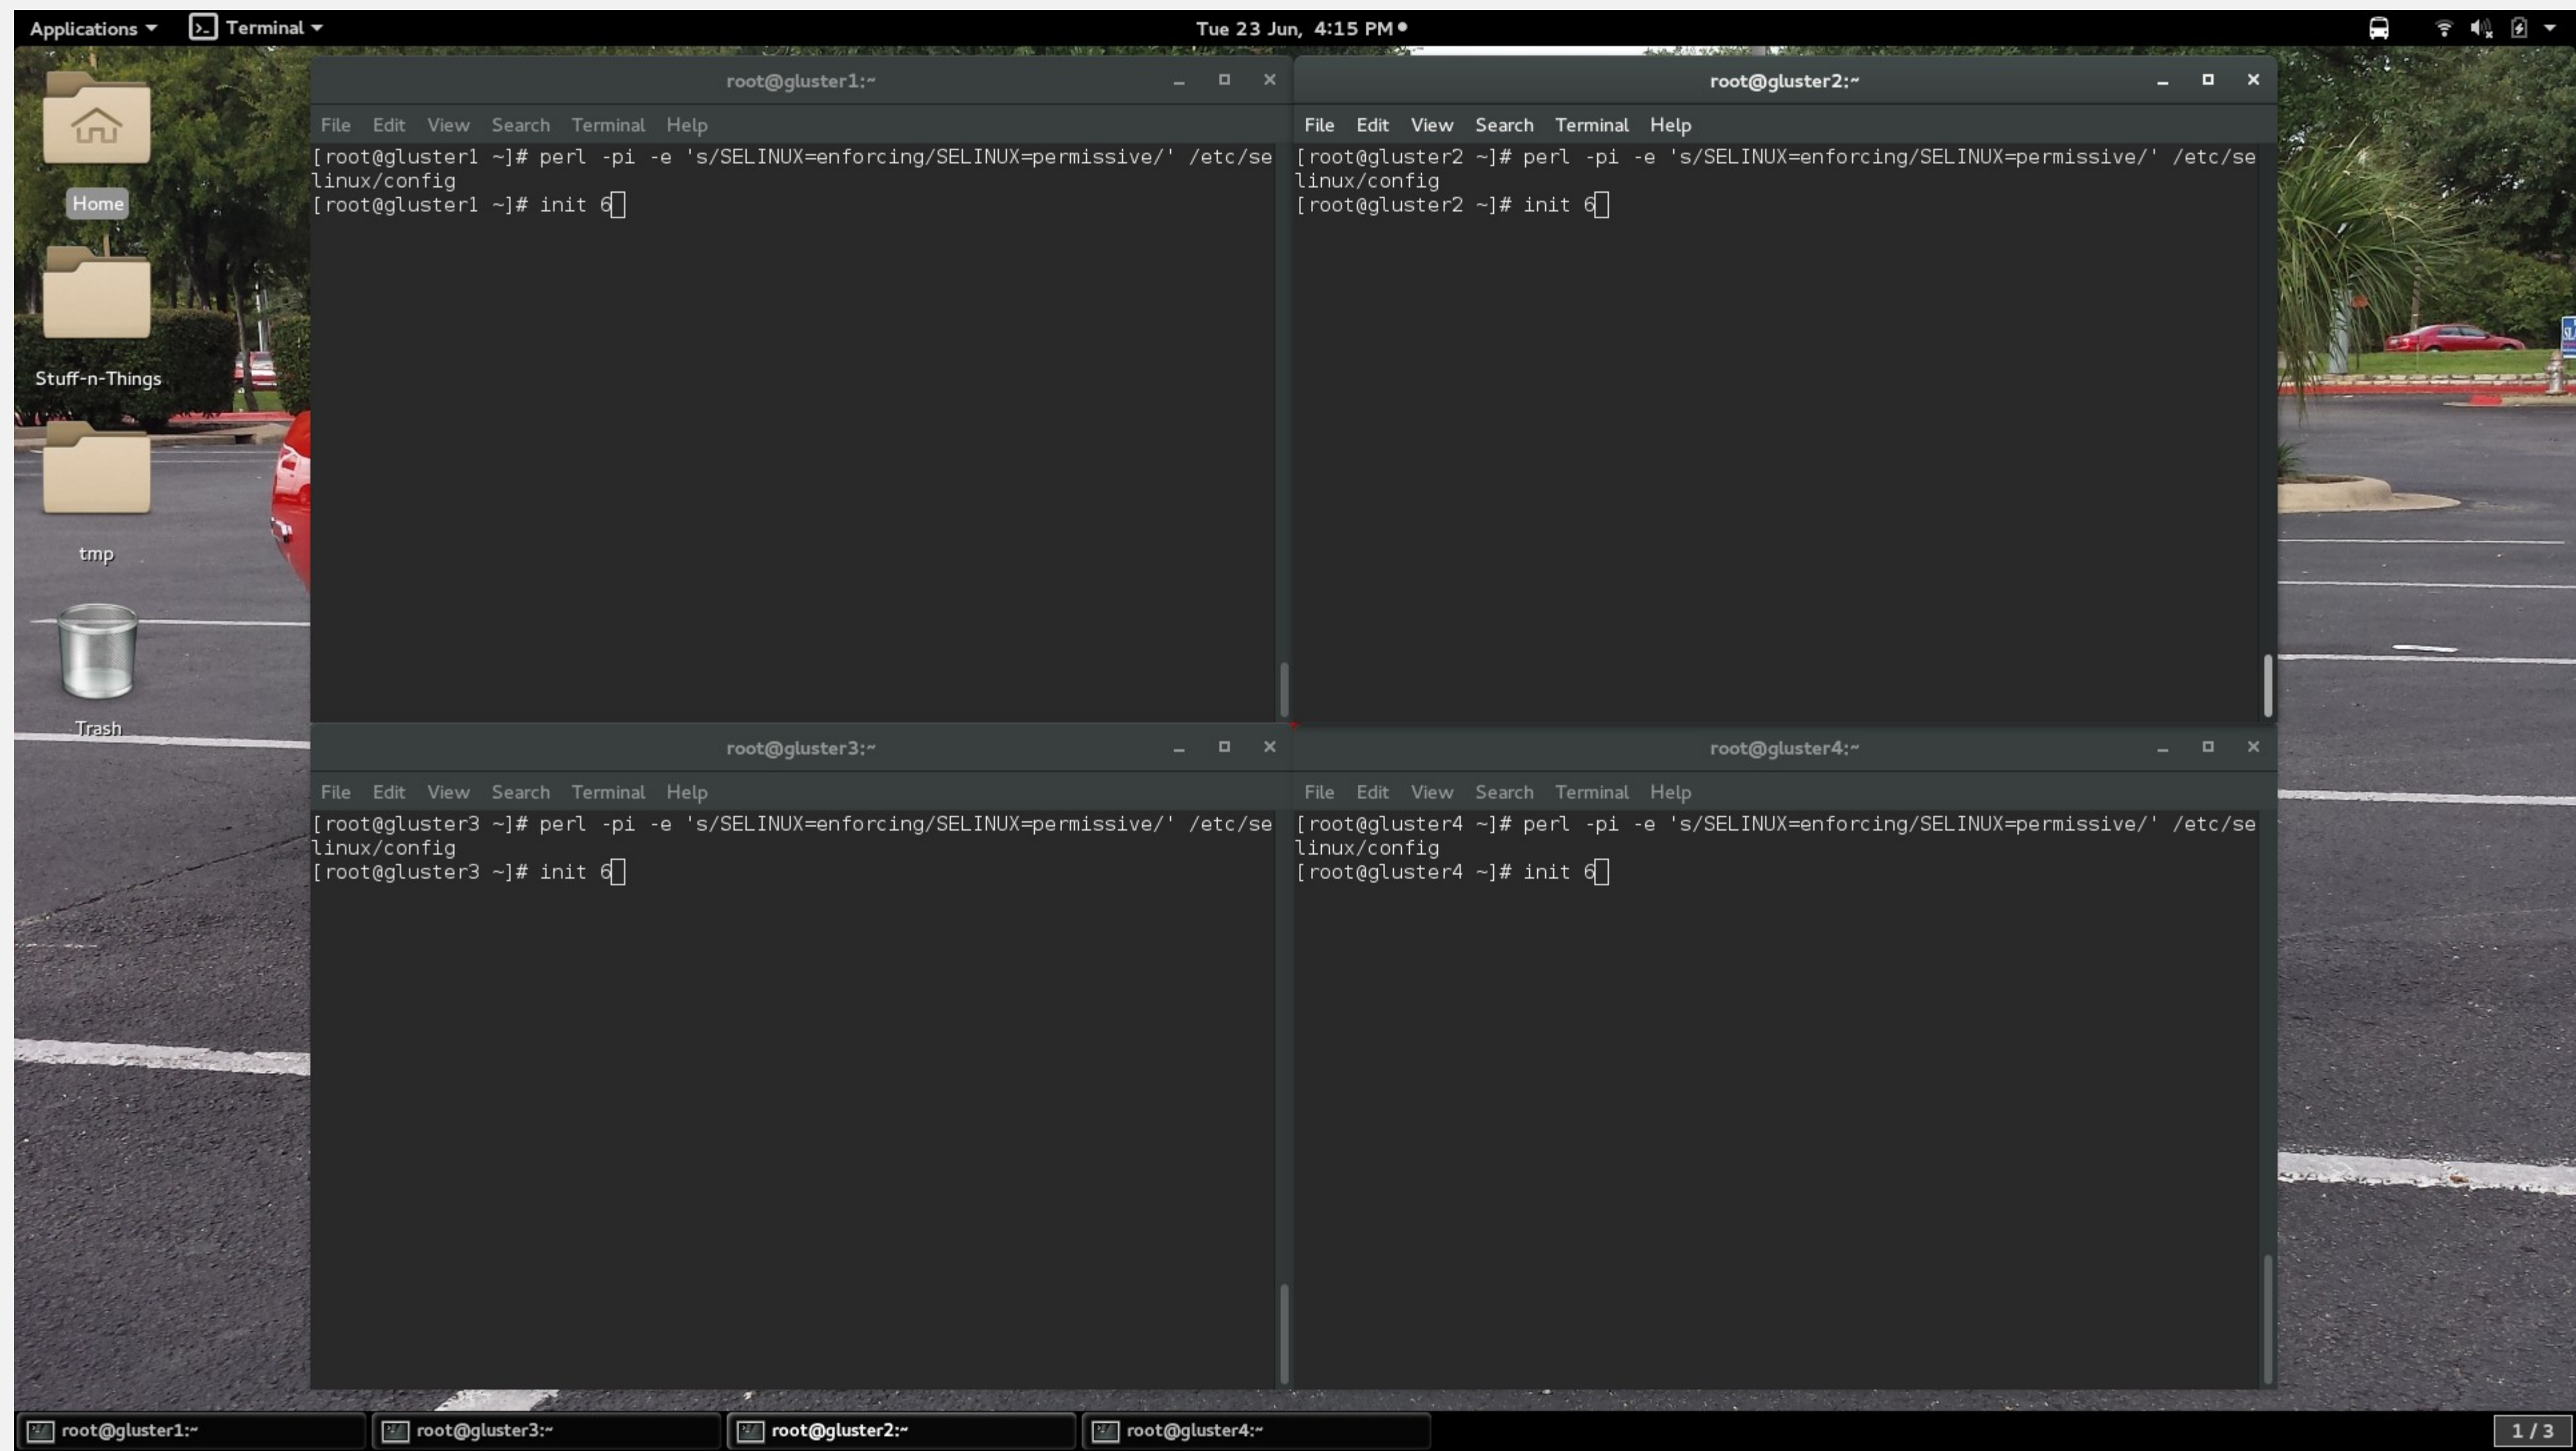
NODE INSTALLATION

NODE INSTALLATION


```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# chkconfig iptables off; service iptables stop  
iptables: Setting chains to policy ACCEPT: filter [ OK ]  
iptables: Flushing firewall rules: [ OK ]  
iptables: Unloading modules: [ OK ]  
[root@gluster1 ~]#
```



```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# perl -pi -e 's/SELINUX=enforcing/SELINUX=permissive/' /etc/se  
linux/config  
[root@gluster1 ~]# init 6
```

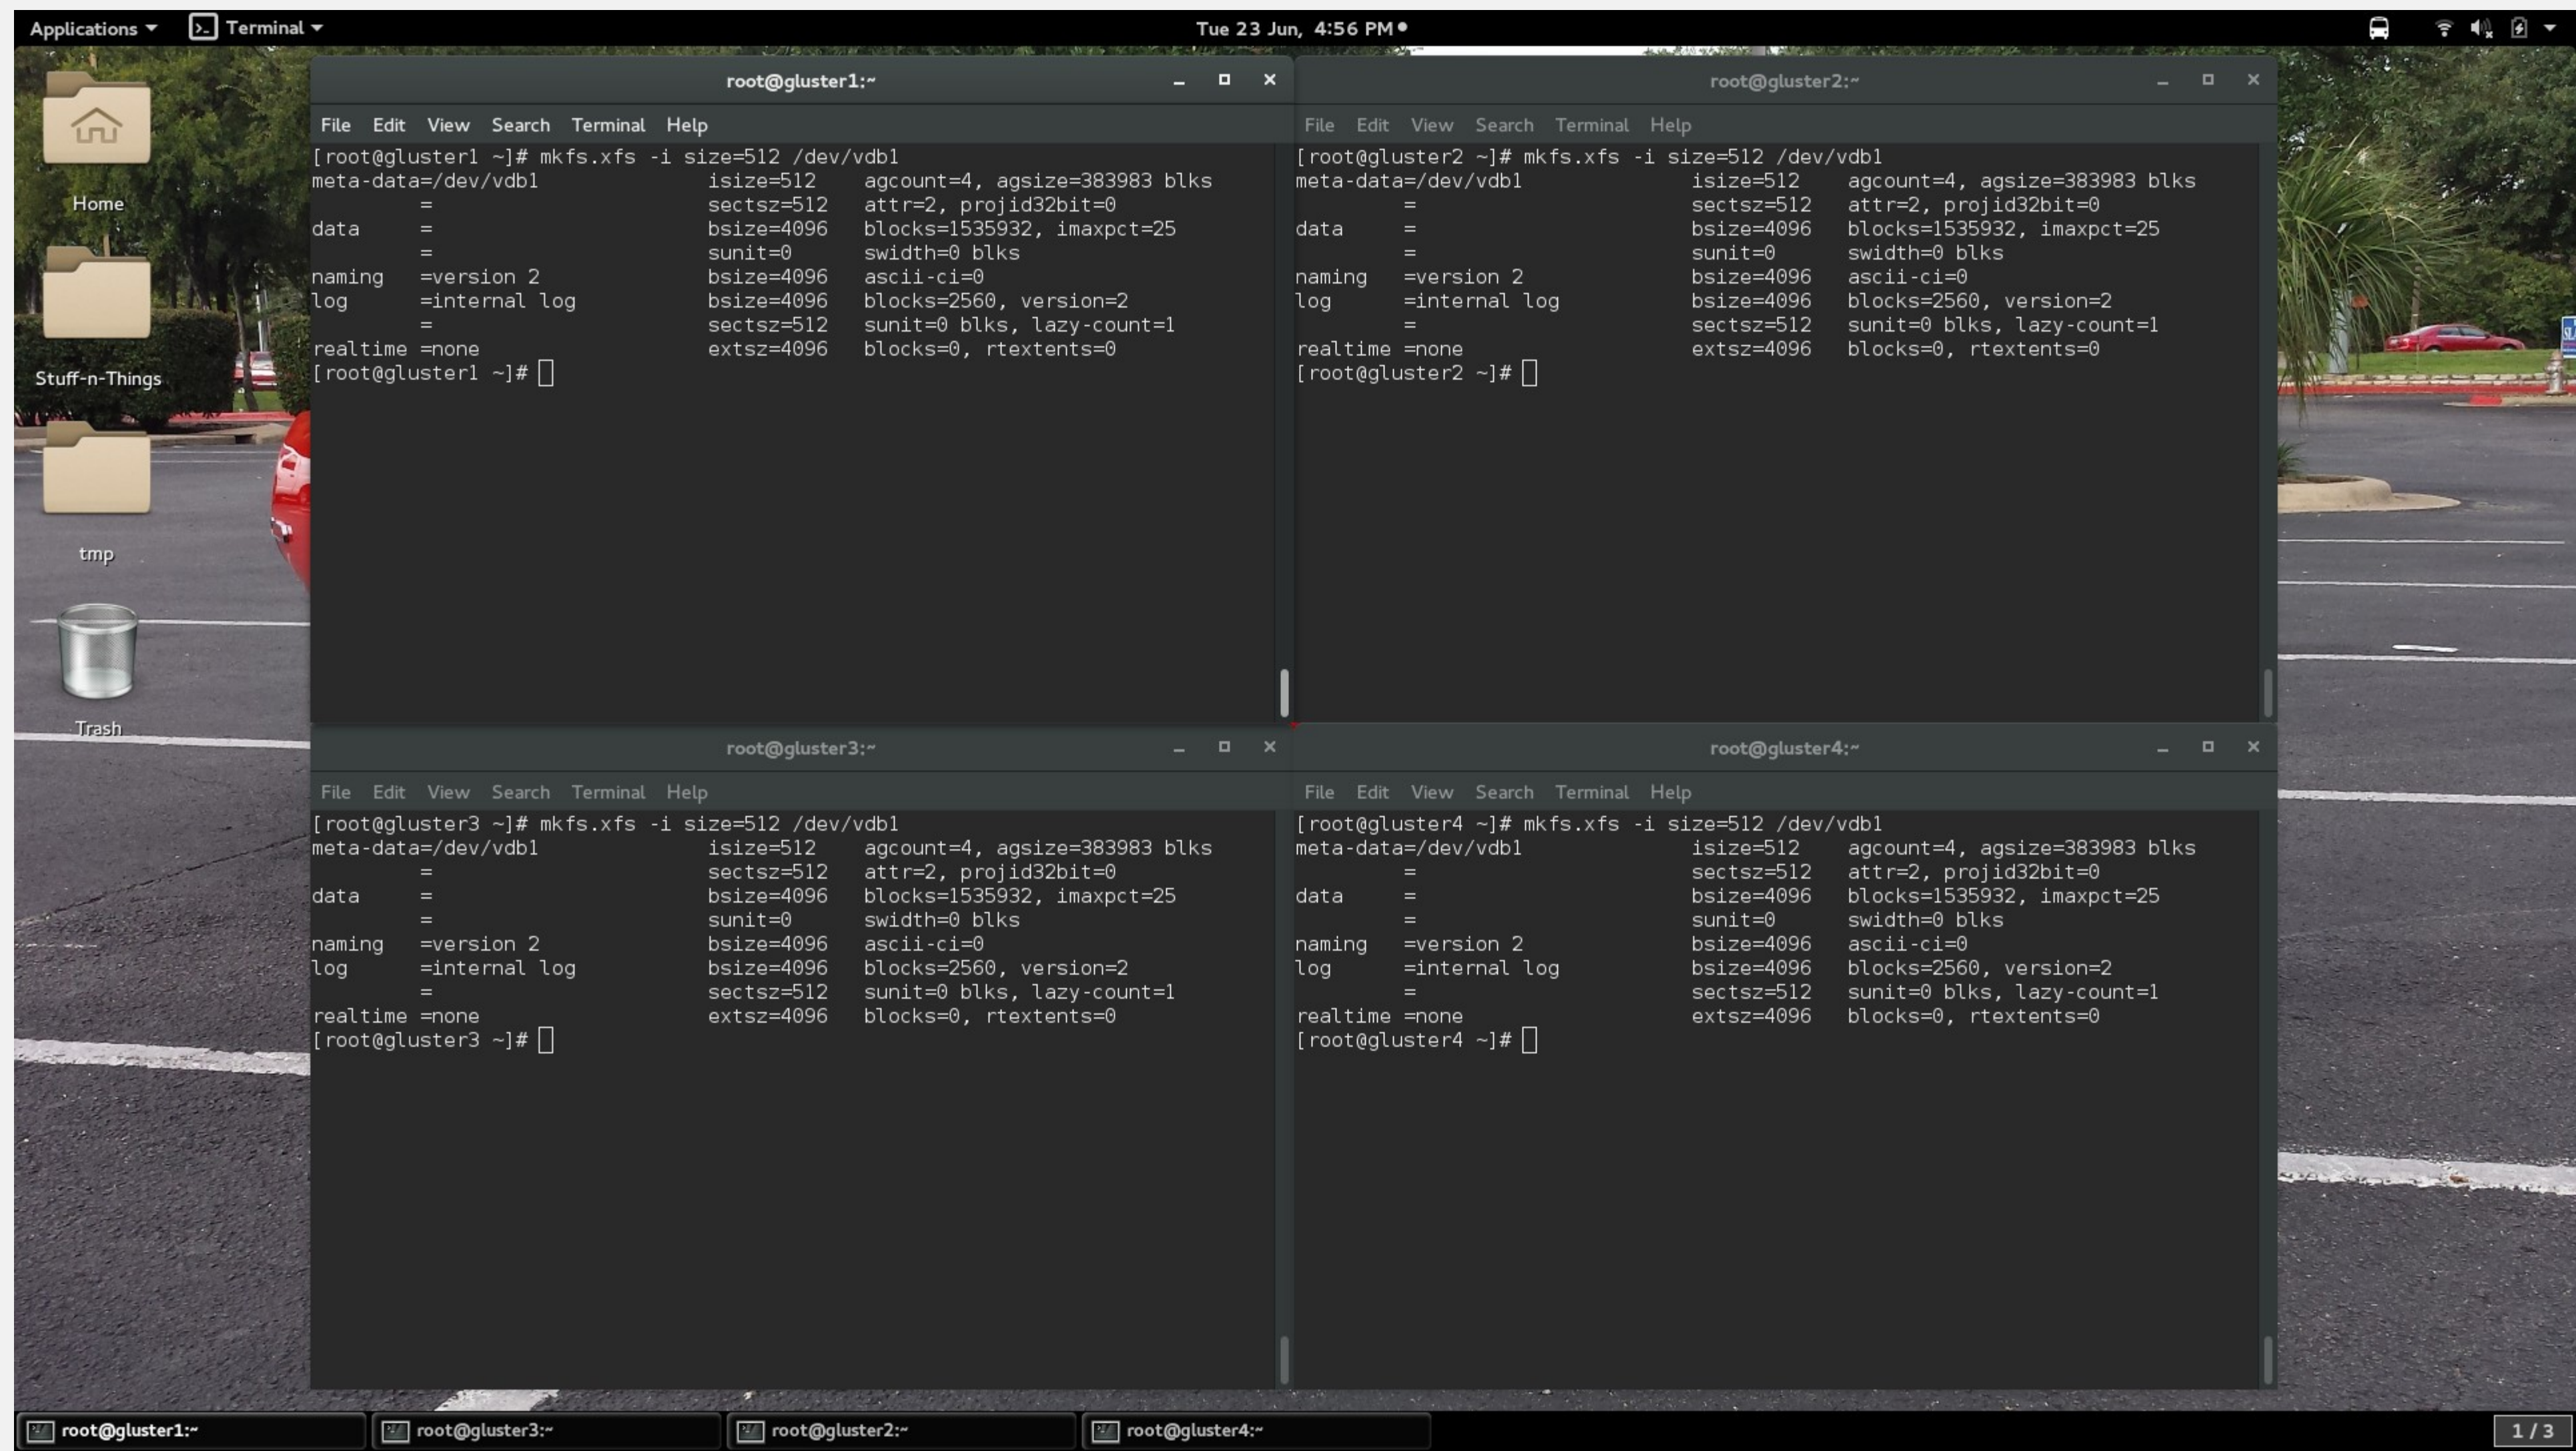



```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# fdisk /dev/vdb  
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel  
Building a new DOS disklabel with disk identifier 0x66e2aa02.  
Changes will remain in memory only, until you decide to write them.  
After that, of course, the previous content won't be recoverable.  
  
Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)  
  
WARNING: DOS-compatible mode is deprecated. It's strongly recommended to  
switch off the mode (command 'c') and change display units to  
sectors (command 'u').  
  
Command (m for help): n  
Command action  
   e   extended  
   p   primary partition (1-4)  
p  
Partition number (1-4): 1  
First cylinder (1-12190, default 1):  
Using default value 1  
Last cylinder, +cylinders or +size{K,M,G} (1-12190, default 12190):  
Using default value 12190  
  
Command (m for help): w  
The partition table has been altered!  
  
Calling ioctl() to re-read partition table.  
Syncing disks.  
[root@gluster1 ~]#
```

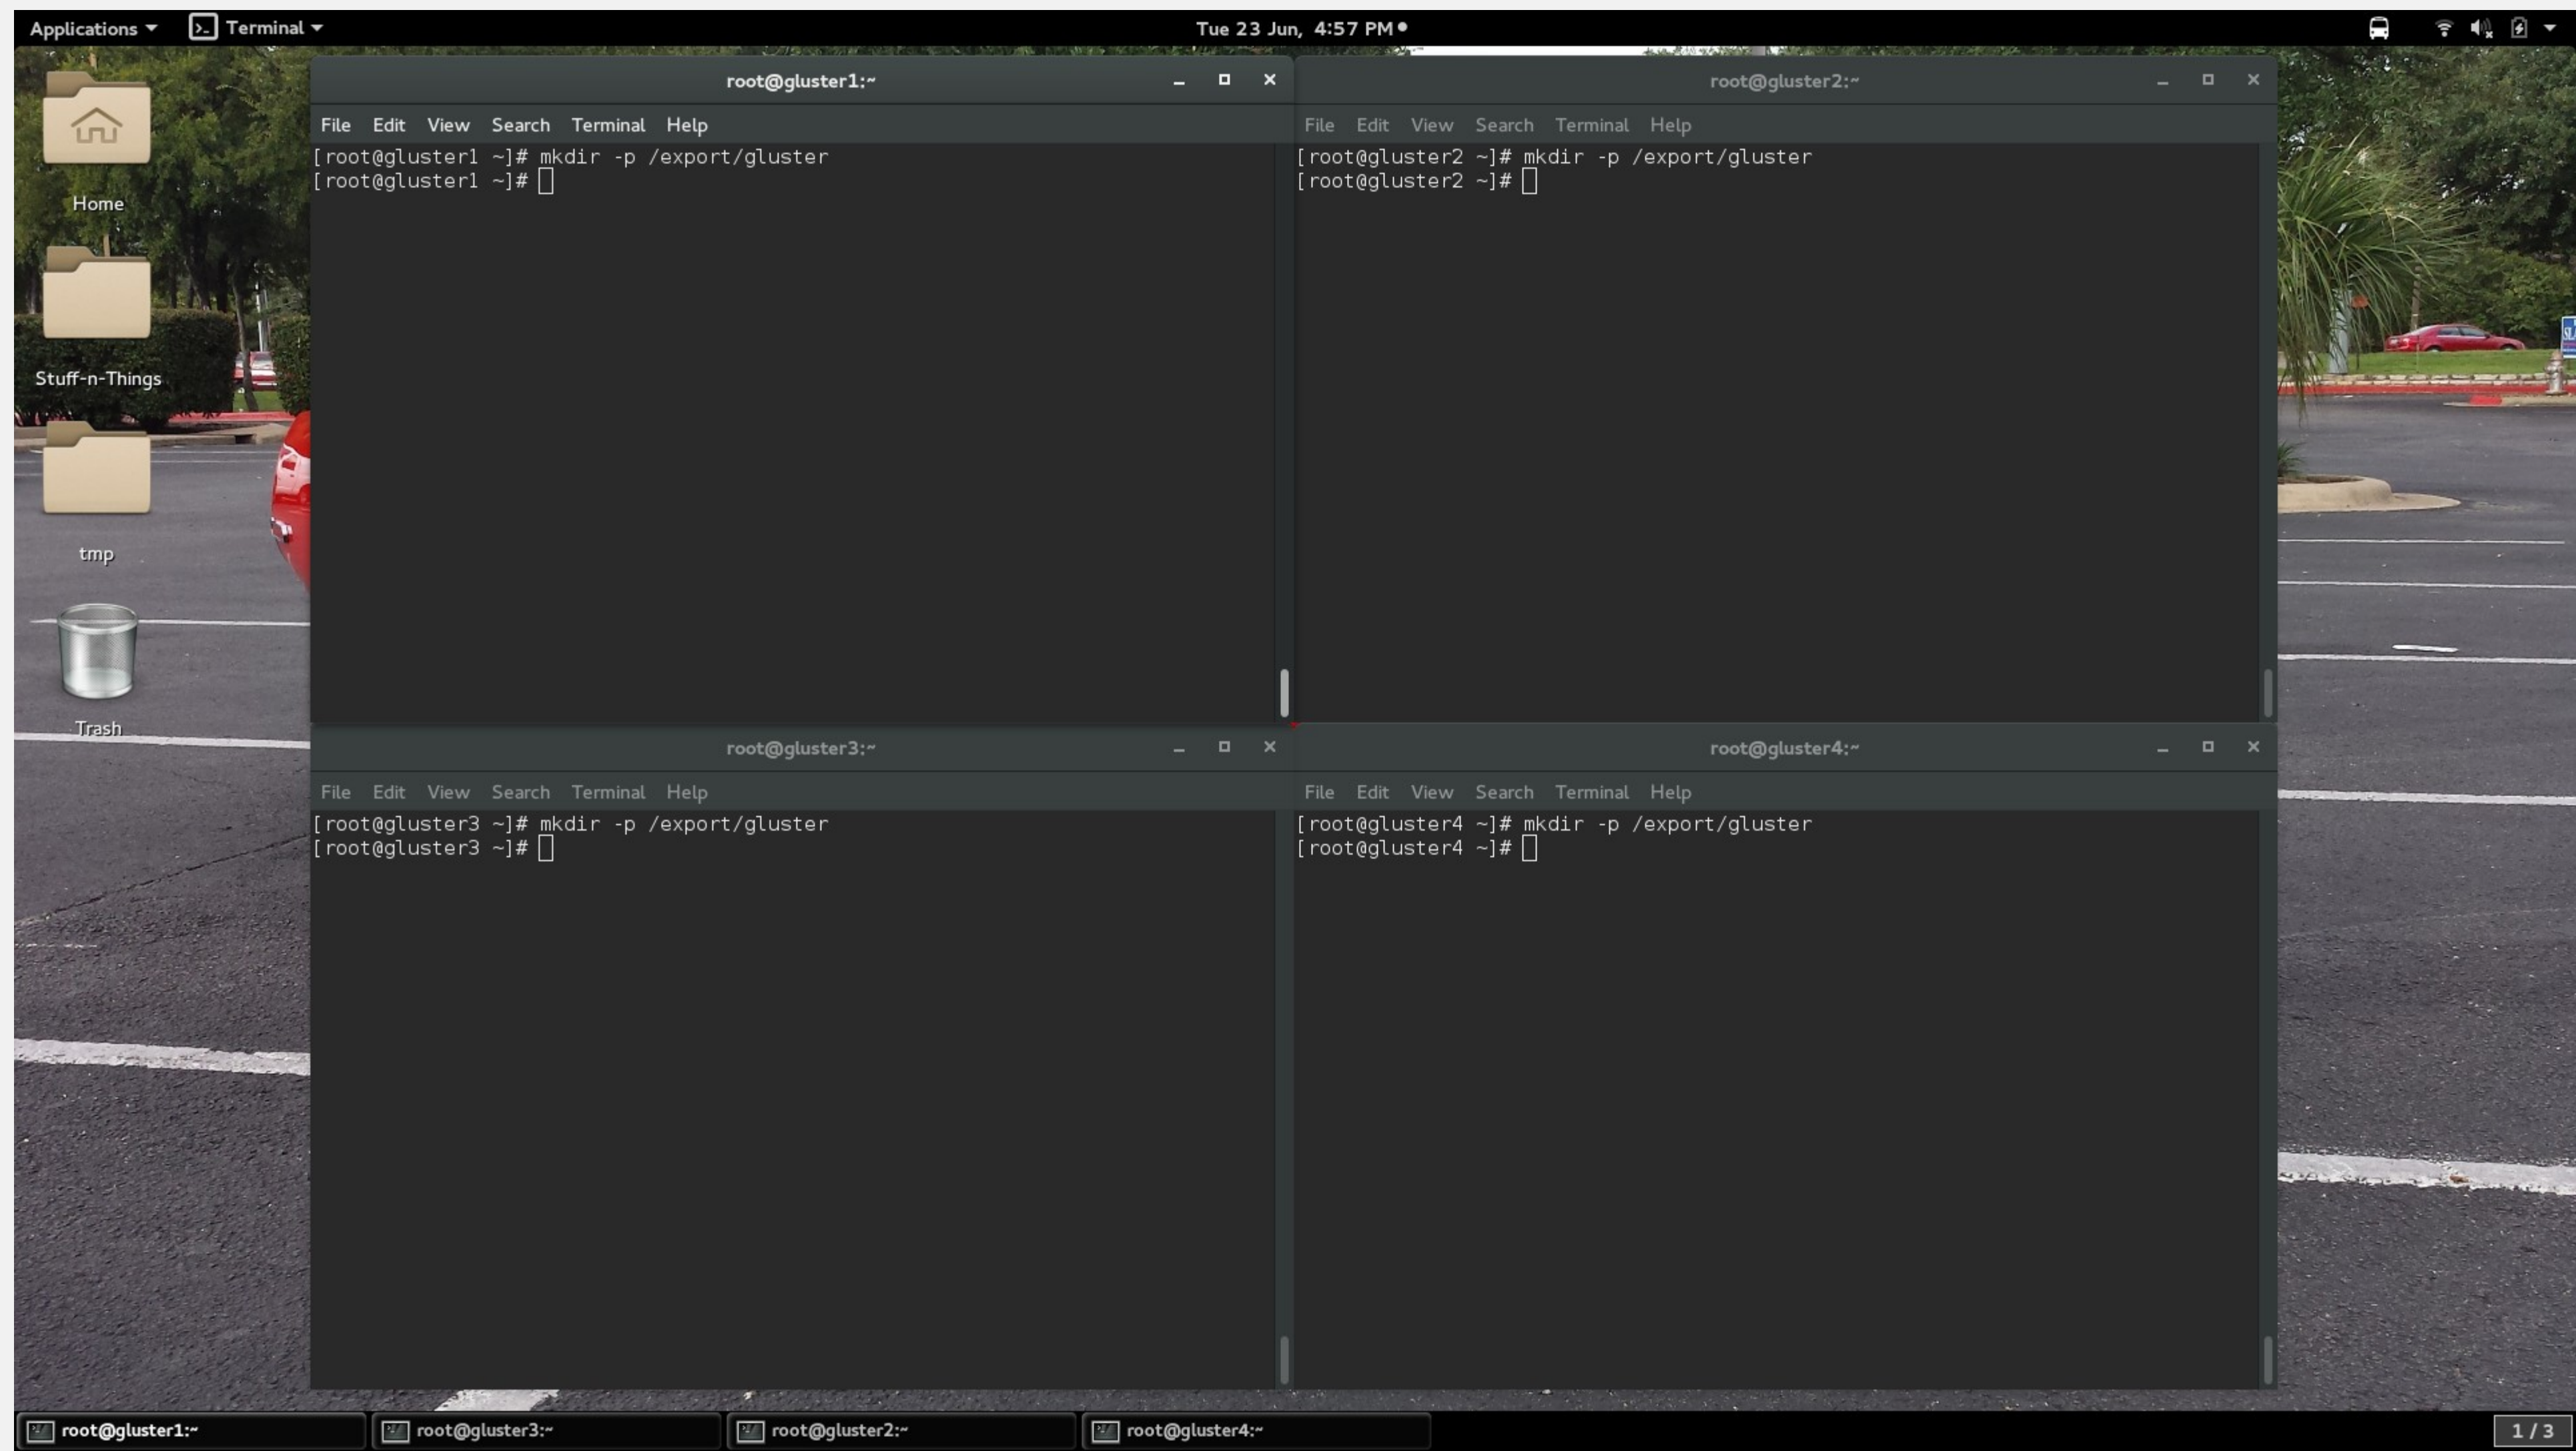

NODE INSTALLATION

- Note the inode option “size” is set to 512 bytes, not the default 256!

```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# mkfs.xfs -i size=512 /dev/vdb1  
meta-data=/dev/vdb1            isize=512    agcount=4, agsize=383983 blks  
      =                       sectsz=512   attr=2, projid32bit=0  
data     =                       bsize=4096  blocks=1535932, imaxpct=25  
      =                       sunit=0      swidth=0 blks  
naming   =version 2           bsize=4096  ascii-ci=0  
log      =internal log       bsize=4096  blocks=2560, version=2  
      =                       sectsz=512   sunit=0 blks, lazy-count=1  
realtime =none               extsz=4096  blocks=0, rtextents=0  
[root@gluster1 ~]#
```

```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# mkdir -p /export/gluster  
[root@gluster1 ~]#
```

```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# mount /dev/vdb1 /export/gluster/  
[root@gluster1 ~]# df -h  
Filesystem      Size  Used Avail Use% Mounted on  
/dev/vda3        4.9G  1.9G  2.8G  40% /  
tmpfs            939M   0  939M   0% /dev/shm  
/dev/vda1        485M   38M  422M   9% /boot  
/dev/vdb1        5.9G   33M  5.9G   1% /export/gluster  
[root@gluster1 ~]#
```


Applications

Terminal

Tue 23 Jun, 4:58 PM

root@gluster1:~

File Edit View Search Terminal Help

[root@gluster1 ~]# mount /dev/vdb1 /export/gluster/

[root@gluster1 ~]# df -h

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/vda3	4.9G	1.9G	2.8G	40%	/
tmpfs	939M	0	939M	0%	/dev/shm
/dev/vda1	485M	38M	422M	9%	/boot
/dev/vdb1	5.9G	33M	5.9G	1%	/export/gluster

[root@gluster1 ~]# df -h

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/vda3	4.9G	1.9G	2.8G	40%	/
tmpfs	939M	0	939M	0%	/dev/shm
/dev/vda1	485M	38M	422M	9%	/boot
/dev/vdb1	5.9G	33M	5.9G	1%	/export/gluster

[root@gluster1 ~]#

root@gluster2:~

File Edit View Search Terminal Help

[root@gluster2 ~]# mkdir -p /export/gluster

[root@gluster2 ~]# mount /dev/vdb1 /export/gluster/

[root@gluster2 ~]# df -h

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/vda3	4.9G	1.8G	2.9G	38%	/
tmpfs	939M	0	939M	0%	/dev/shm
/dev/vda1	485M	38M	422M	9%	/boot
/dev/vdb1	5.9G	33M	5.9G	1%	/export/gluster

[root@gluster2 ~]#

root@gluster3:~

File Edit View Search Terminal Help

[root@gluster3 ~]# mkdir -p /export/gluster

[root@gluster3 ~]# mount /dev/vdb1 /export/gluster/

[root@gluster3 ~]# df -h

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/vda3	4.9G	1.7G	2.9G	38%	/
tmpfs	939M	0	939M	0%	/dev/shm
/dev/vda1	485M	38M	422M	9%	/boot
/dev/vdb1	5.9G	33M	5.9G	1%	/export/gluster

[root@gluster3 ~]#

root@gluster4:~

File Edit View Search Terminal Help

[root@gluster4 ~]# mkdir -p /export/gluster

[root@gluster4 ~]# mount /dev/vdb1 /export/gluster/

[root@gluster4 ~]# df -h

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/vda3	4.9G	1.7G	2.9G	38%	/
tmpfs	939M	0	939M	0%	/dev/shm
/dev/vda1	485M	38M	422M	9%	/boot
/dev/vdb1	5.9G	33M	5.9G	1%	/export/gluster

[root@gluster4 ~]#

root@gluster1:~


root@gluster3:~

root@gluster2:~

root@gluster4:~

1 / 3

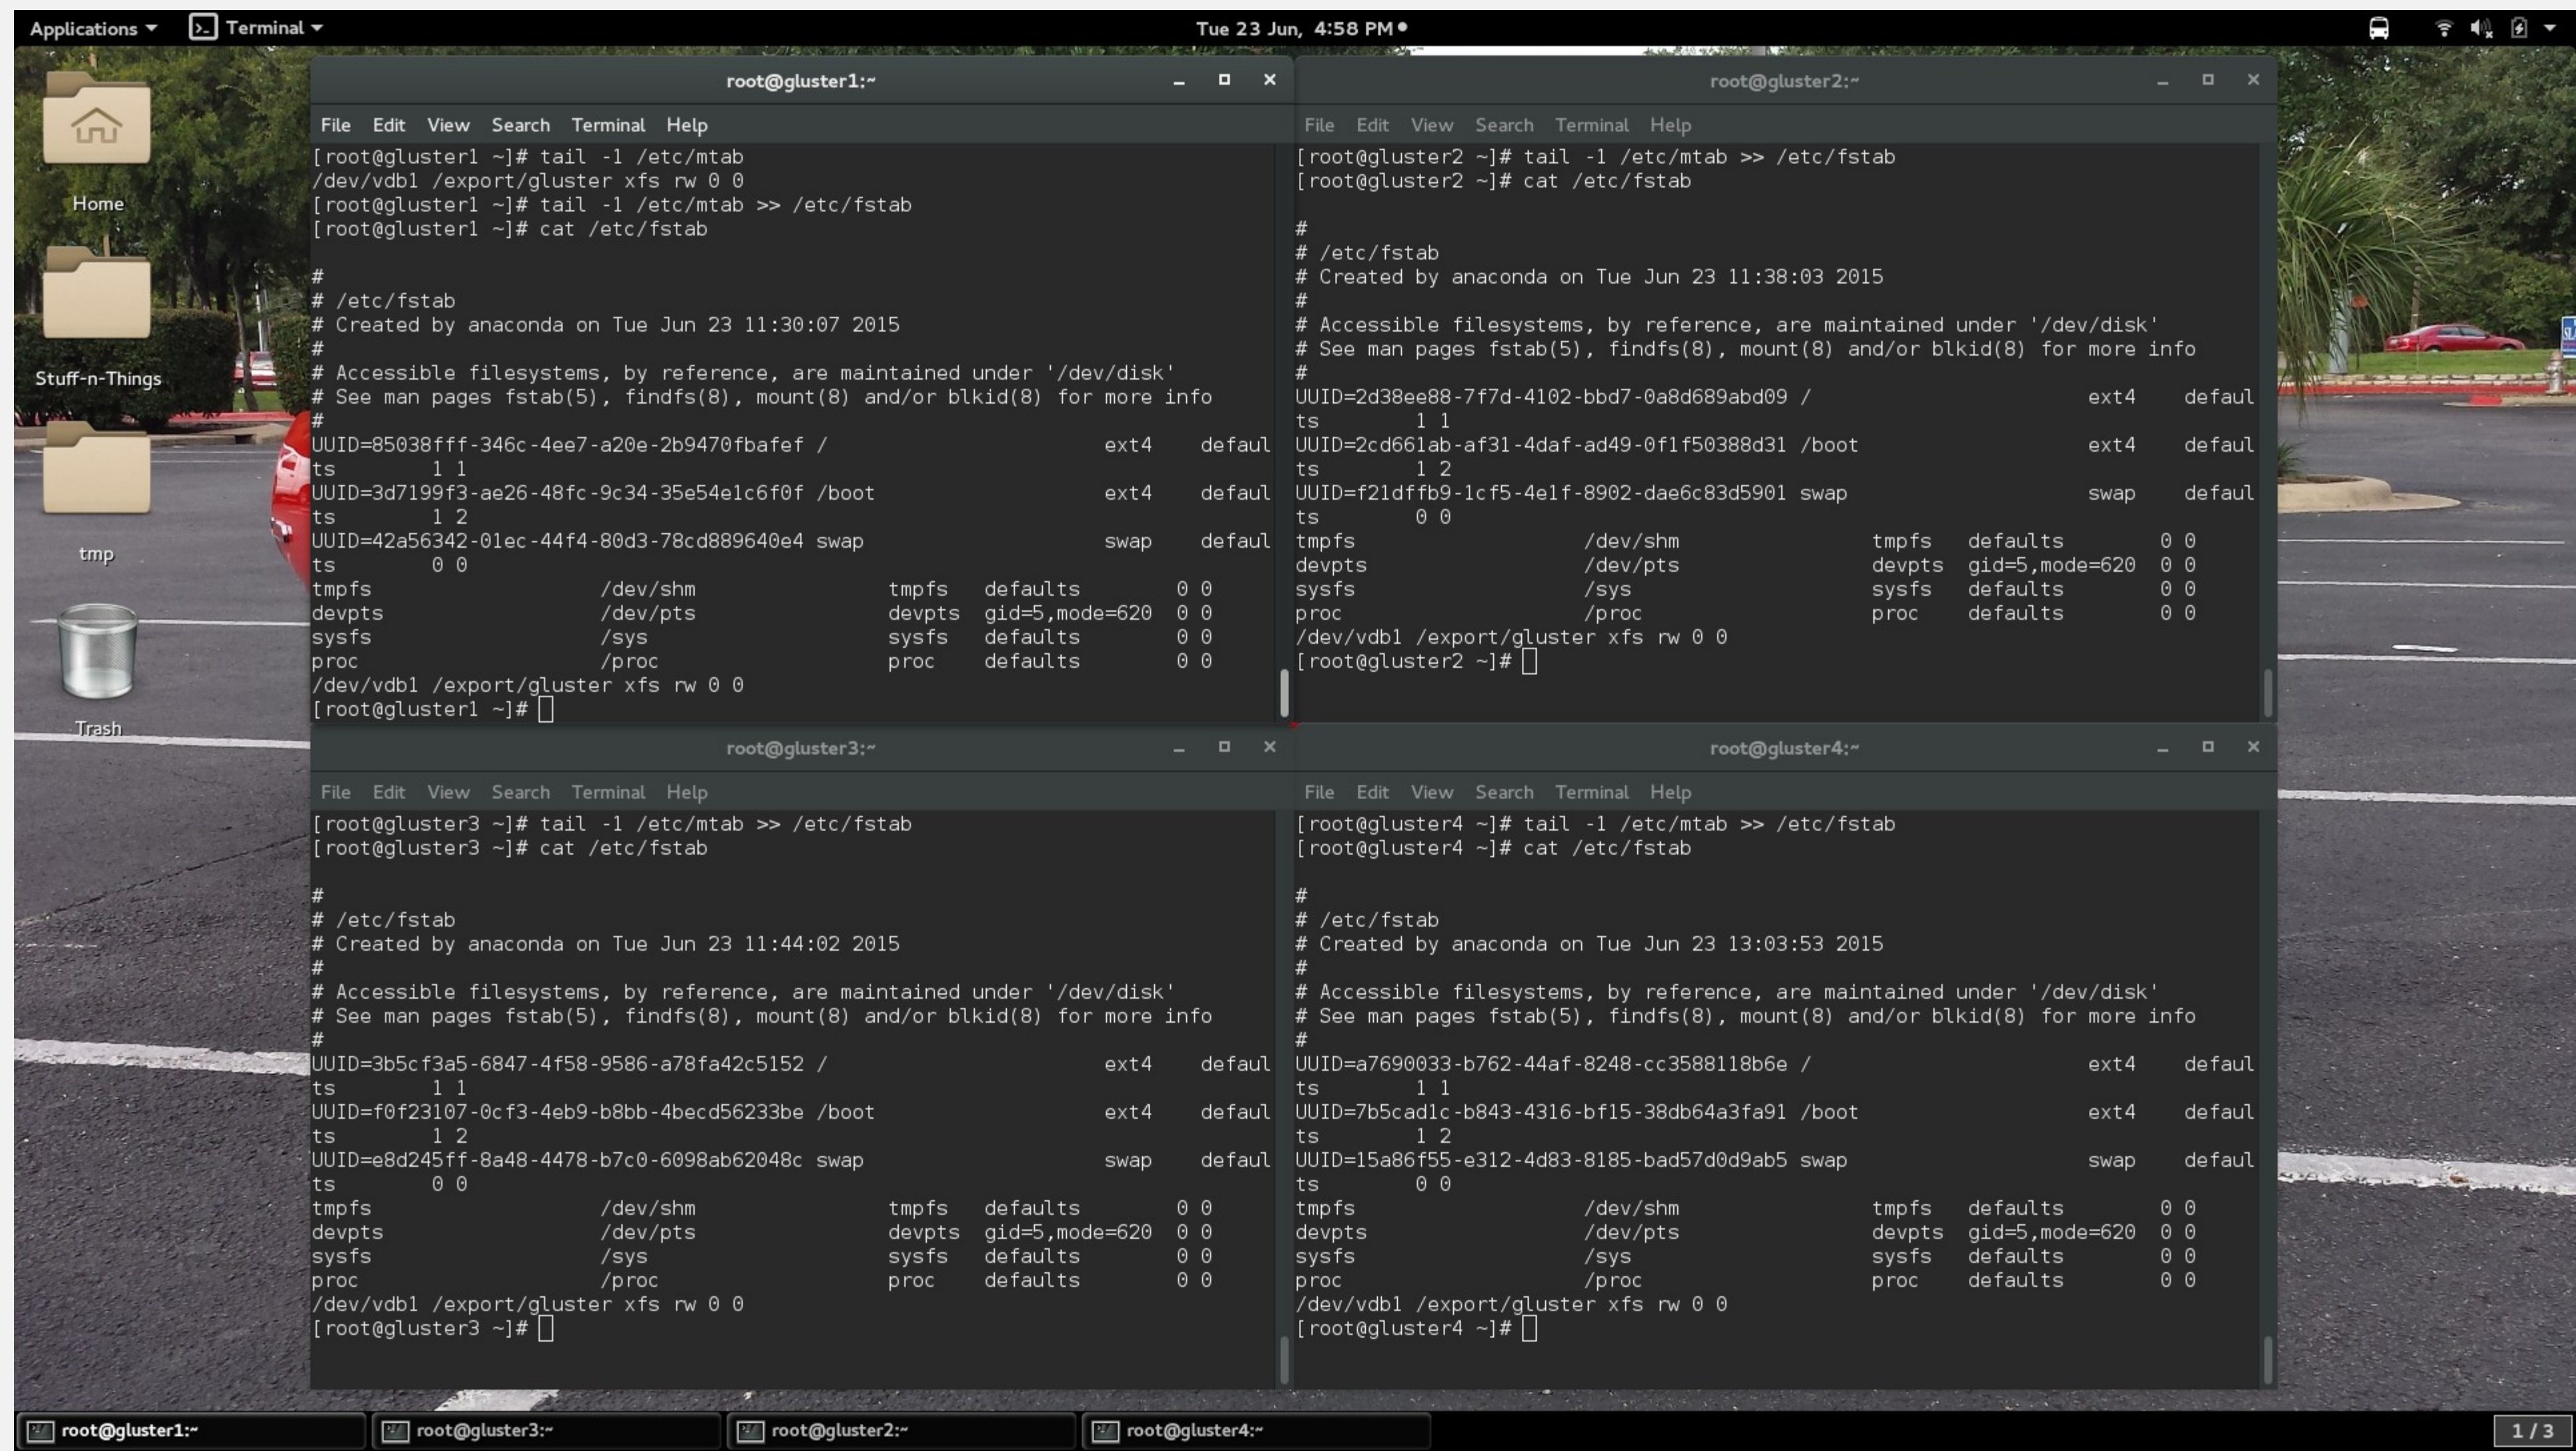
#redhat #rhsummit



NODE INSTALLATION

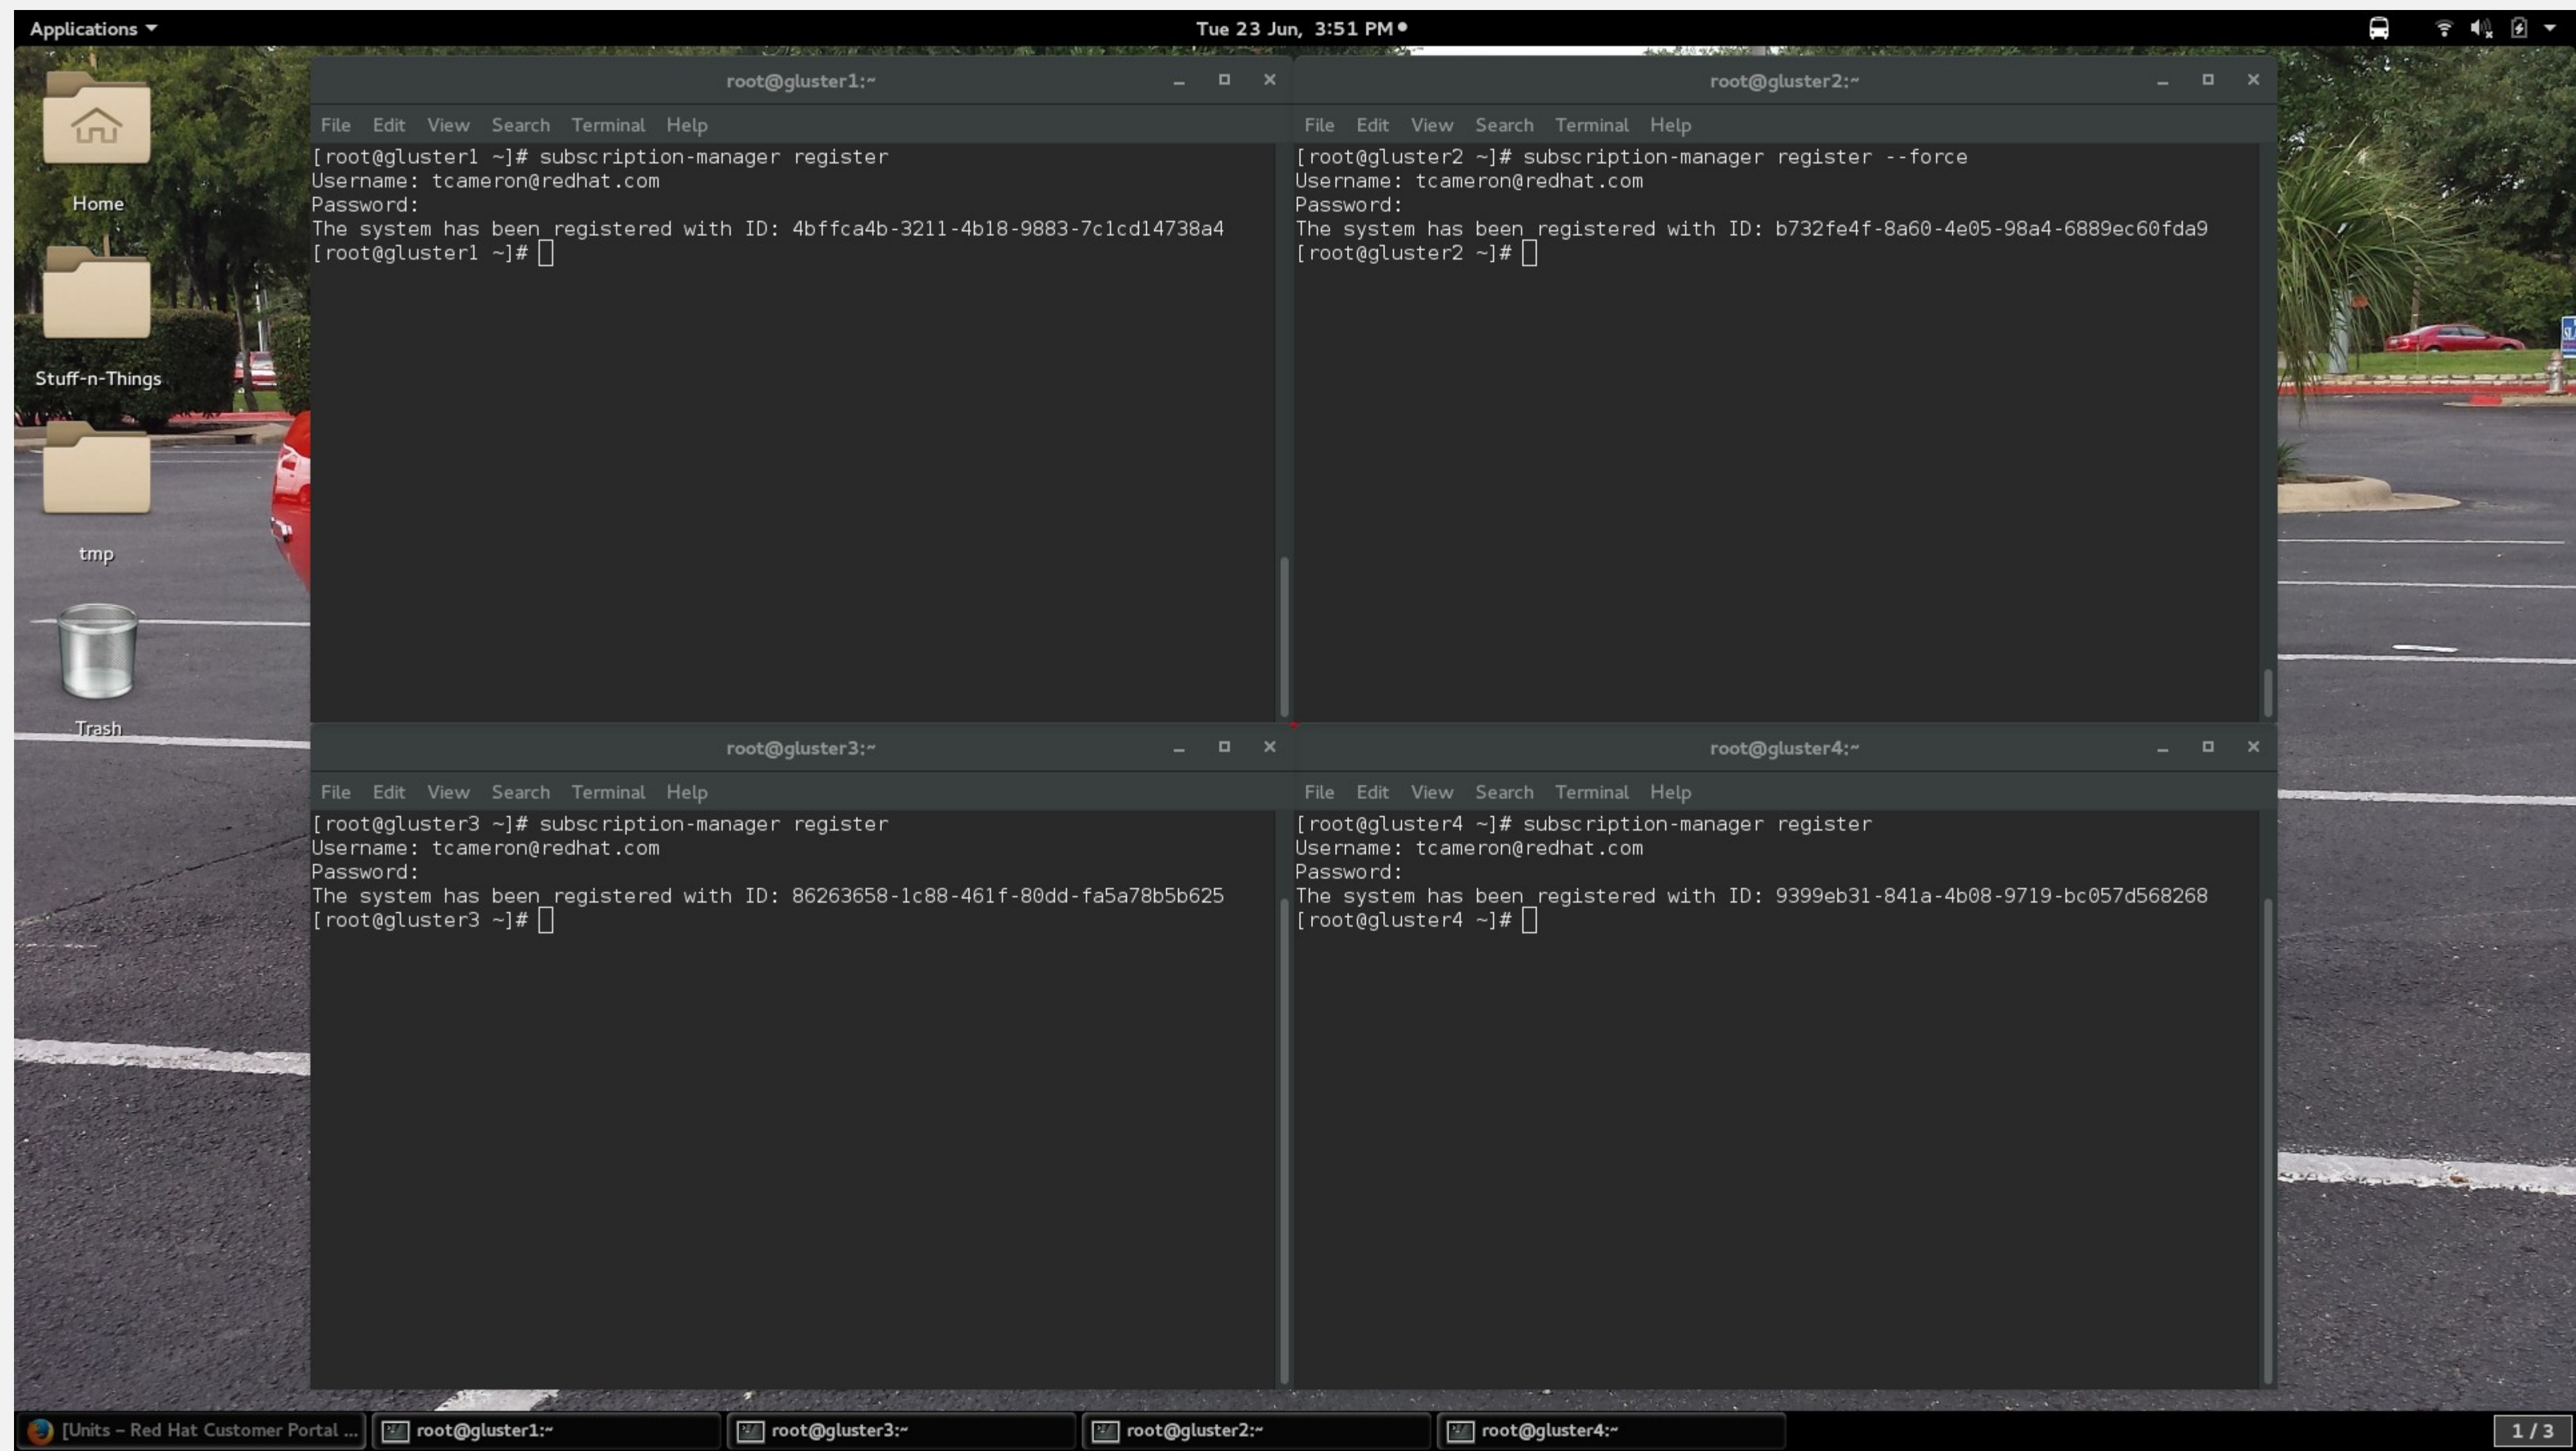
- Once you've mounted the filesystem, you can “cheat” and get the correct syntax for `/etc/fstab` from `/etc/mtab`.

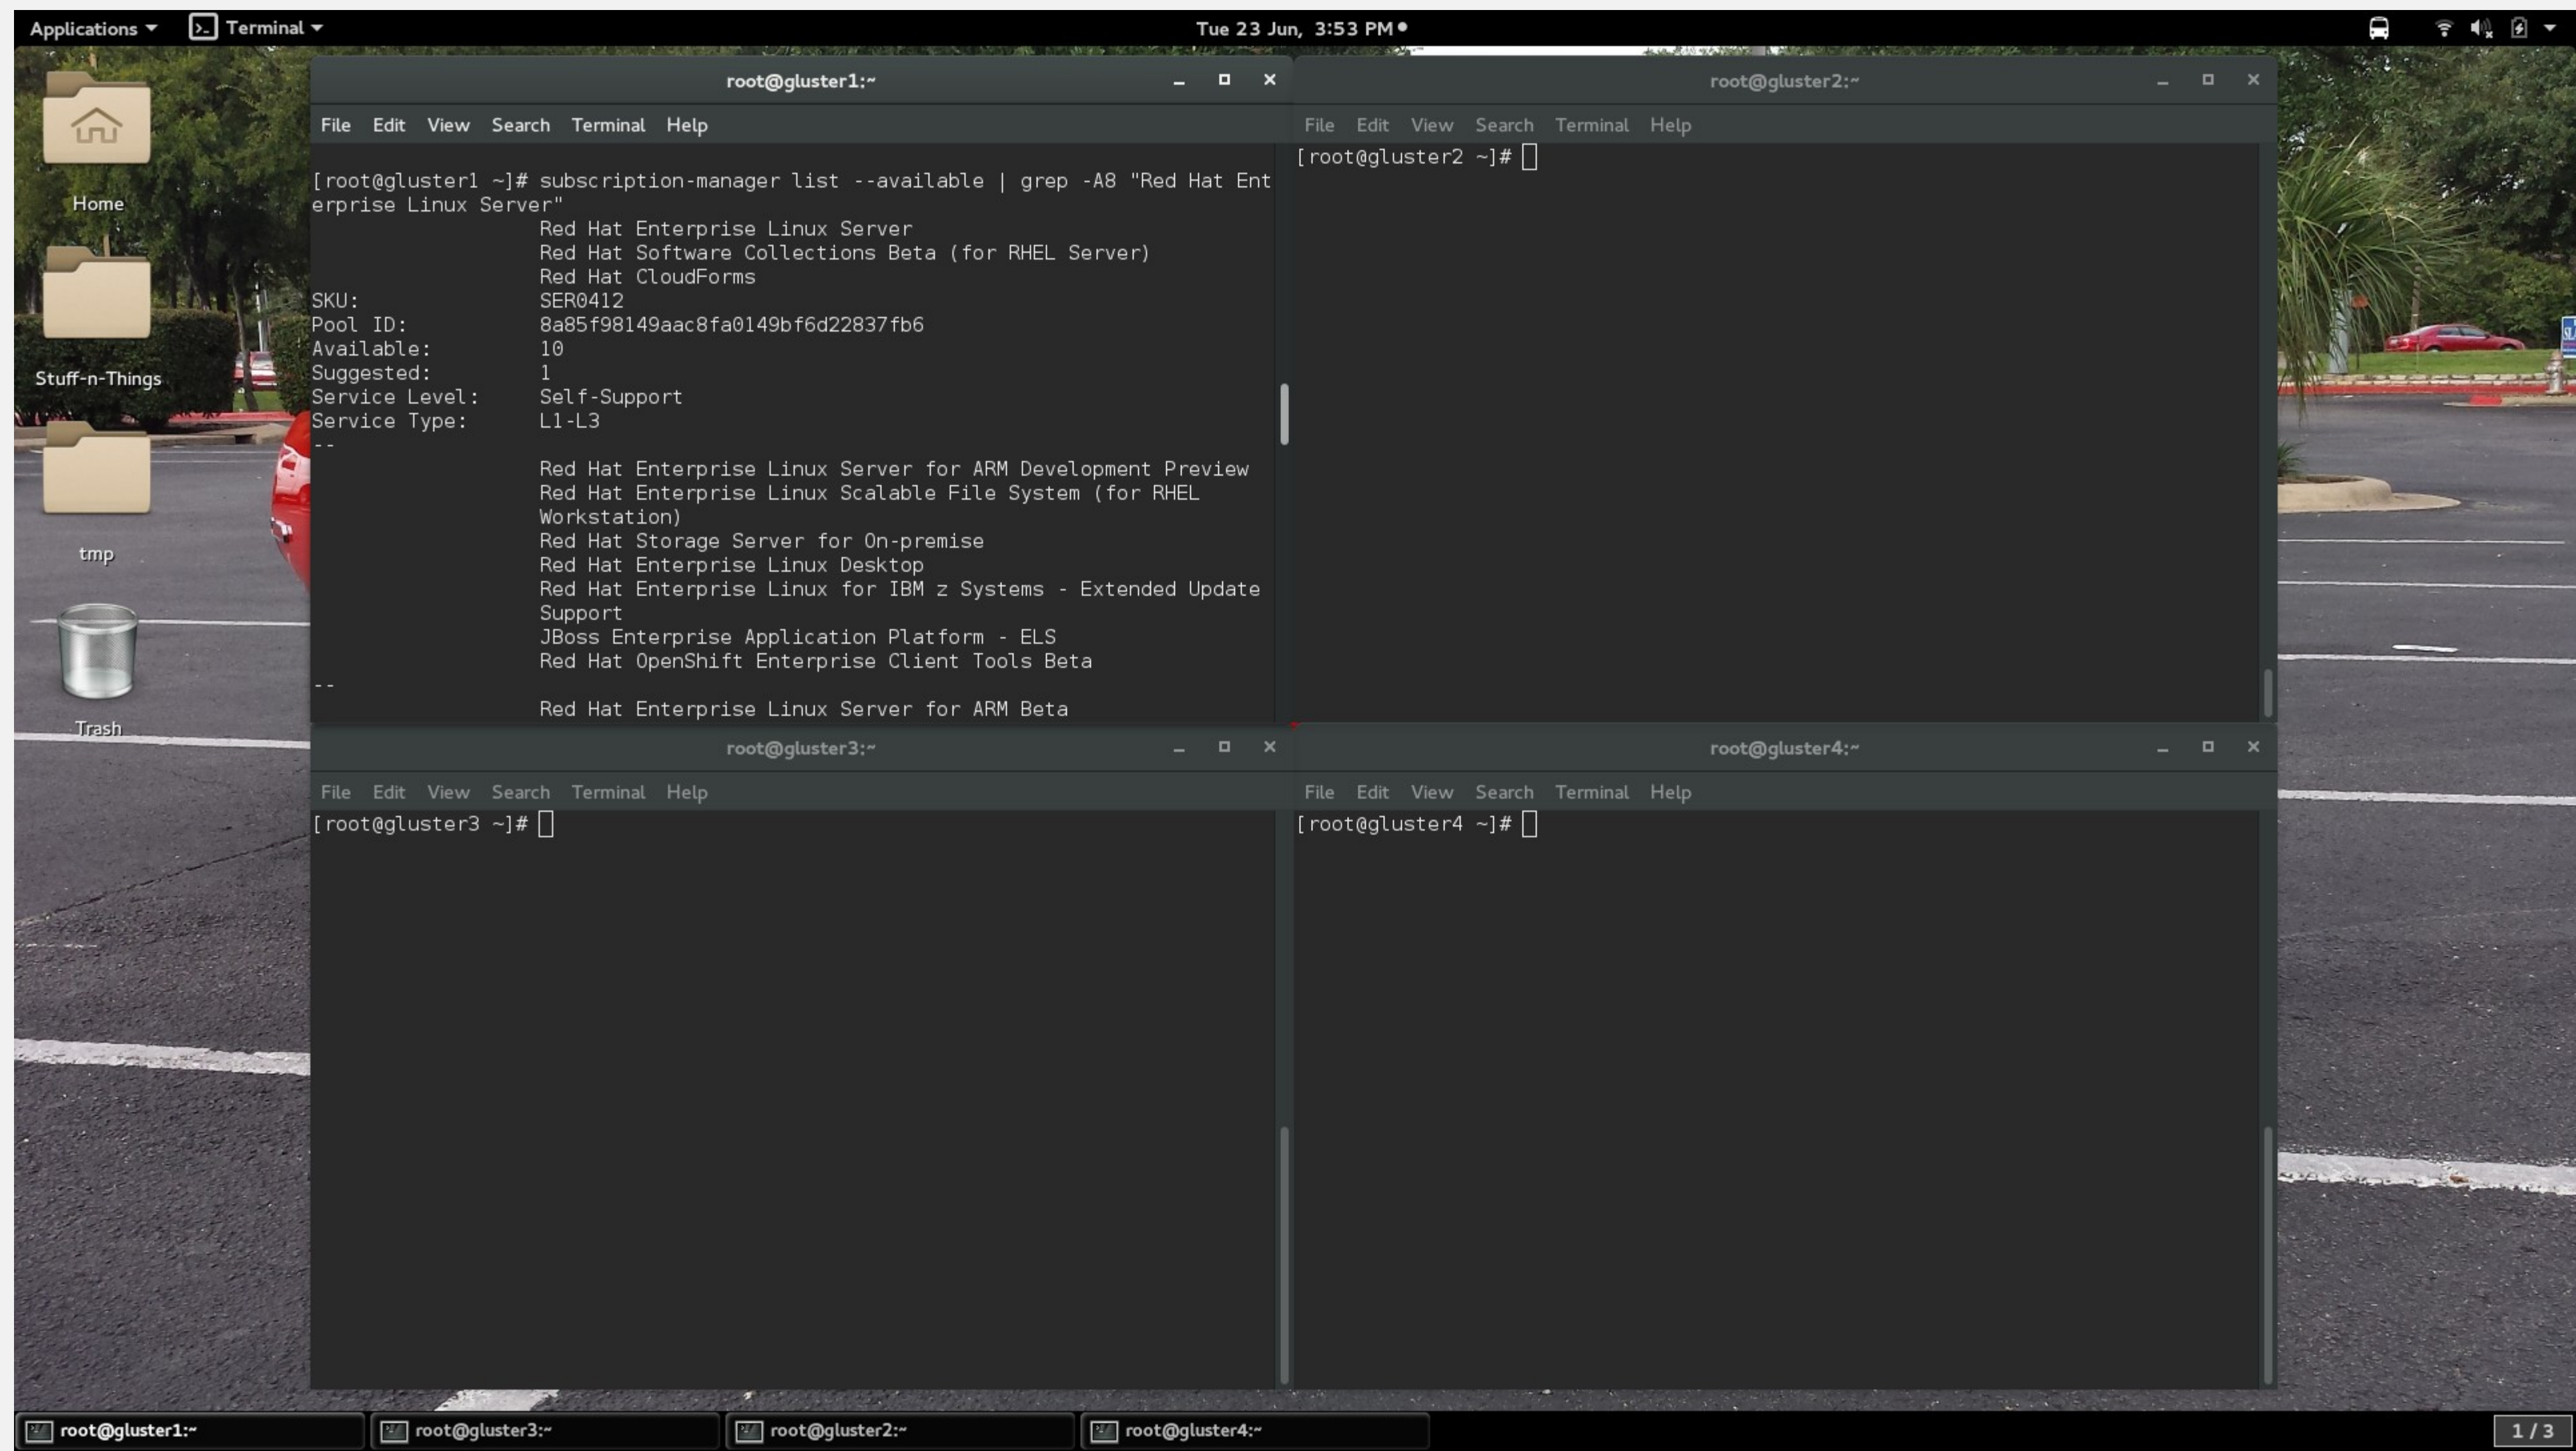

```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# tail -1 /etc/mtab  
/dev/vdb1 /export/gluster xfs rw 0 0  
[root@gluster1 ~]# tail -1 /etc/mtab >> /etc/fstab  
[root@gluster1 ~]# cat /etc/fstab  
  
#  
# /etc/fstab  
# Created by anaconda on Tue Jun 23 11:30:07 2015  
#  
# Accessible filesystems, by reference, are maintained under '/dev/disk'  
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info  
#  
UUID=85038fff-346c-4ee7-a20e-2b9470fbafef / ext4 default  
ts 1 1  
UUID=3d7199f3-ae26-48fc-9c34-35e54e1c6f0f /boot ext4 default  
ts 1 2  
UUID=42a56342-01ec-44f4-80d3-78cd889640e4 swap swap default  
ts 0 0  
tmpfs /dev/shm tmpfs defaults 0 0  
devpts /dev/pts devpts gid=5,mode=620 0 0  
sysfs /sys sysfs defaults 0 0  
proc /proc proc defaults 0 0  
/dev/vdb1 /export/gluster xfs rw 0 0  
[root@gluster1 ~]#
```

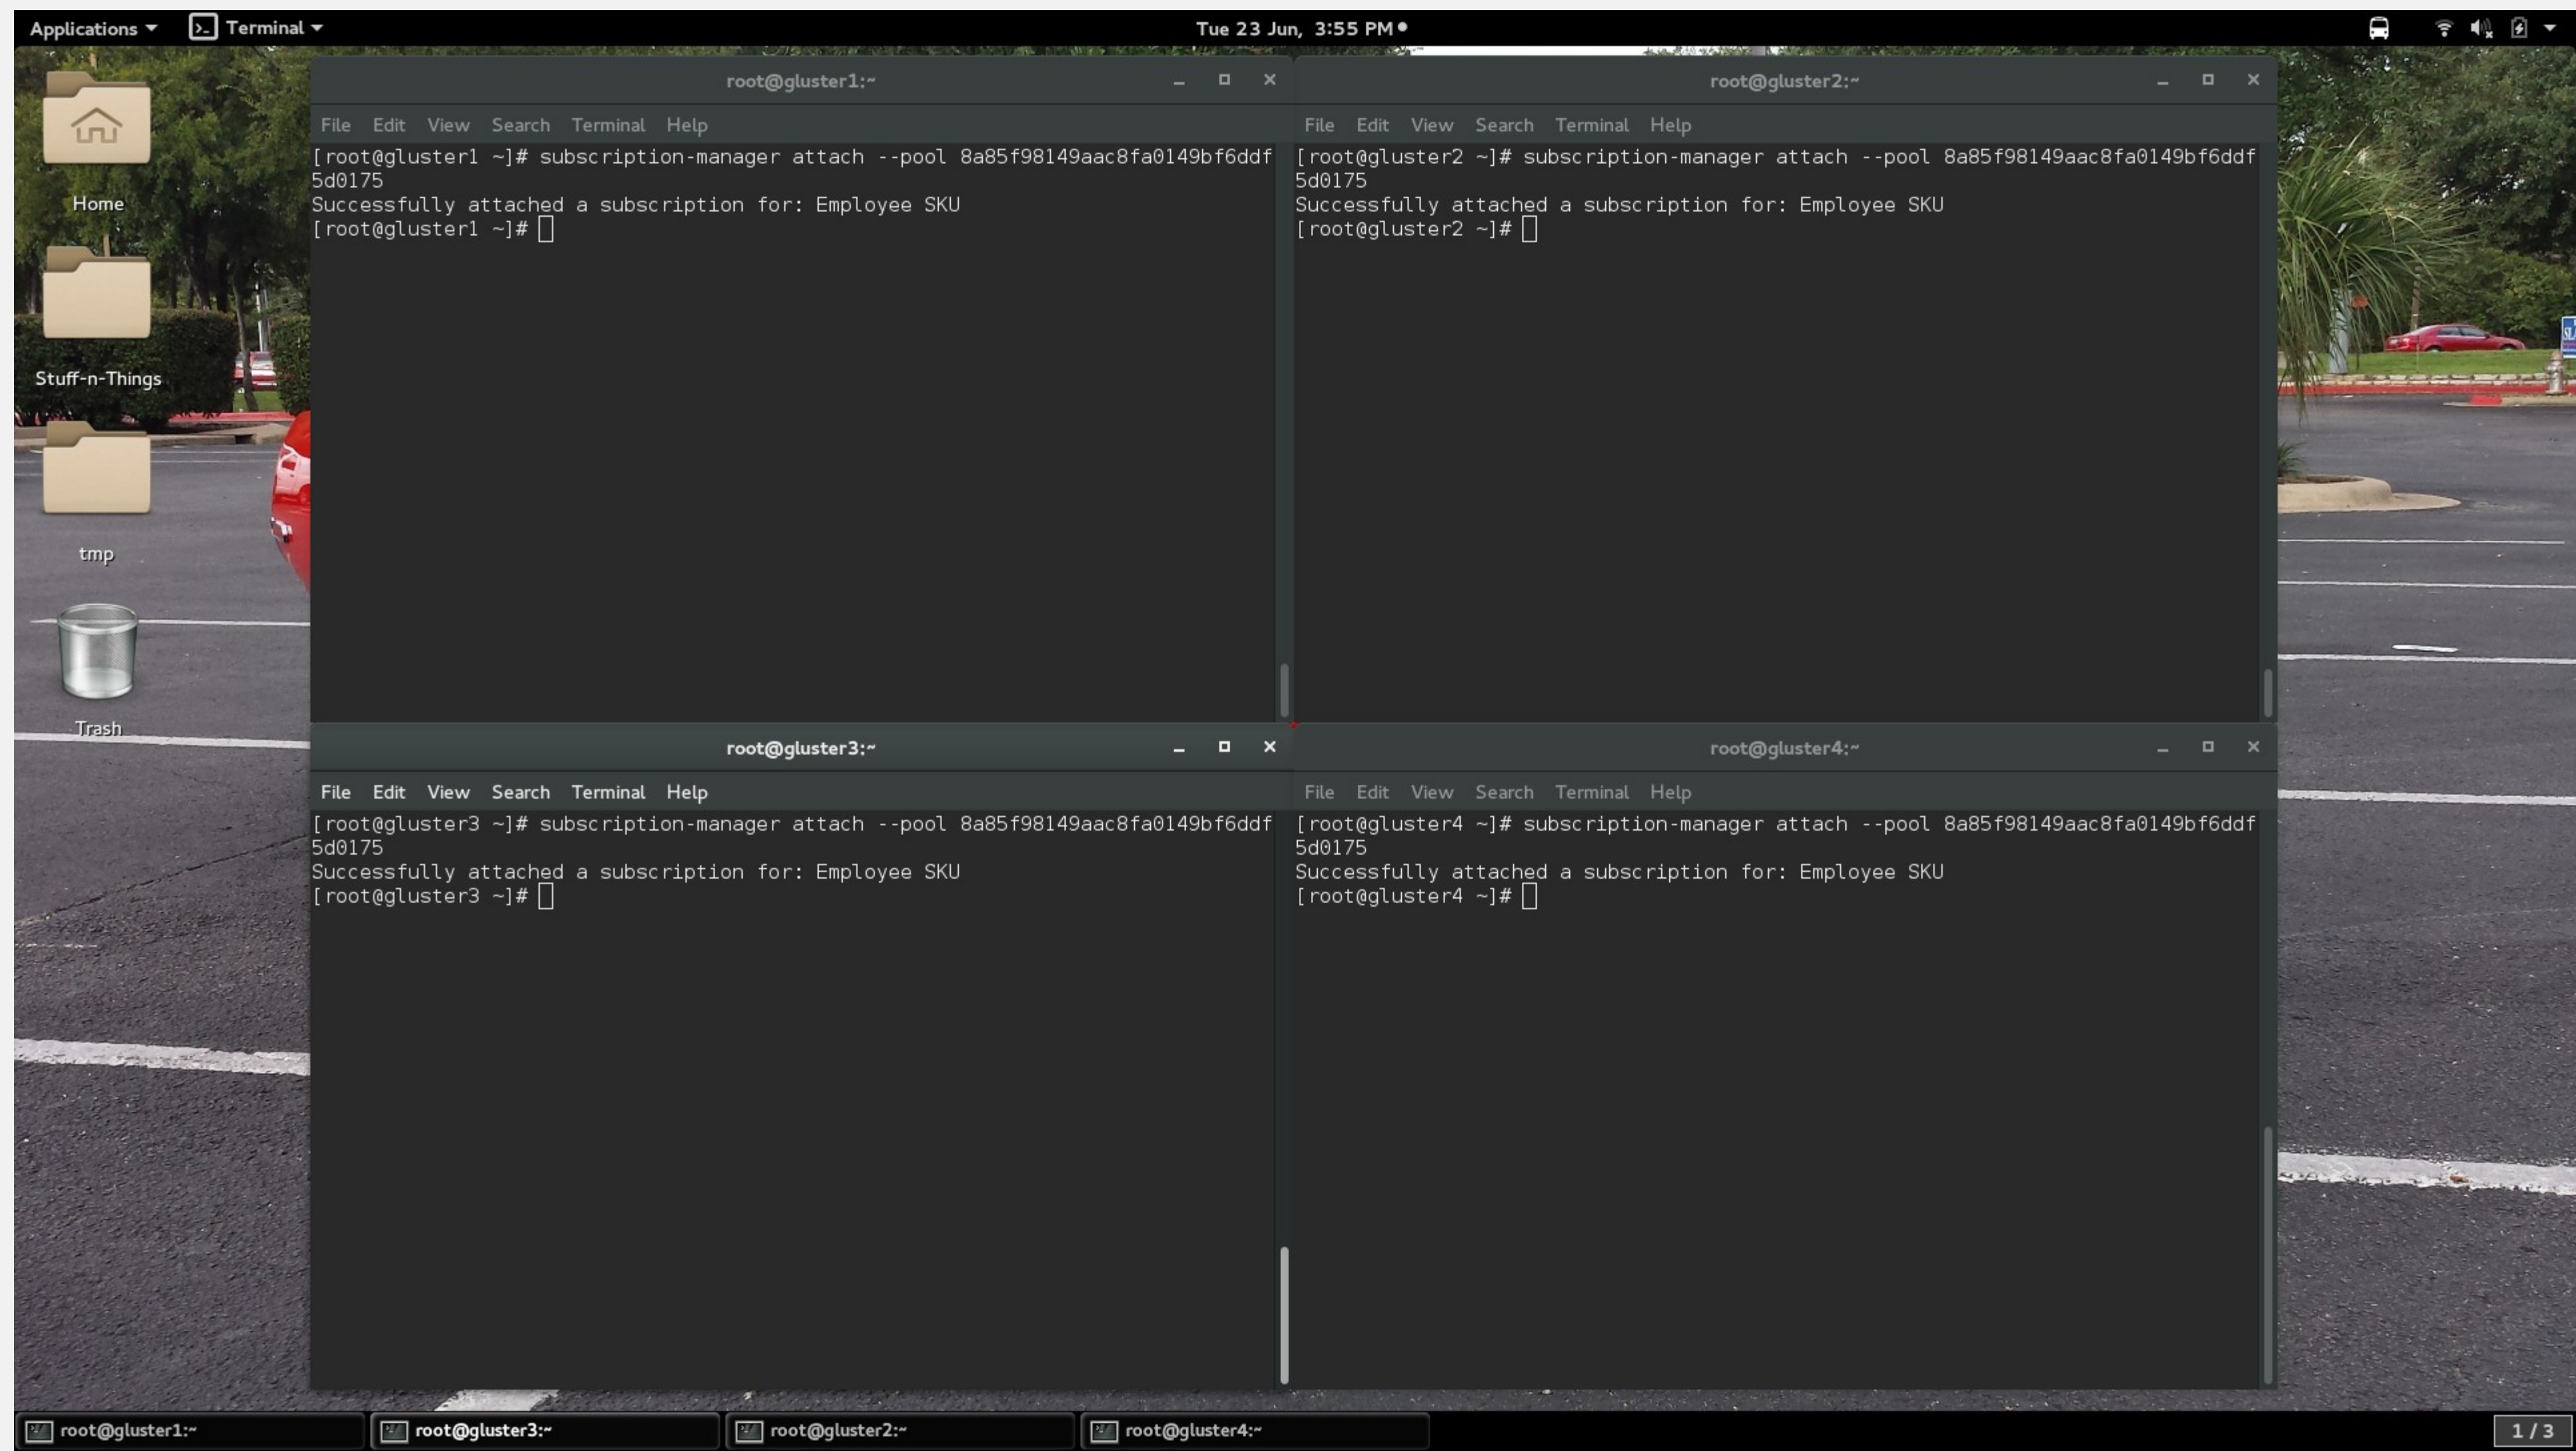



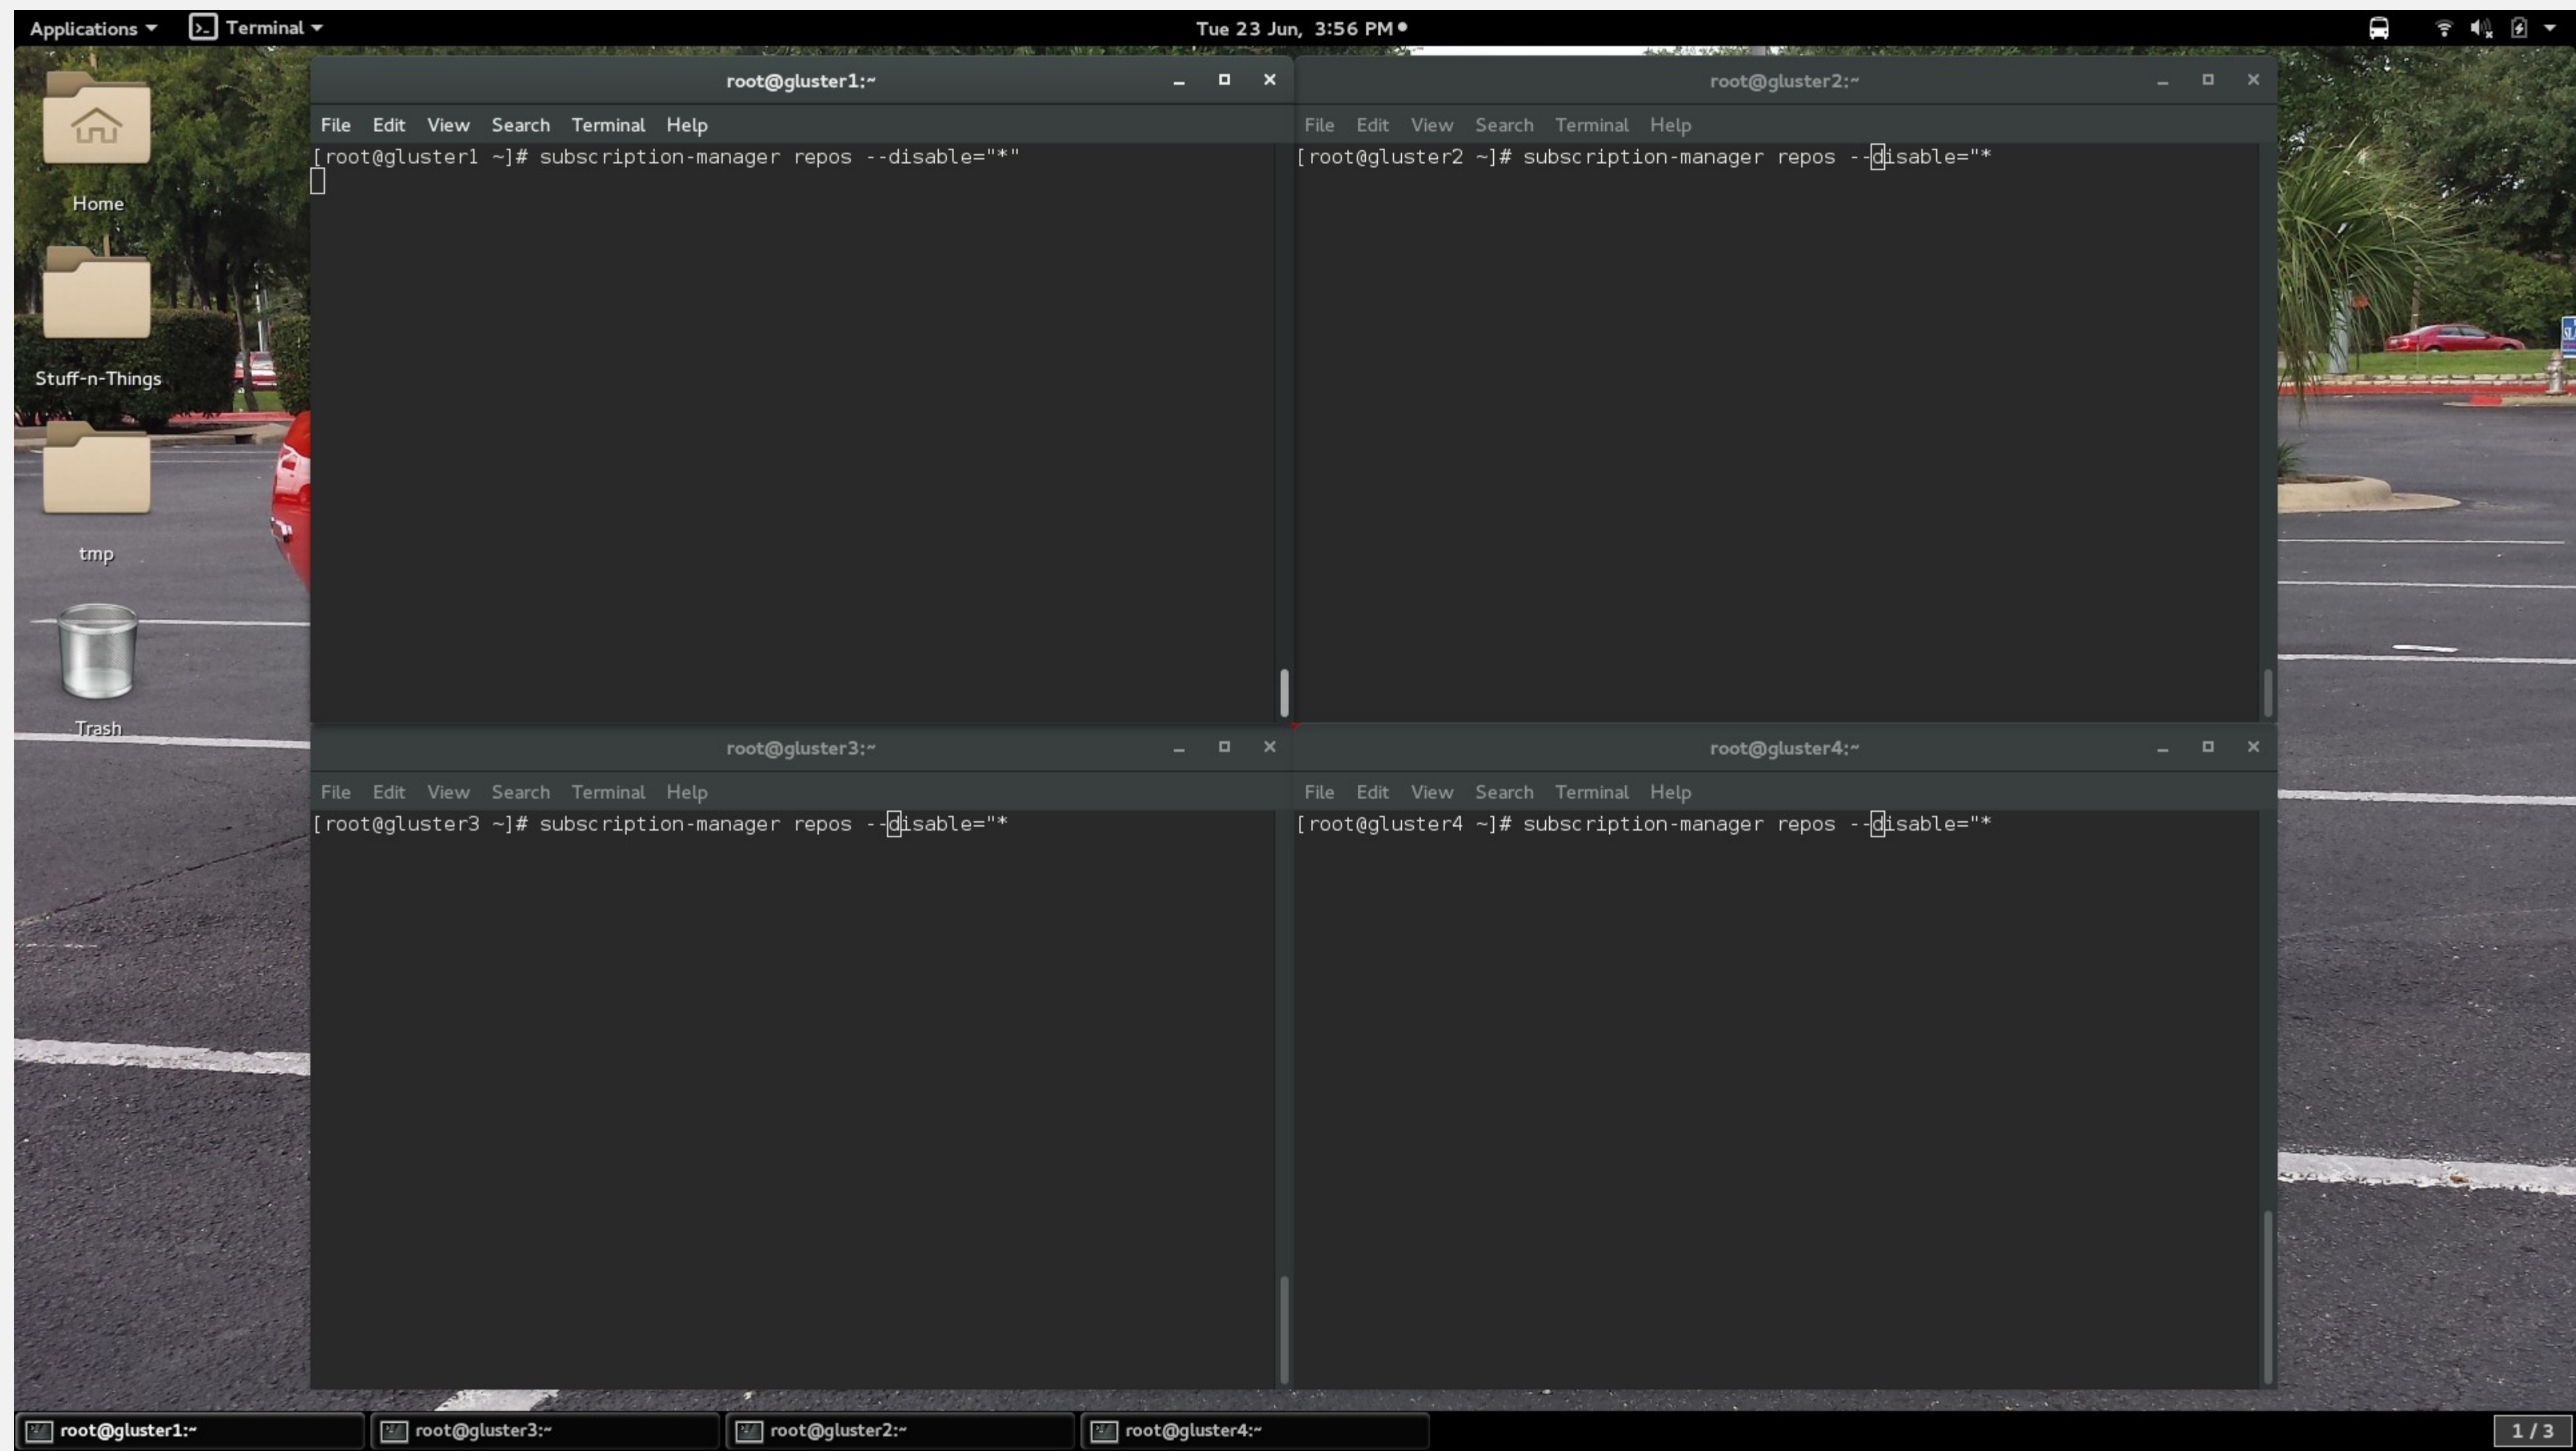
NODE INSTALLATION

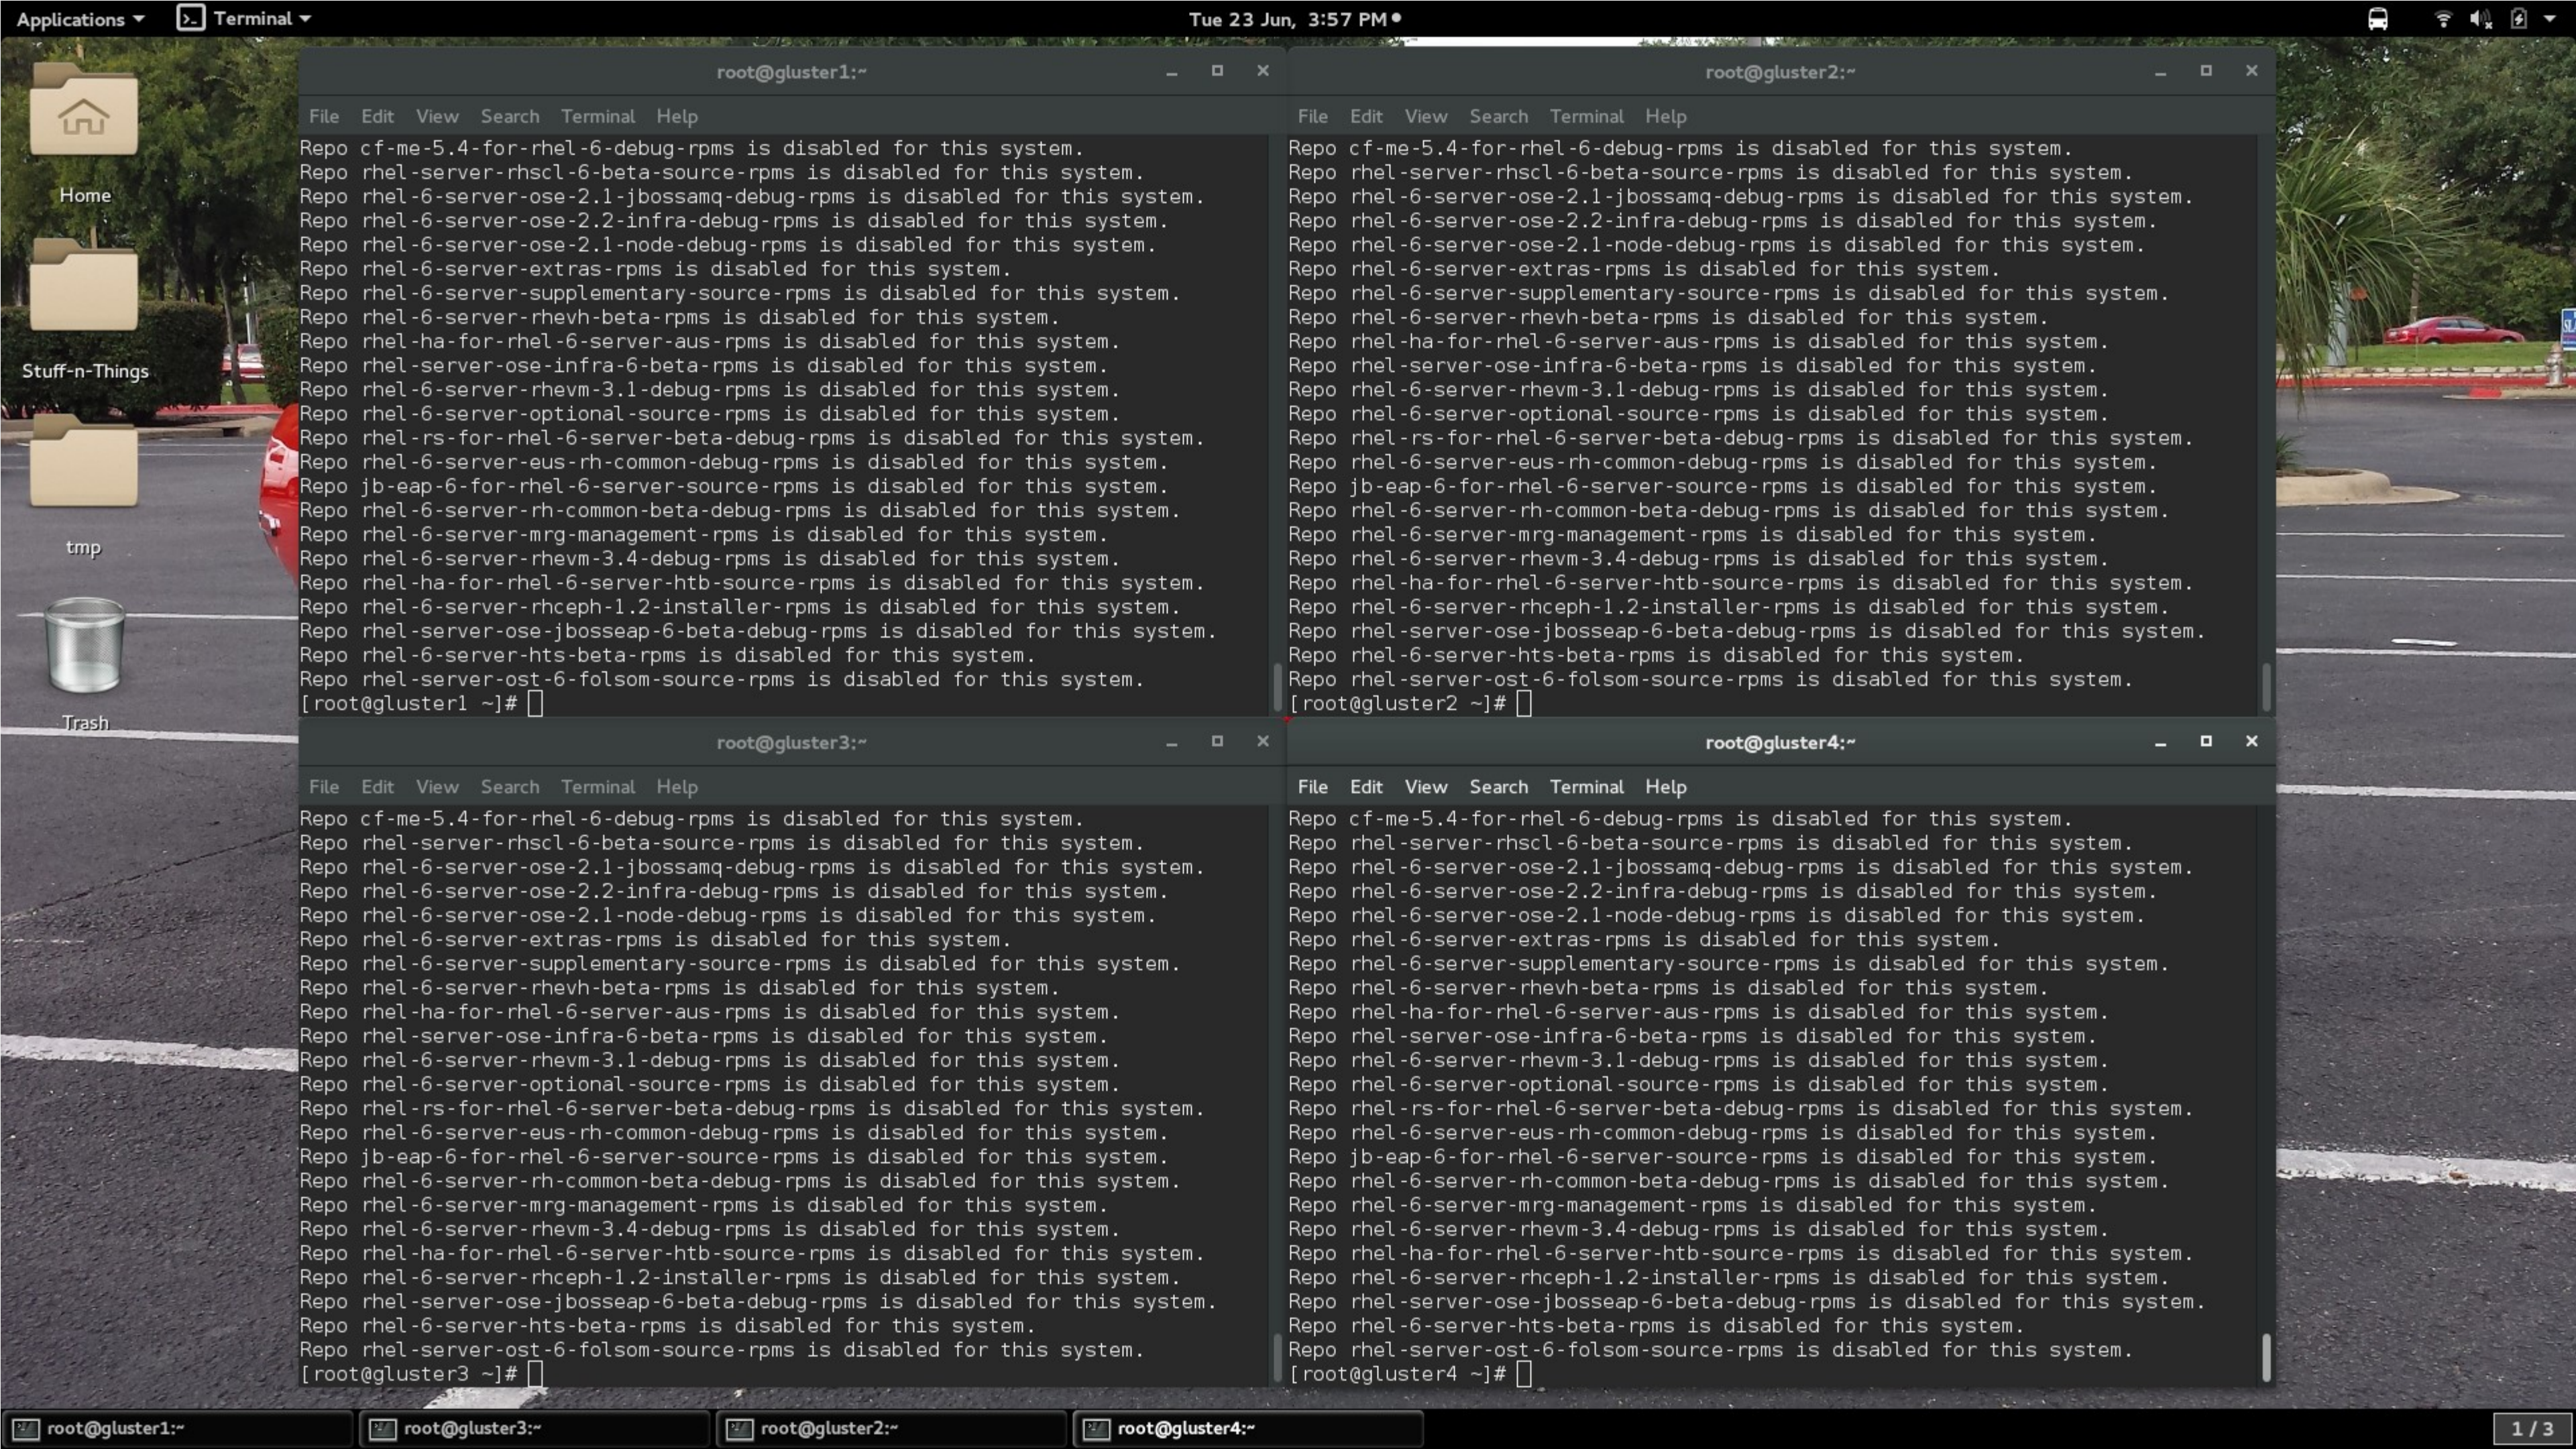
- Register the systems to Red Hat



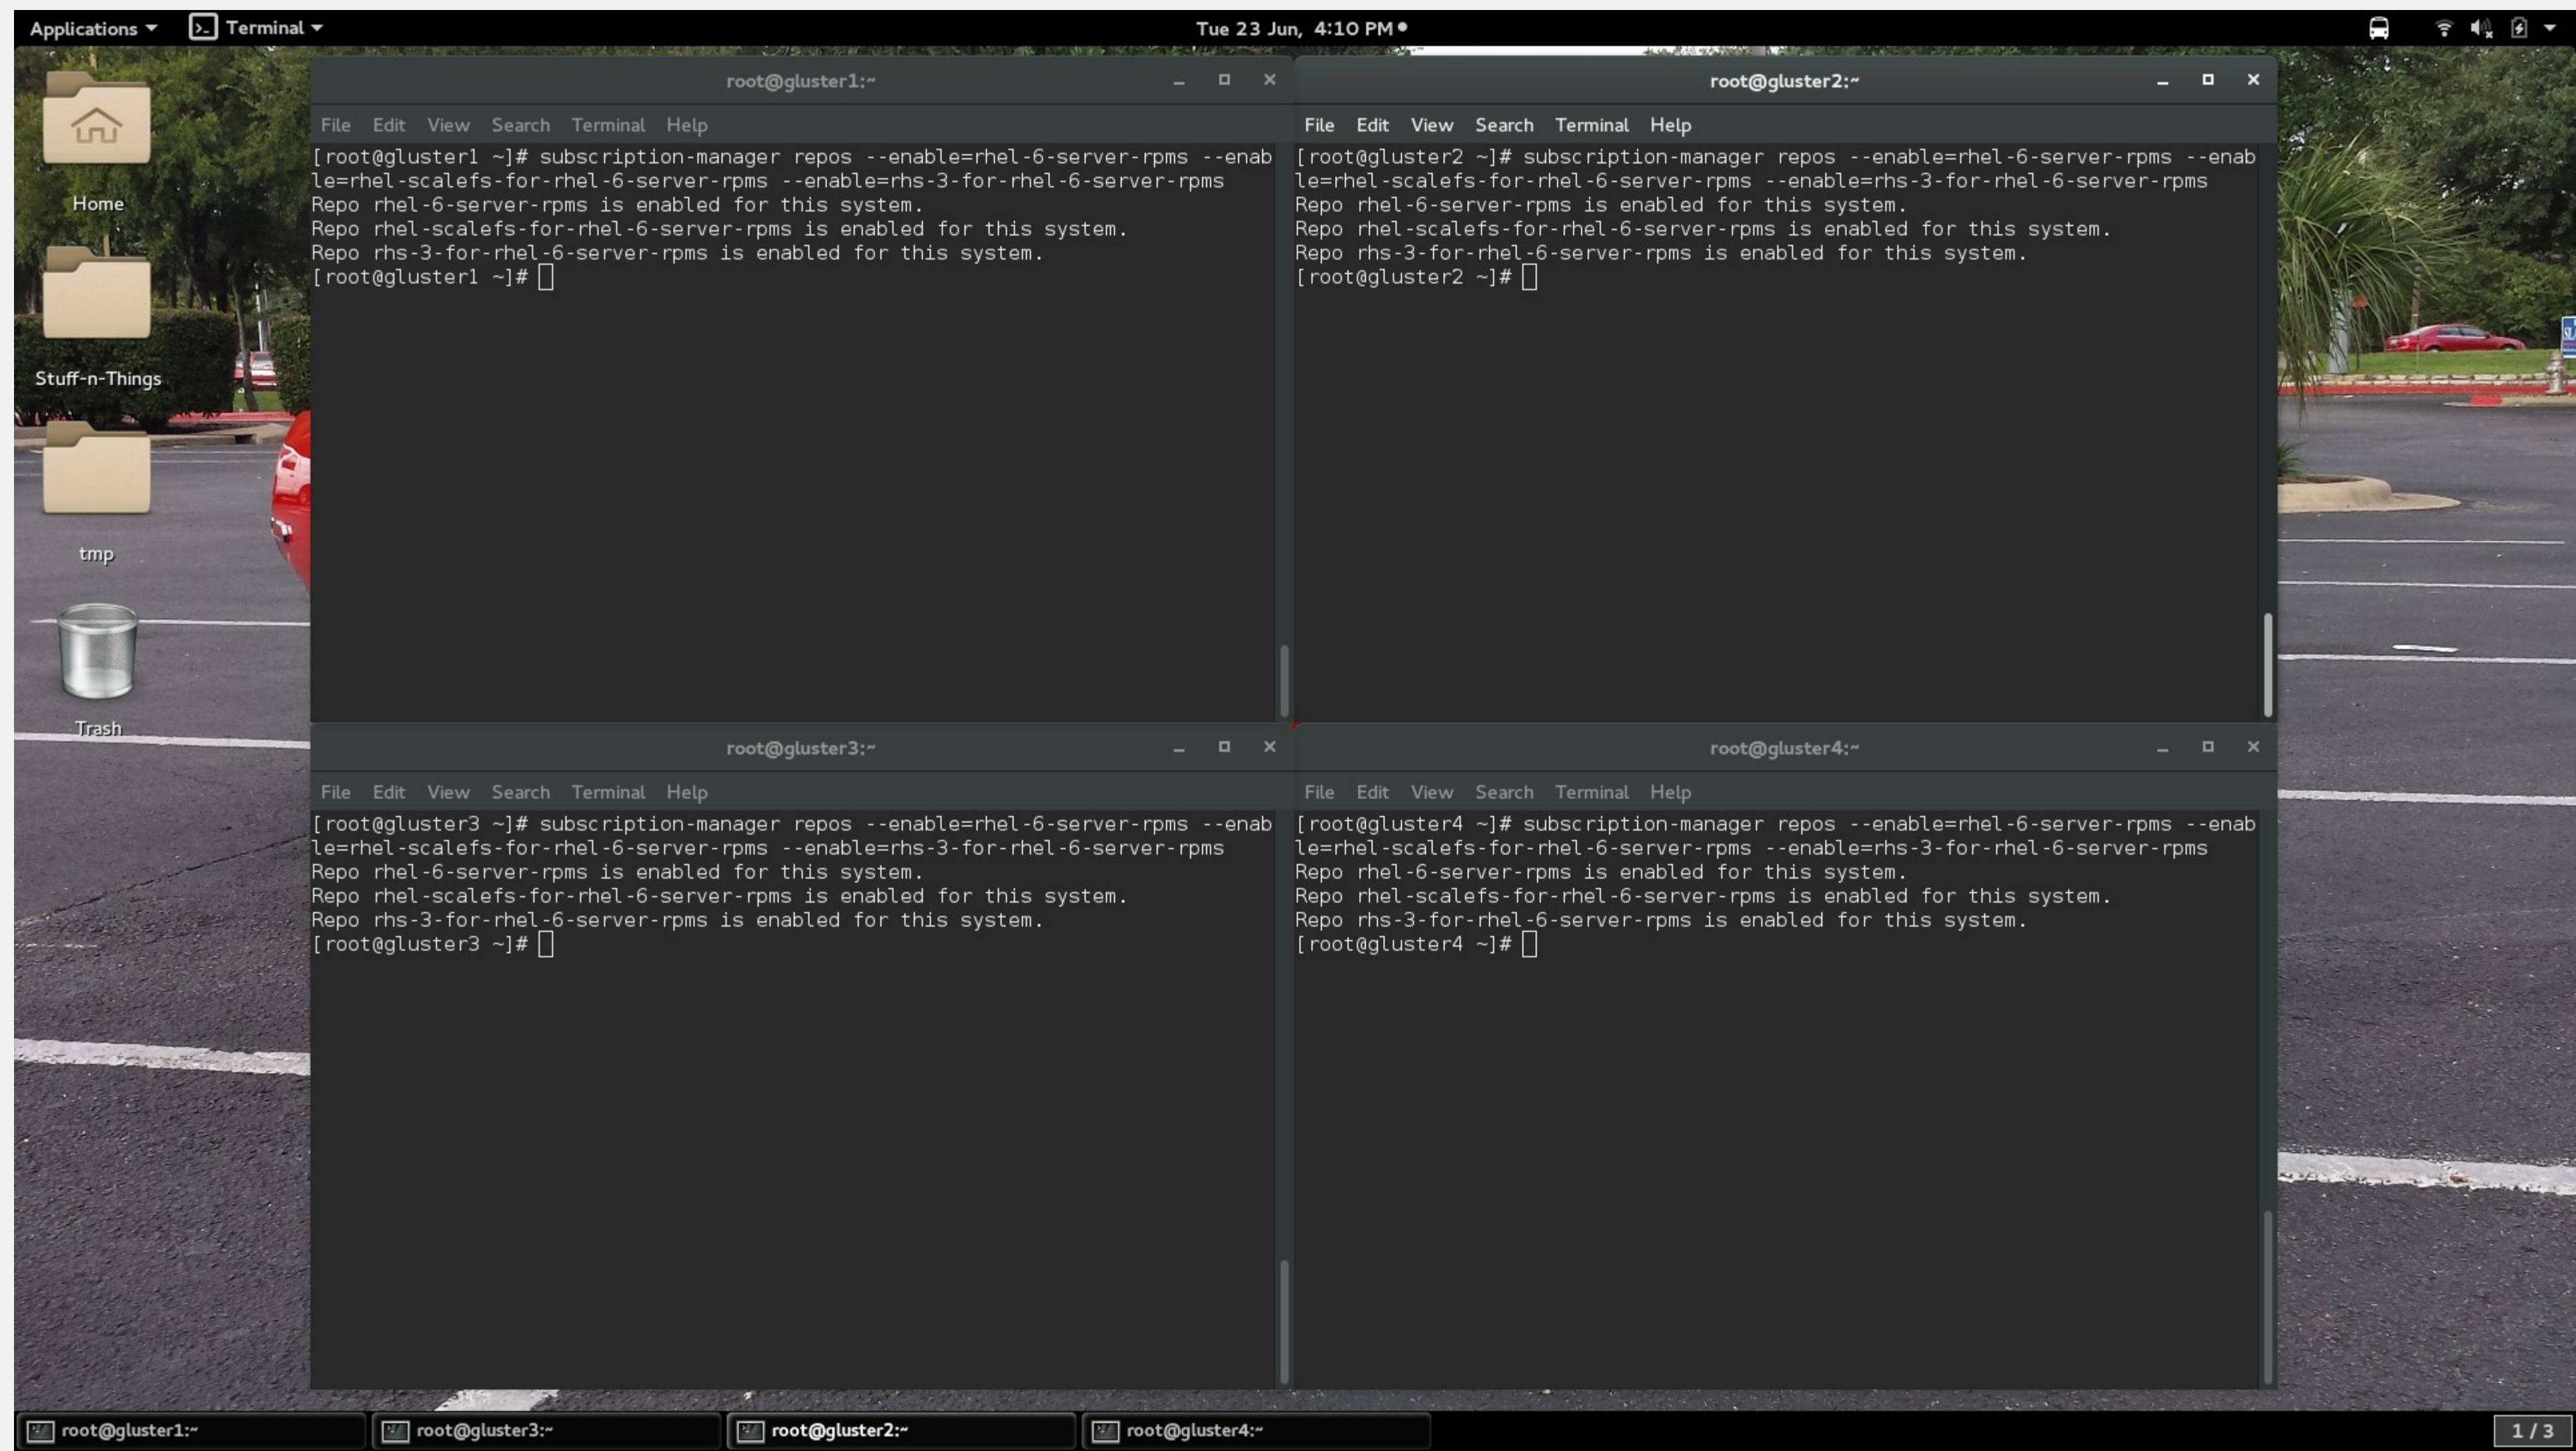









```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# subscription-manager repos --enable=rhel-6-server-rpms --enable=rhel-scalefs-for-rhel-6-server-rpms --enable=rhs-3-for-rhel-6-server-rpms  
Repo rhel-6-server-rpms is enabled for this system.  
Repo rhel-scalefs-for-rhel-6-server-rpms is enabled for this system.  
Repo rhs-3-for-rhel-6-server-rpms is enabled for this system.  
[root@gluster1 ~]#
```




```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# yum install redhat-storage-server  
Loaded plugins: product-id, security, subscription-manager  
This system is receiving updates from Red Hat Subscription Management.  
█
```

```
root@gluster1:~
File Edit View Search Terminal Help
audit-libs      x86_64 2.3.7-5.el6      rhel-6-server-rpms      71 k
cyrus-sasl      x86_64 2.1.23-15.el6_6.2 rhel-6-server-rpms      78 k
cyrus-sasl-lib  x86_64 2.1.23-15.el6_6.2 rhel-6-server-rpms     136 k
cyrus-sasl-plain x86_64 2.1.23-15.el6_6.2 rhel-6-server-rpms      31 k
file            x86_64 5.04-21.el6       rhel-6-server-rpms      47 k
file-libs       x86_64 5.04-21.el6       rhel-6-server-rpms     313 k
gnutls          x86_64 2.8.5-14.el6_5    rhel-6-server-rpms     346 k
keyutils-libs  x86_64 1.4-5.el6         rhel-6-server-rpms      20 k
kpartx          x86_64 0.4.9-80.el6_6.3  rhel-6-server-rpms      63 k
libX11          x86_64 1.6.0-2.2.el6     rhel-6-server-rpms     586 k
libX11-common  noarch 1.6.0-2.2.el6     rhel-6-server-rpms     192 k
libselinux      x86_64 2.0.94-5.8.el6    rhel-6-server-rpms     109 k
libselinux-utils x86_64 2.0.94-5.8.el6    rhel-6-server-rpms      82 k
policycoreutils x86_64 2.0.83-19.47.el6_6.1
sg3_utils-libs  x86_64 1.28-6.el6        rhel-6-server-rpms     680 k
                x86_64 1.28-6.el6        rhel-6-server-rpms      52 k

Transaction Summary
=====
Install      179 Package(s)
Upgrade      16 Package(s)

Total download size: 122 M
Is this ok [y/N]:
```



```
root@gluster1:~  
File Edit View Search Terminal Help  
Dependency Updated:  
  audit.x86_64 0:2.3.7-5.el6  
  audit-libs.x86_64 0:2.3.7-5.el6  
  cyrus-sasl.x86_64 0:2.1.23-15.el6_6.2  
  cyrus-sasl-lib.x86_64 0:2.1.23-15.el6_6.2  
  cyrus-sasl-plain.x86_64 0:2.1.23-15.el6_6.2  
  file.x86_64 0:5.04-21.el6  
  file-libs.x86_64 0:5.04-21.el6  
  gnutls.x86_64 0:2.8.5-14.el6_5  
  keyutils-libs.x86_64 0:1.4-5.el6  
  kpartx.x86_64 0:0.4.9-80.el6_6.3  
  libX11.x86_64 0:1.6.0-2.2.el6  
  libX11-common.noarch 0:1.6.0-2.2.el6  
  libselinux.x86_64 0:2.0.94-5.8.el6  
  libselinux-utils.x86_64 0:2.0.94-5.8.el6  
  policycoreutils.x86_64 0:2.0.83-19.47.el6_6.1  
  sg3_utils-libs.x86_64 0:1.28-6.el6  
  
Replaced:  
  redhat-logos.noarch 0:60.0.14-1.el6  
  
Complete!  
[root@gluster1 ~]#
```


Applications

Terminal

Home

Stuff-n-Things

tmp

Trash

root@gluster1:~

File Edit View Search Terminal Help

Dependency Updated:
audit.x86_64 0:2.3.7-5.el6
audit-libs.x86_64 0:2.3.7-5.el6
cyrus-sasl.x86_64 0:2.1.23-15.el6_6.2
cyrus-sasl-lib.x86_64 0:2.1.23-15.el6_6.2
cyrus-sasl-plain.x86_64 0:2.1.23-15.el6_6.2
file.x86_64 0:5.04-21.el6
file-libs.x86_64 0:5.04-21.el6
gnutls.x86_64 0:2.8.5-14.el6_5
keyutils-libs.x86_64 0:1.4-5.el6
kpartx.x86_64 0:0.4.9-80.el6_6.3
libX11.x86_64 0:1.6.0-2.2.el6
libX11-common.noarch 0:1.6.0-2.2.el6
libselinux.x86_64 0:2.0.94-5.8.el6
libselinux-utils.x86_64 0:2.0.94-5.8.el6
policycoreutils.x86_64 0:2.0.83-19.47.el6_6.1
sg3_utils-libs.x86_64 0:1.28-6.el6

Replaced:
redhat-logos.noarch 0:60.0.14-1.el6

Complete!
[root@gluster1 ~]# yum install redhat-storage-server

root@gluster3:~

File Edit View Search Terminal Help

Downloading Packages:
warning: rpmts_HdrFromFdno: Header V3 RSA/SHA256 Signature, key ID fd431d51: NOKEY
Retrieving key from file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
Importing GPG key 0xFD431D51:
Userid : Red Hat, Inc. (release key 2) <security@redhat.com>
Package: redhat-release-server-6Server-6.5.0.1.el6.x86_64 (@anaconda-RedHatEnterpriseLinux-201311111358.x86_64/6.5)
From : /etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
Importing GPG key 0x2FA658E0:
Userid : Red Hat, Inc. (auxiliary key) <security@redhat.com>
Package: redhat-release-server-6Server-6.5.0.1.el6.x86_64 (@anaconda-RedHatEnterpriseLinux-201311111358.x86_64/6.5)
From : /etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
Updating : cyrus-sasl-lib-2.1.23-15.el6_6.2.x86_64 1/212
Updating : libselinux-2.0.94-5.8.el6.x86_64 2/212
Updating : audit-libs-2.3.7-5.el6.x86_64 3/212
Installing : glusterfs-libs-3.6.0.53-1.el6rhs.x86_64 4/212
Installing : glusterfs-3.6.0.53-1.el6rhs.x86_64 5/212

root@gluster2:~

File Edit View Search Terminal Help

sg3_utils-libs x86_64 1.28-6.el6 rhel-6-server-rpms 52 k

Transaction Summary
=====

Install 179 Package(s)
Upgrade 16 Package(s)

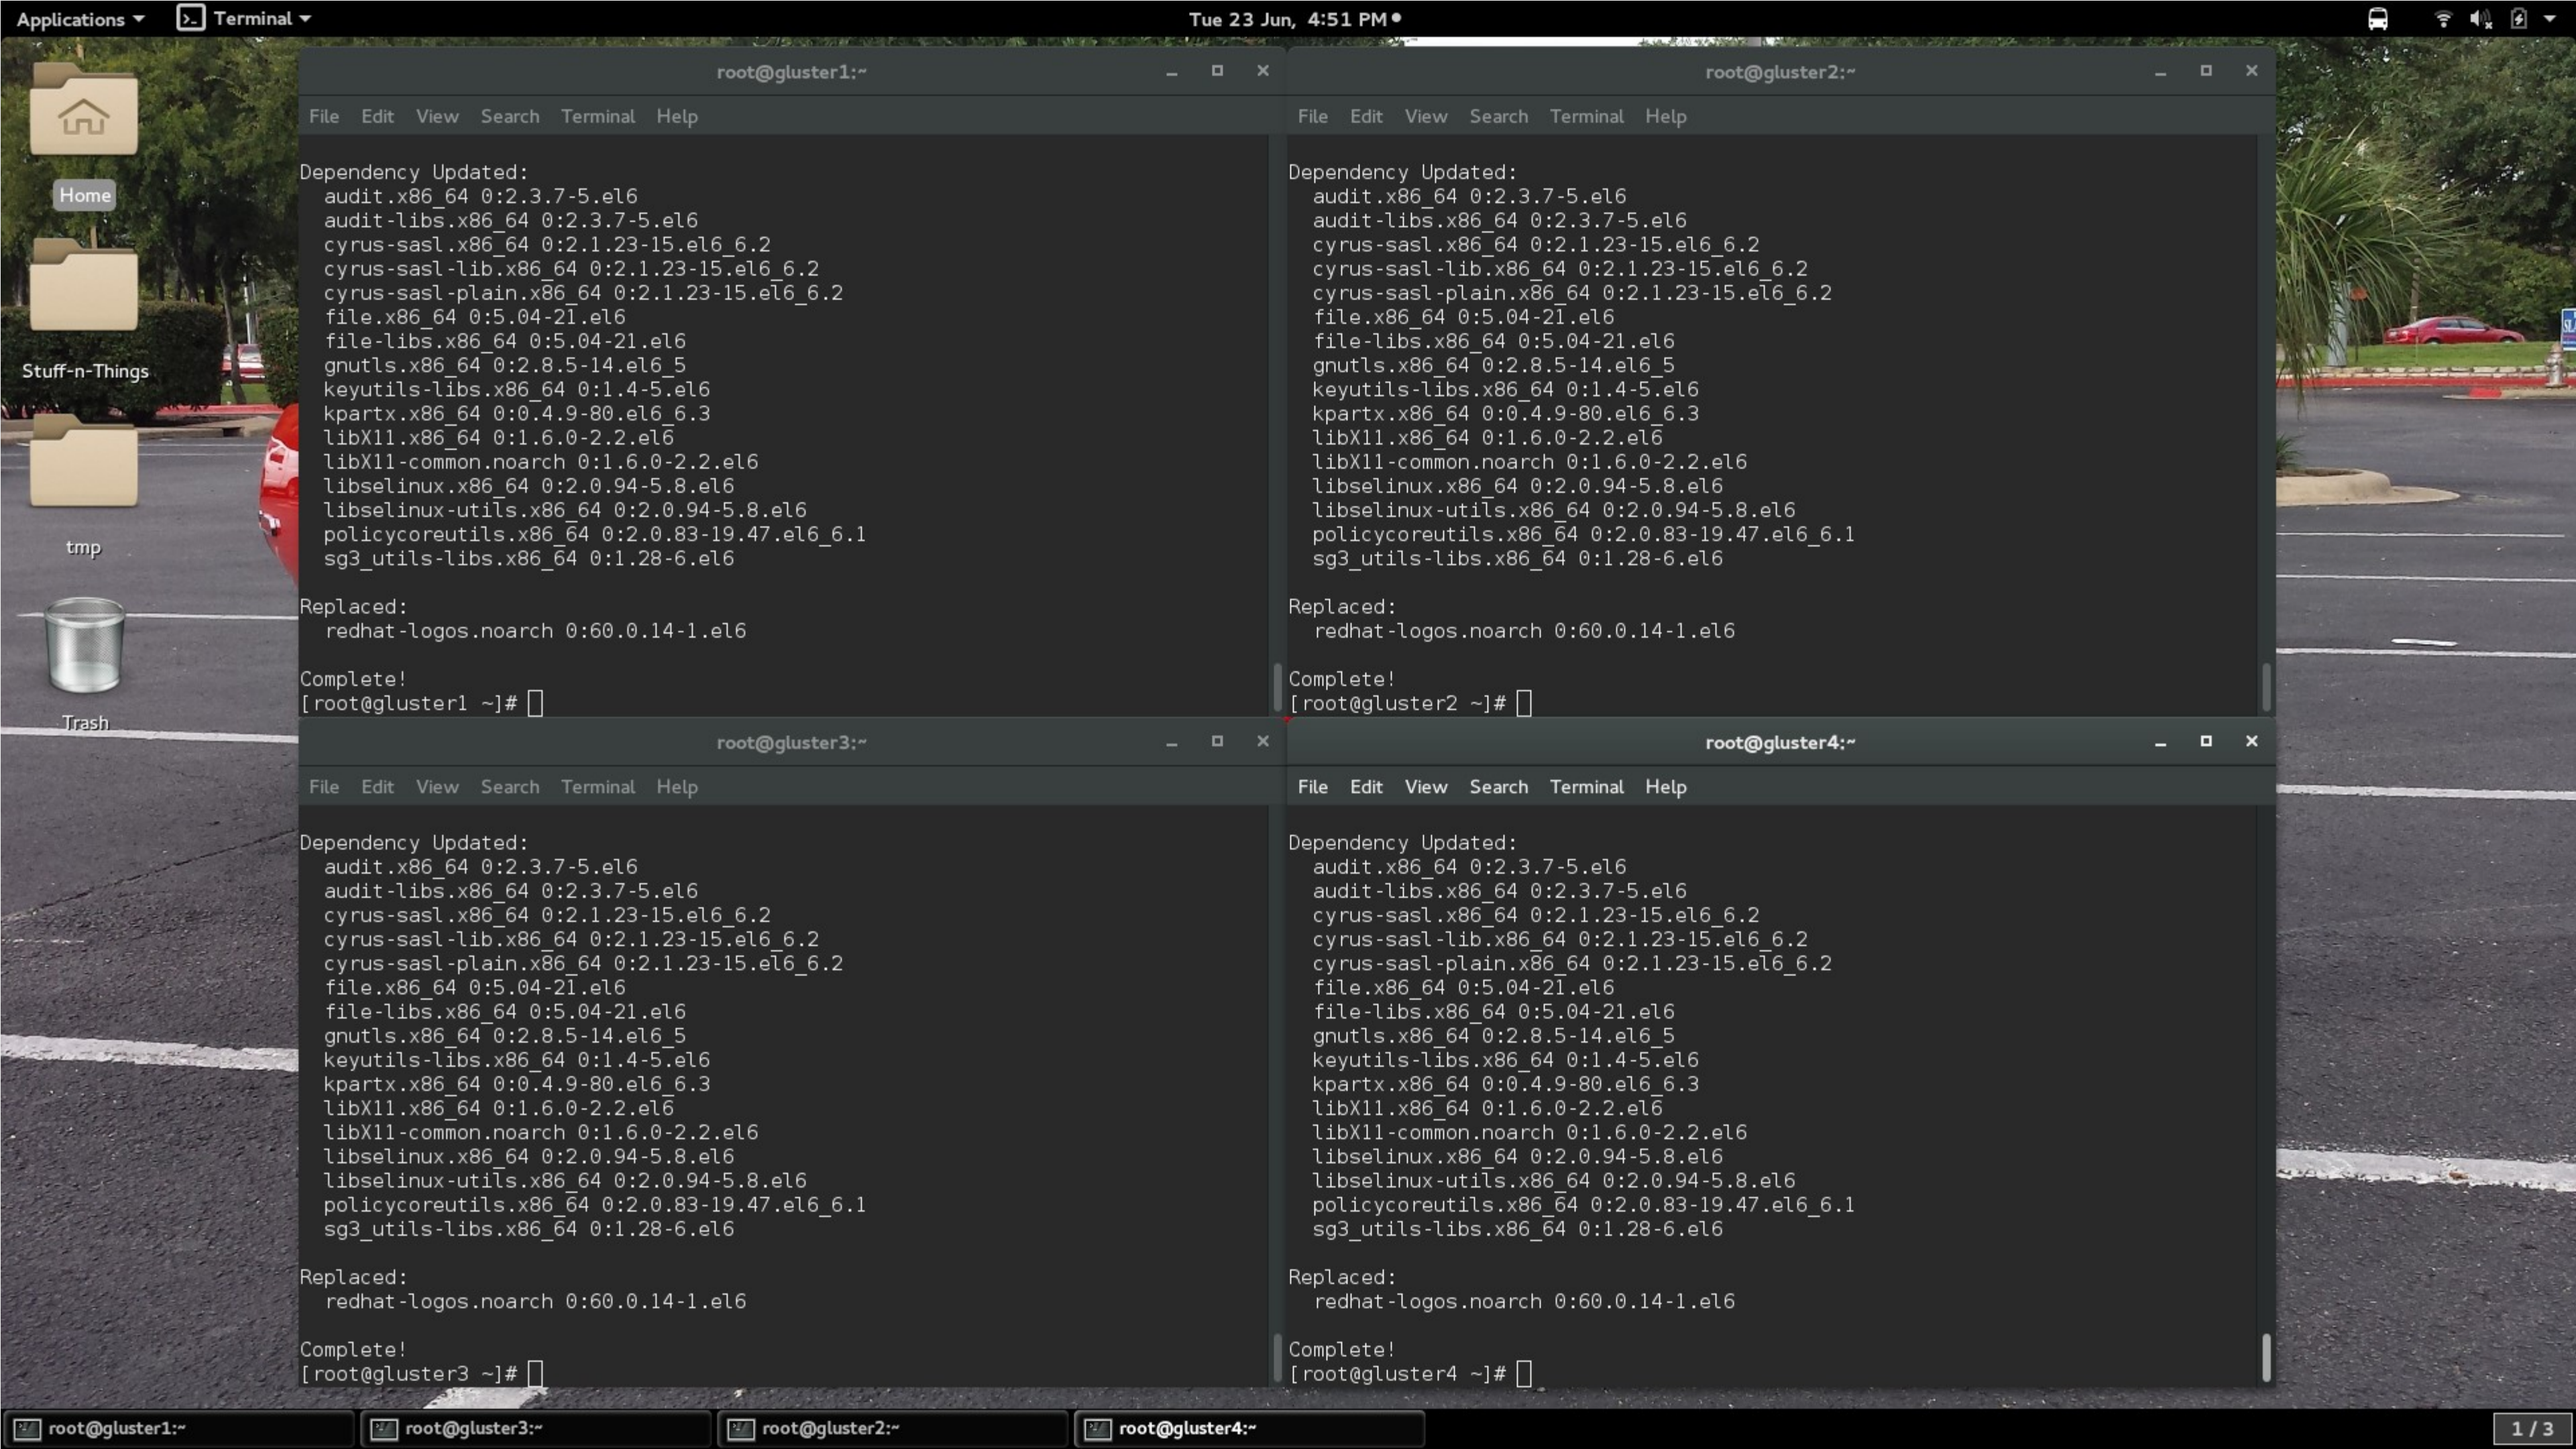
Total size: 122 M
Downloading Packages:
warning: rpmts_HdrFromFdno: Header V3 RSA/SHA256 Signature, key ID fd431d51: NOKEY
Retrieving key from file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
Importing GPG key 0xFD431D51:
Userid : Red Hat, Inc. (release key 2) <security@redhat.com>
Package: redhat-release-server-6Server-6.5.0.1.el6.x86_64 (@anaconda-RedHatEnterpriseLinux-201311111358.x86_64/6.5)
From : /etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
Importing GPG key 0x2FA658E0:
Userid : Red Hat, Inc. (auxiliary key) <security@redhat.com>
Package: redhat-release-server-6Server-6.5.0.1.el6.x86_64 (@anaconda-RedHatEnterpriseLinux-201311111358.x86_64/6.5)
From : /etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
Running rpm_check_debug

root@gluster4:~

File Edit View Search Terminal Help

Downloading Packages:
warning: rpmts_HdrFromFdno: Header V3 RSA/SHA256 Signature, key ID fd431d51: NOKEY
Retrieving key from file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
Importing GPG key 0xFD431D51:
Userid : Red Hat, Inc. (release key 2) <security@redhat.com>
Package: redhat-release-server-6Server-6.5.0.1.el6.x86_64 (@anaconda-RedHatEnterpriseLinux-201311111358.x86_64/6.5)
From : /etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
Importing GPG key 0x2FA658E0:
Userid : Red Hat, Inc. (auxiliary key) <security@redhat.com>
Package: redhat-release-server-6Server-6.5.0.1.el6.x86_64 (@anaconda-RedHatEnterpriseLinux-201311111358.x86_64/6.5)
From : /etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
Updating : cyrus-sasl-lib-2.1.23-15.el6_6.2.x86_64 1/212
Updating : libselinux-2.0.94-5.8.el6.x86_64 2/212
Updating : audit-libs-2.3.7-5.el6.x86_64 3/212
Installing : glusterfs-libs-3.6.0.53-1.el6rhs.x86_64 4/212
Installing : glusterfs-3.6.0.53-1.el6rhs.x86_64 5/212

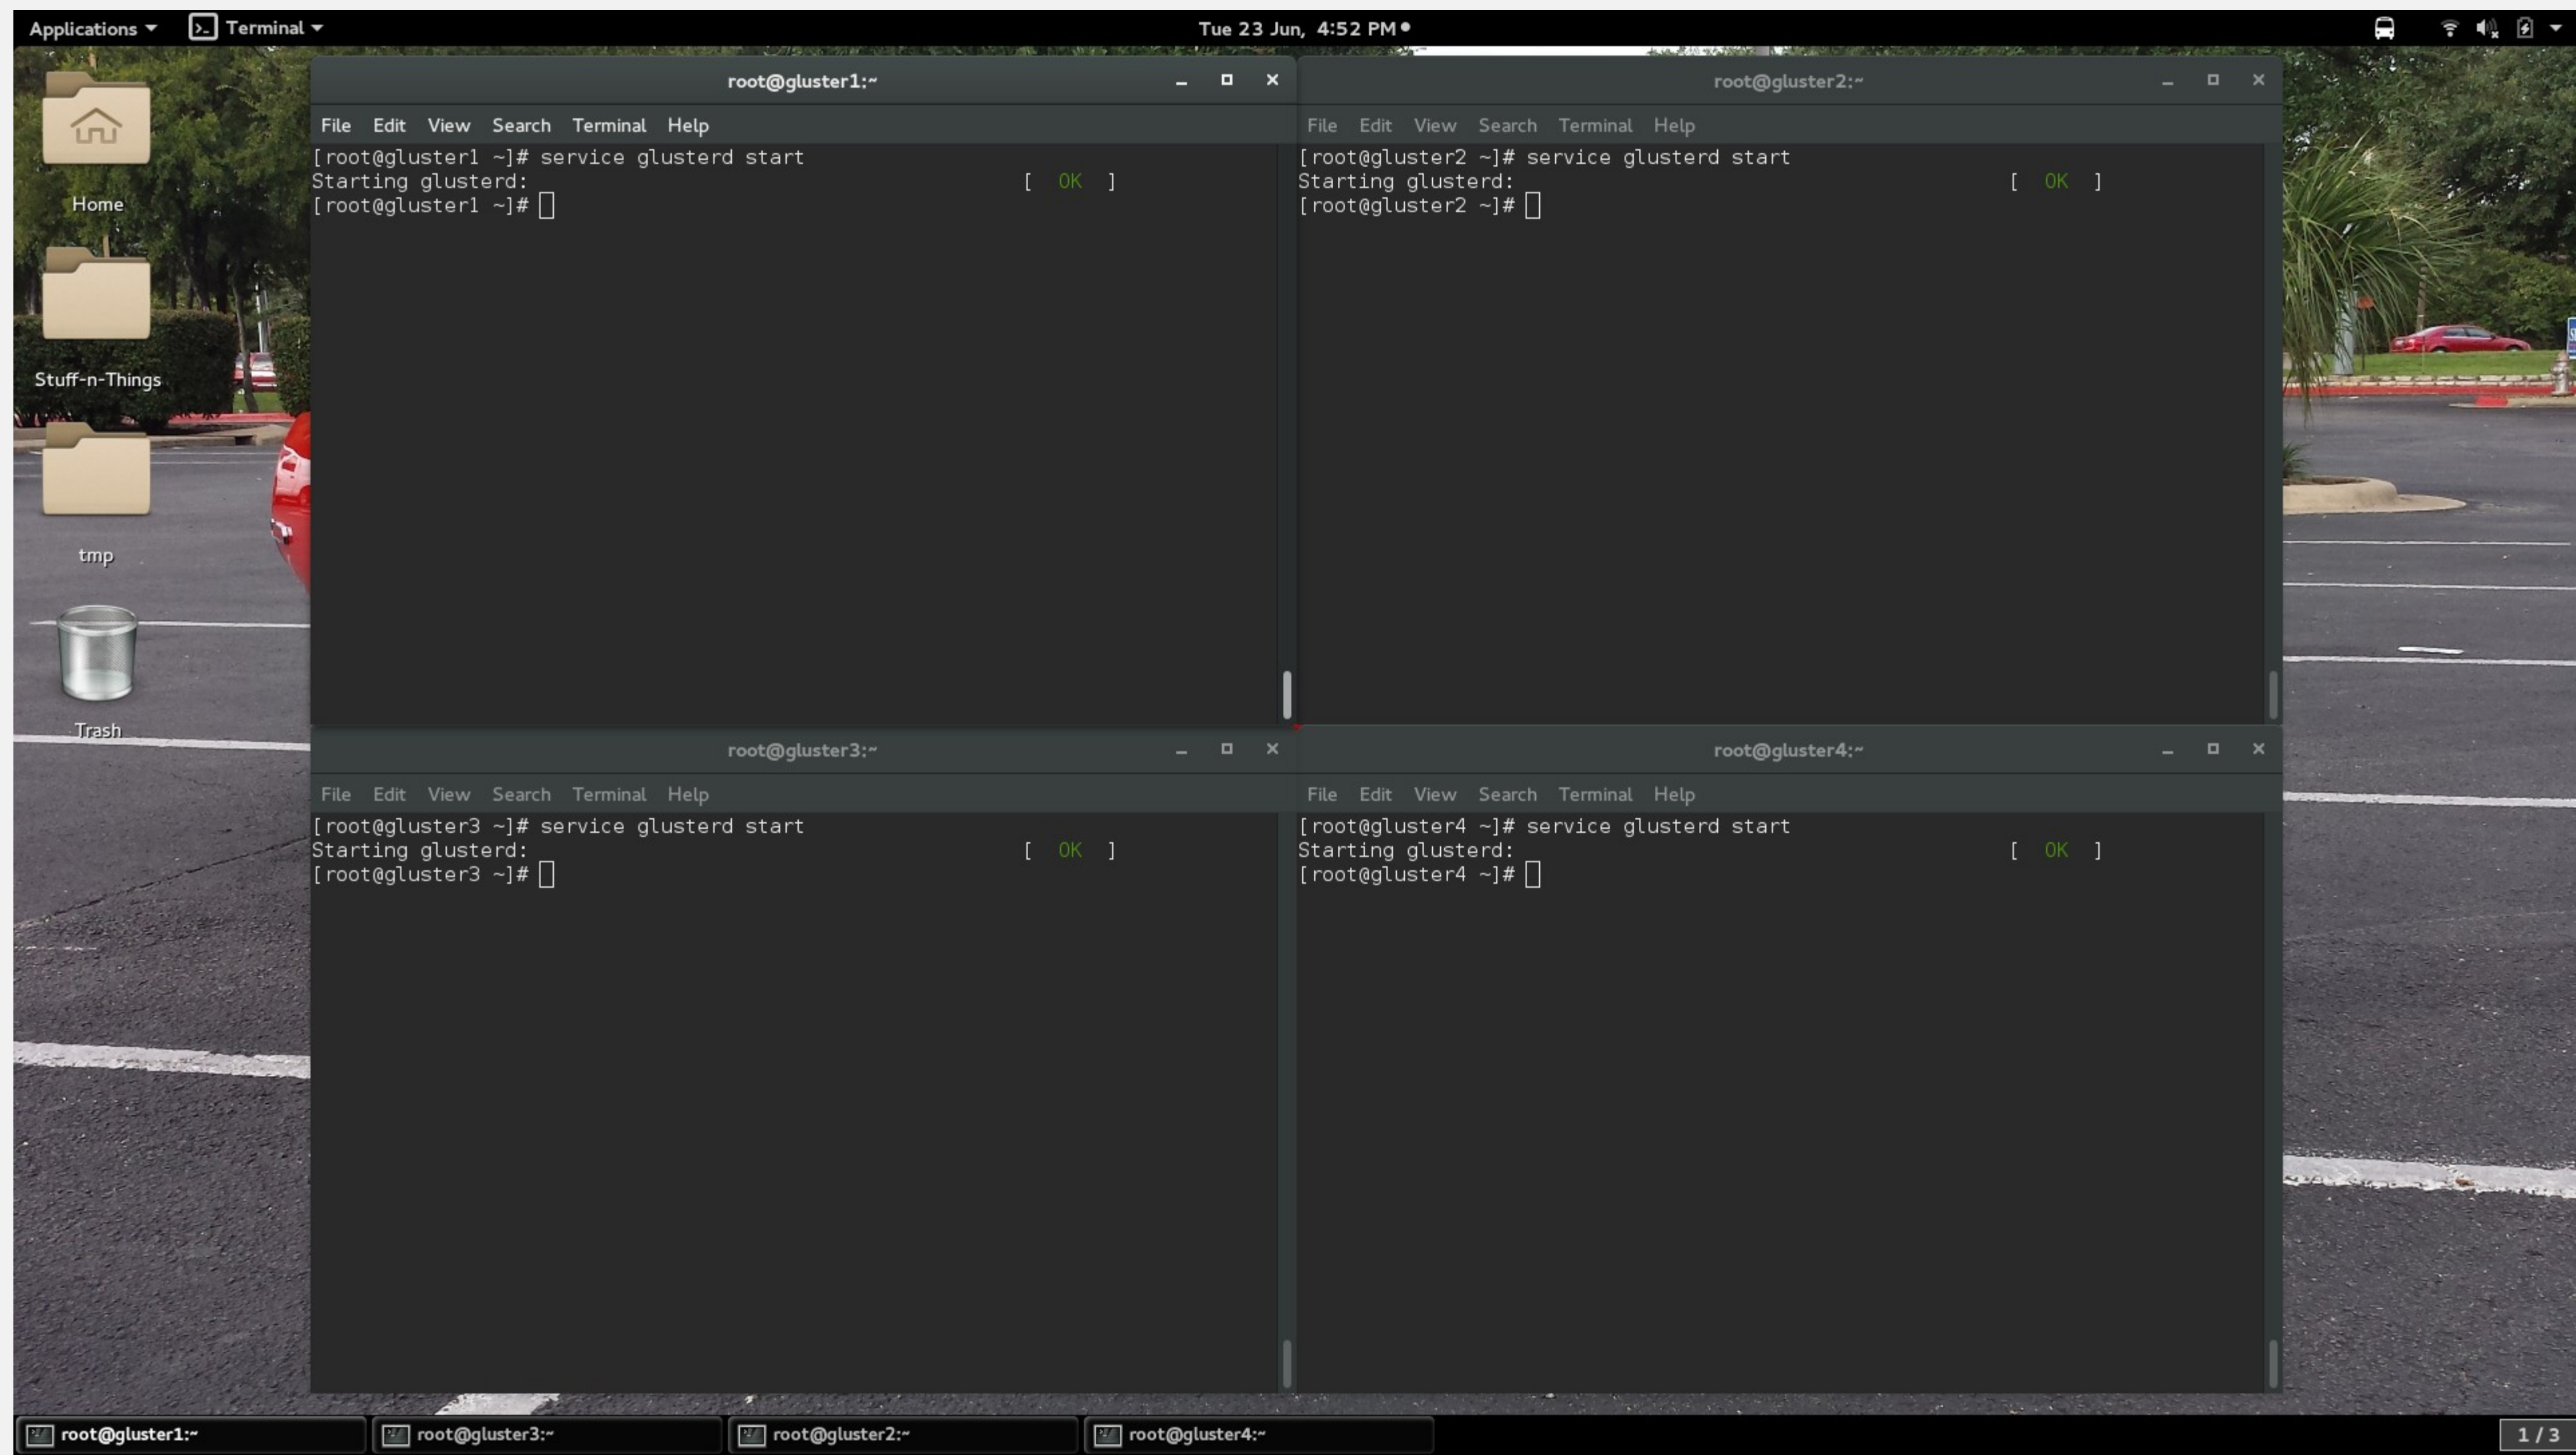
root@gluster1:~ root@gluster3:~ root@gluster2:~ root@gluster4:~ 1 / 3




```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# chkconfig glusterd --list  
glusterd      0:off    1:off    2:on     3:on     4:on     5:on     6:off  
[root@gluster1 ~]#
```



```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# service glusterd start  
Starting glusterd: [ OK ]  
[root@gluster1 ~]#
```



CLIENT INSTALLATION


```
root@host189:~  
File Edit View Search Terminal Help  
[root@host189 ~]# subscription-manager register --auto-attach  
Username: tcameron@redhat.com  
Password:  
The system has been registered with ID: 5c203798-4c40-4cc4-a48b-9f7451a8131e  
Installed Product Current Status:  
Product Name: Red Hat Enterprise Linux Server  
Status: Subscribed  
[root@host189 ~]#
```



```
root@host189:~  
File Edit View Search Terminal Help  
[root@host189 ~]# subscription-manager repos --disable="*"  
[
```

```
root@host189:~  
File Edit View Search Terminal Help  
Repo cf-me-5.4-for-rhel-6-debug-rpms is disabled for this system.  
Repo rhel-server-rhscl-6-beta-source-rpms is disabled for this system.  
Repo rhel-6-server-ose-2.1-jbossamq-debug-rpms is disabled for this system.  
Repo rhel-6-server-ose-2.2-infra-debug-rpms is disabled for this system.  
Repo rhel-6-server-ose-2.1-node-debug-rpms is disabled for this system.  
Repo rhel-6-server-extras-rpms is disabled for this system.  
Repo rhel-6-server-supplementary-source-rpms is disabled for this system.  
Repo rhel-6-server-rhev-beta-rpms is disabled for this system.  
Repo rhel-ha-for-rhel-6-server-aus-rpms is disabled for this system.  
Repo rhel-server-ose-infra-6-beta-rpms is disabled for this system.  
Repo rhel-6-server-rhev-3.1-debug-rpms is disabled for this system.  
Repo rhel-6-server-optional-source-rpms is disabled for this system.  
Repo rhel-rs-for-rhel-6-server-beta-debug-rpms is disabled for this system.  
Repo rhel-6-server-eus-rh-common-debug-rpms is disabled for this system.  
Repo jb-eap-6-for-rhel-6-server-source-rpms is disabled for this system.  
Repo rhel-6-server-rh-common-beta-debug-rpms is disabled for this system.  
Repo rhel-6-server-mrg-management-rpms is disabled for this system.  
Repo rhel-6-server-rhev-3.4-debug-rpms is disabled for this system.  
Repo rhel-ha-for-rhel-6-server-htb-source-rpms is disabled for this system.  
Repo rhel-6-server-rhceph-1.2-installer-rpms is disabled for this system.  
Repo rhel-server-ose-jbosseap-6-beta-debug-rpms is disabled for this system.  
Repo rhel-6-server-hts-beta-rpms is disabled for this system.  
Repo rhel-server-ost-6-folsom-source-rpms is disabled for this system.  
[root@host189 ~]#
```



```
root@host189:~  
File Edit View Search Terminal Help  
[root@host189 ~]# subscription-manager repos --enable=rhel-6-server-rpms --enable=rhel-6-server-rhs-client-1-rpms  
Repo rhel-6-server-rpms is enabled for this system.  
Repo rhel-6-server-rhs-client-1-rpms is enabled for this system.  
[root@host189 ~]#
```

```
root@host189:~  
File Edit View Search Terminal Help  
[root@host189 ~]# yum install glusterfs glusterfs-fuse
```

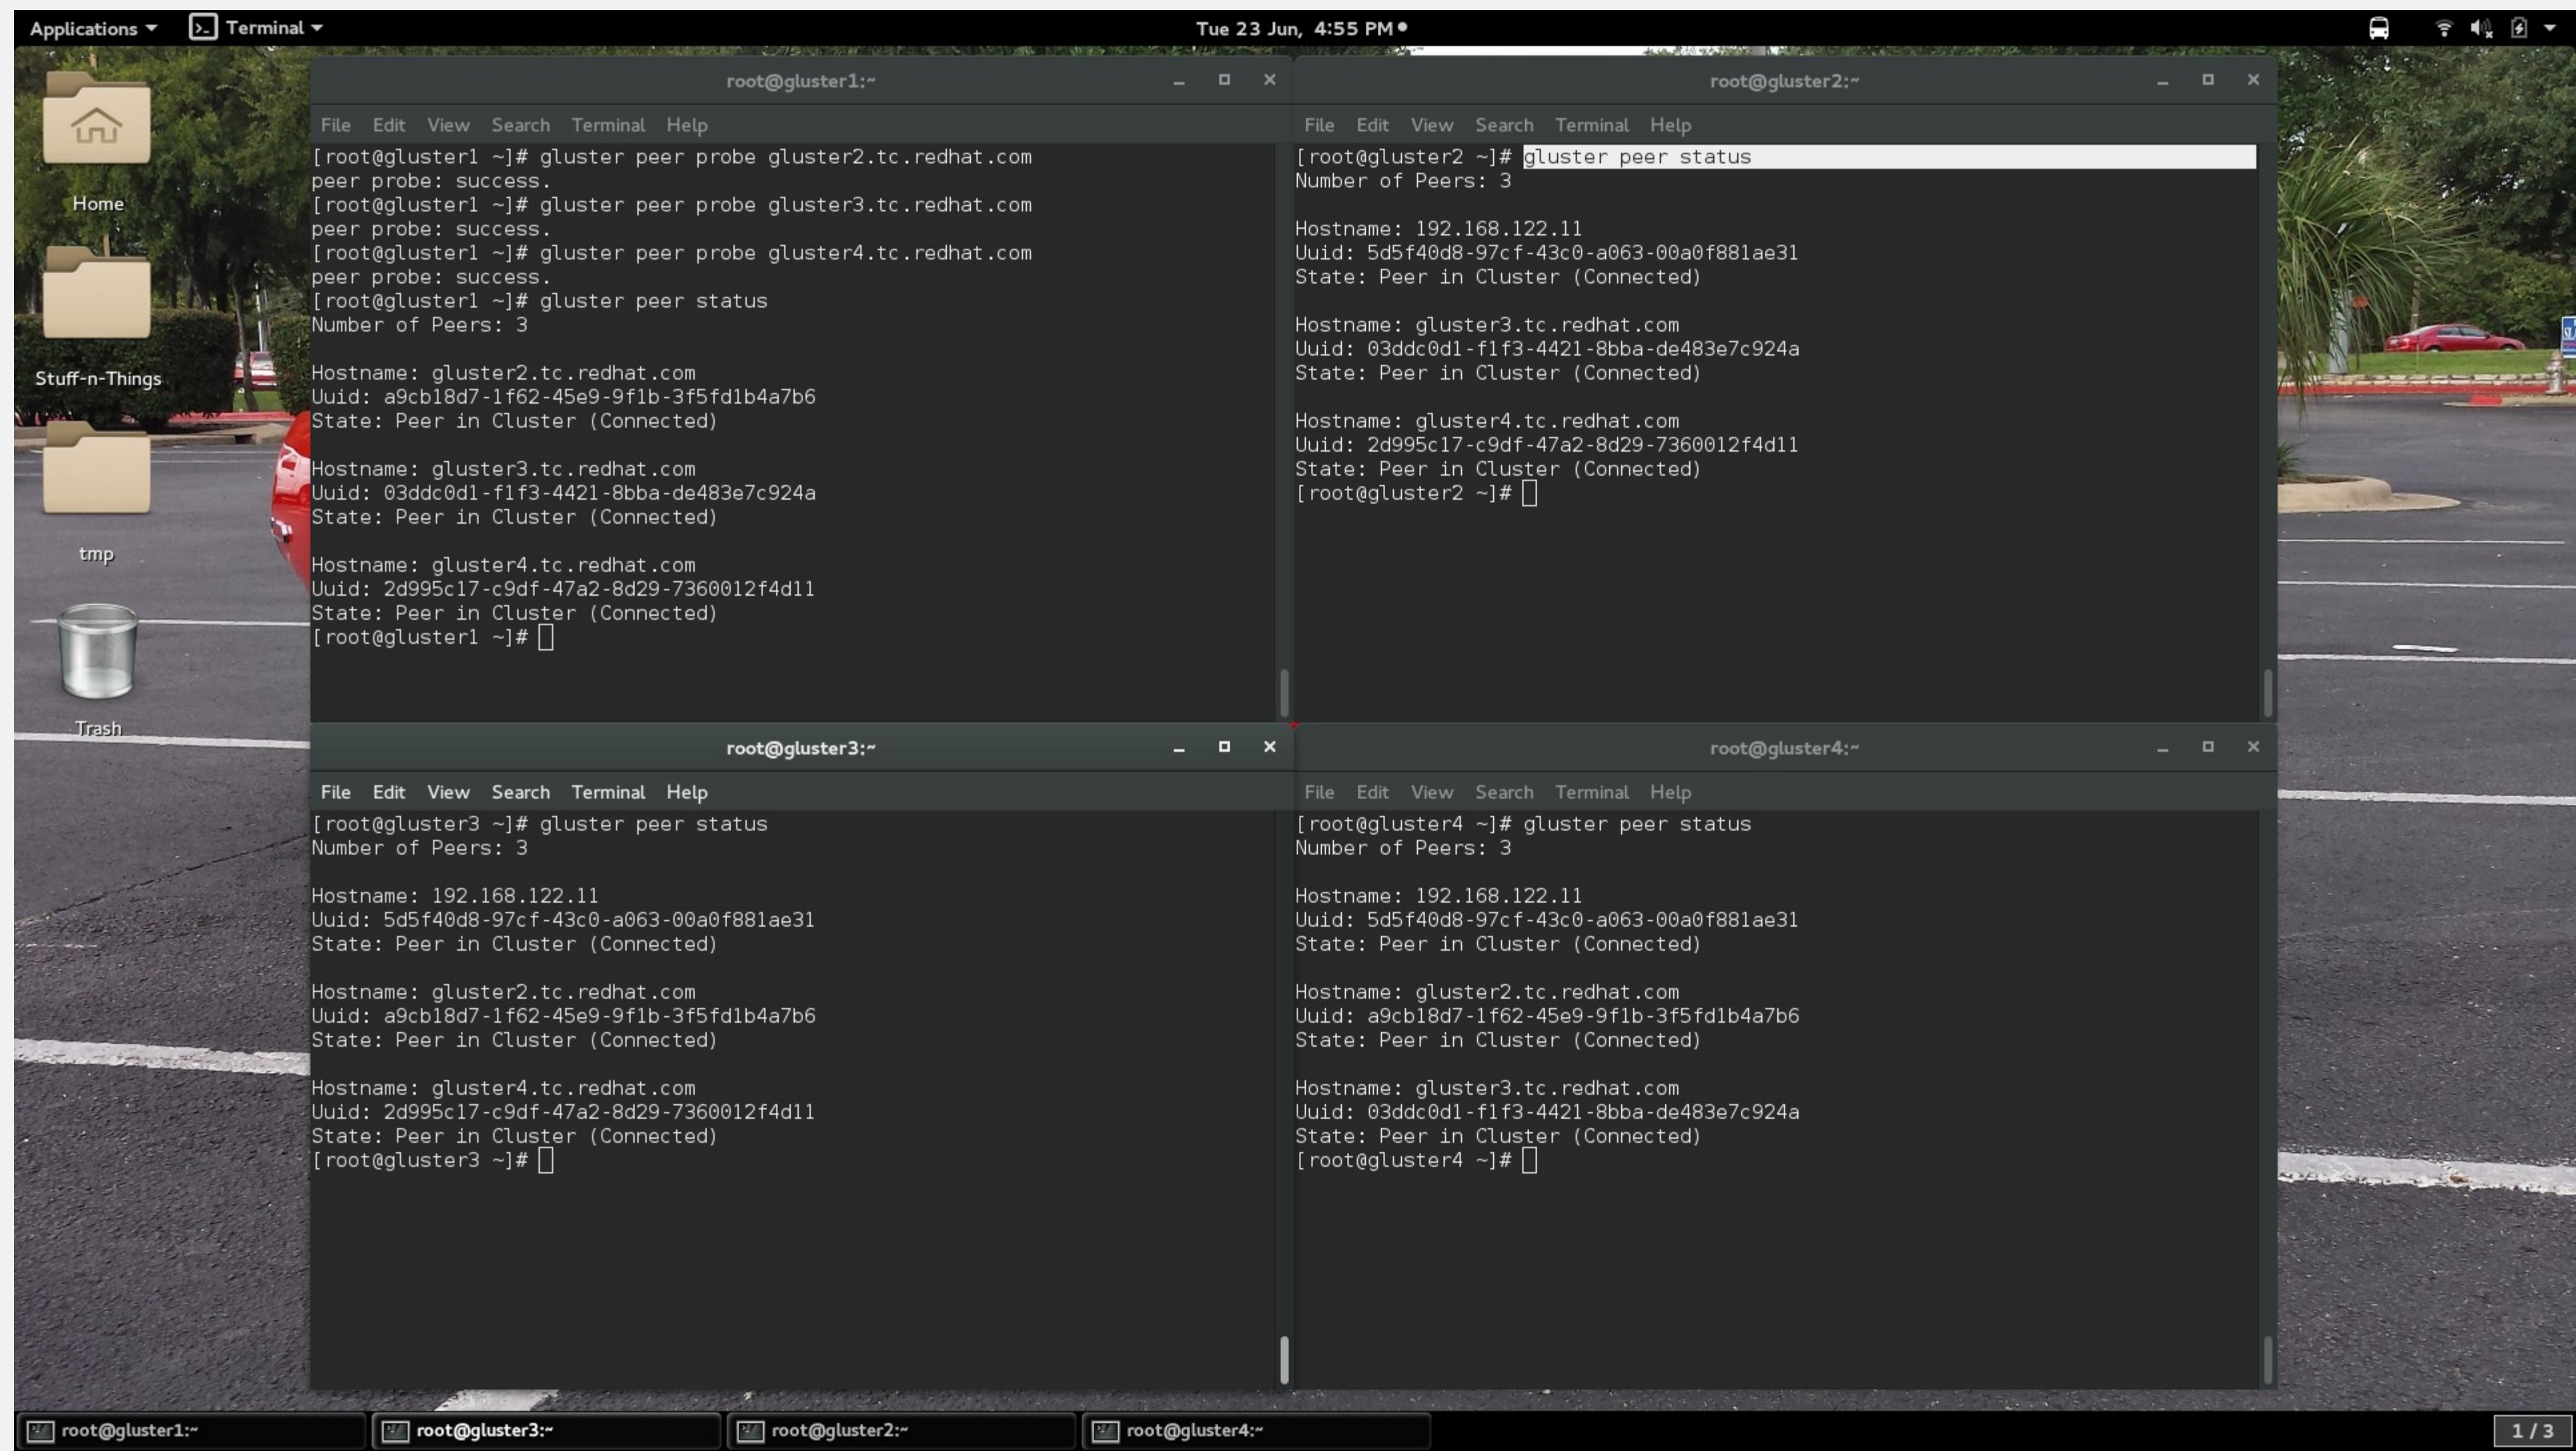


```
root@host189:~  
File Edit View Search Terminal Help  
rpriseLinux-201311111358.x86_64/6.5)  
From : /etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release  
Running rpm_check_debug  
Running Transaction Test  
Transaction Test Succeeded  
Running Transaction  
  Installing : glusterfs-libs-3.6.0.53-1.el6.x86_64 1/4  
  Installing : glusterfs-3.6.0.53-1.el6.x86_64 2/4  
  Installing : glusterfs-api-3.6.0.53-1.el6.x86_64 3/4  
  Installing : glusterfs-fuse-3.6.0.53-1.el6.x86_64 4/4  
rhel-6-server-rpms/productid | 1.7 kB 00:00  
  Verifying : glusterfs-api-3.6.0.53-1.el6.x86_64 1/4  
  Verifying : glusterfs-libs-3.6.0.53-1.el6.x86_64 2/4  
  Verifying : glusterfs-fuse-3.6.0.53-1.el6.x86_64 3/4  
  Verifying : glusterfs-3.6.0.53-1.el6.x86_64 4/4  
  
Installed:  
  glusterfs.x86_64 0:3.6.0.53-1.el6 glusterfs-fuse.x86_64 0:3.6.0.53-1.el6  
  
Dependency Installed:  
  glusterfs-api.x86_64 0:3.6.0.53-1.el6 glusterfs-libs.x86_64 0:3.6.0.53-1.el6  
  
Complete!  
[root@host189 ~]#
```


ADDING SERVERS TO THE TRUSTED POOL

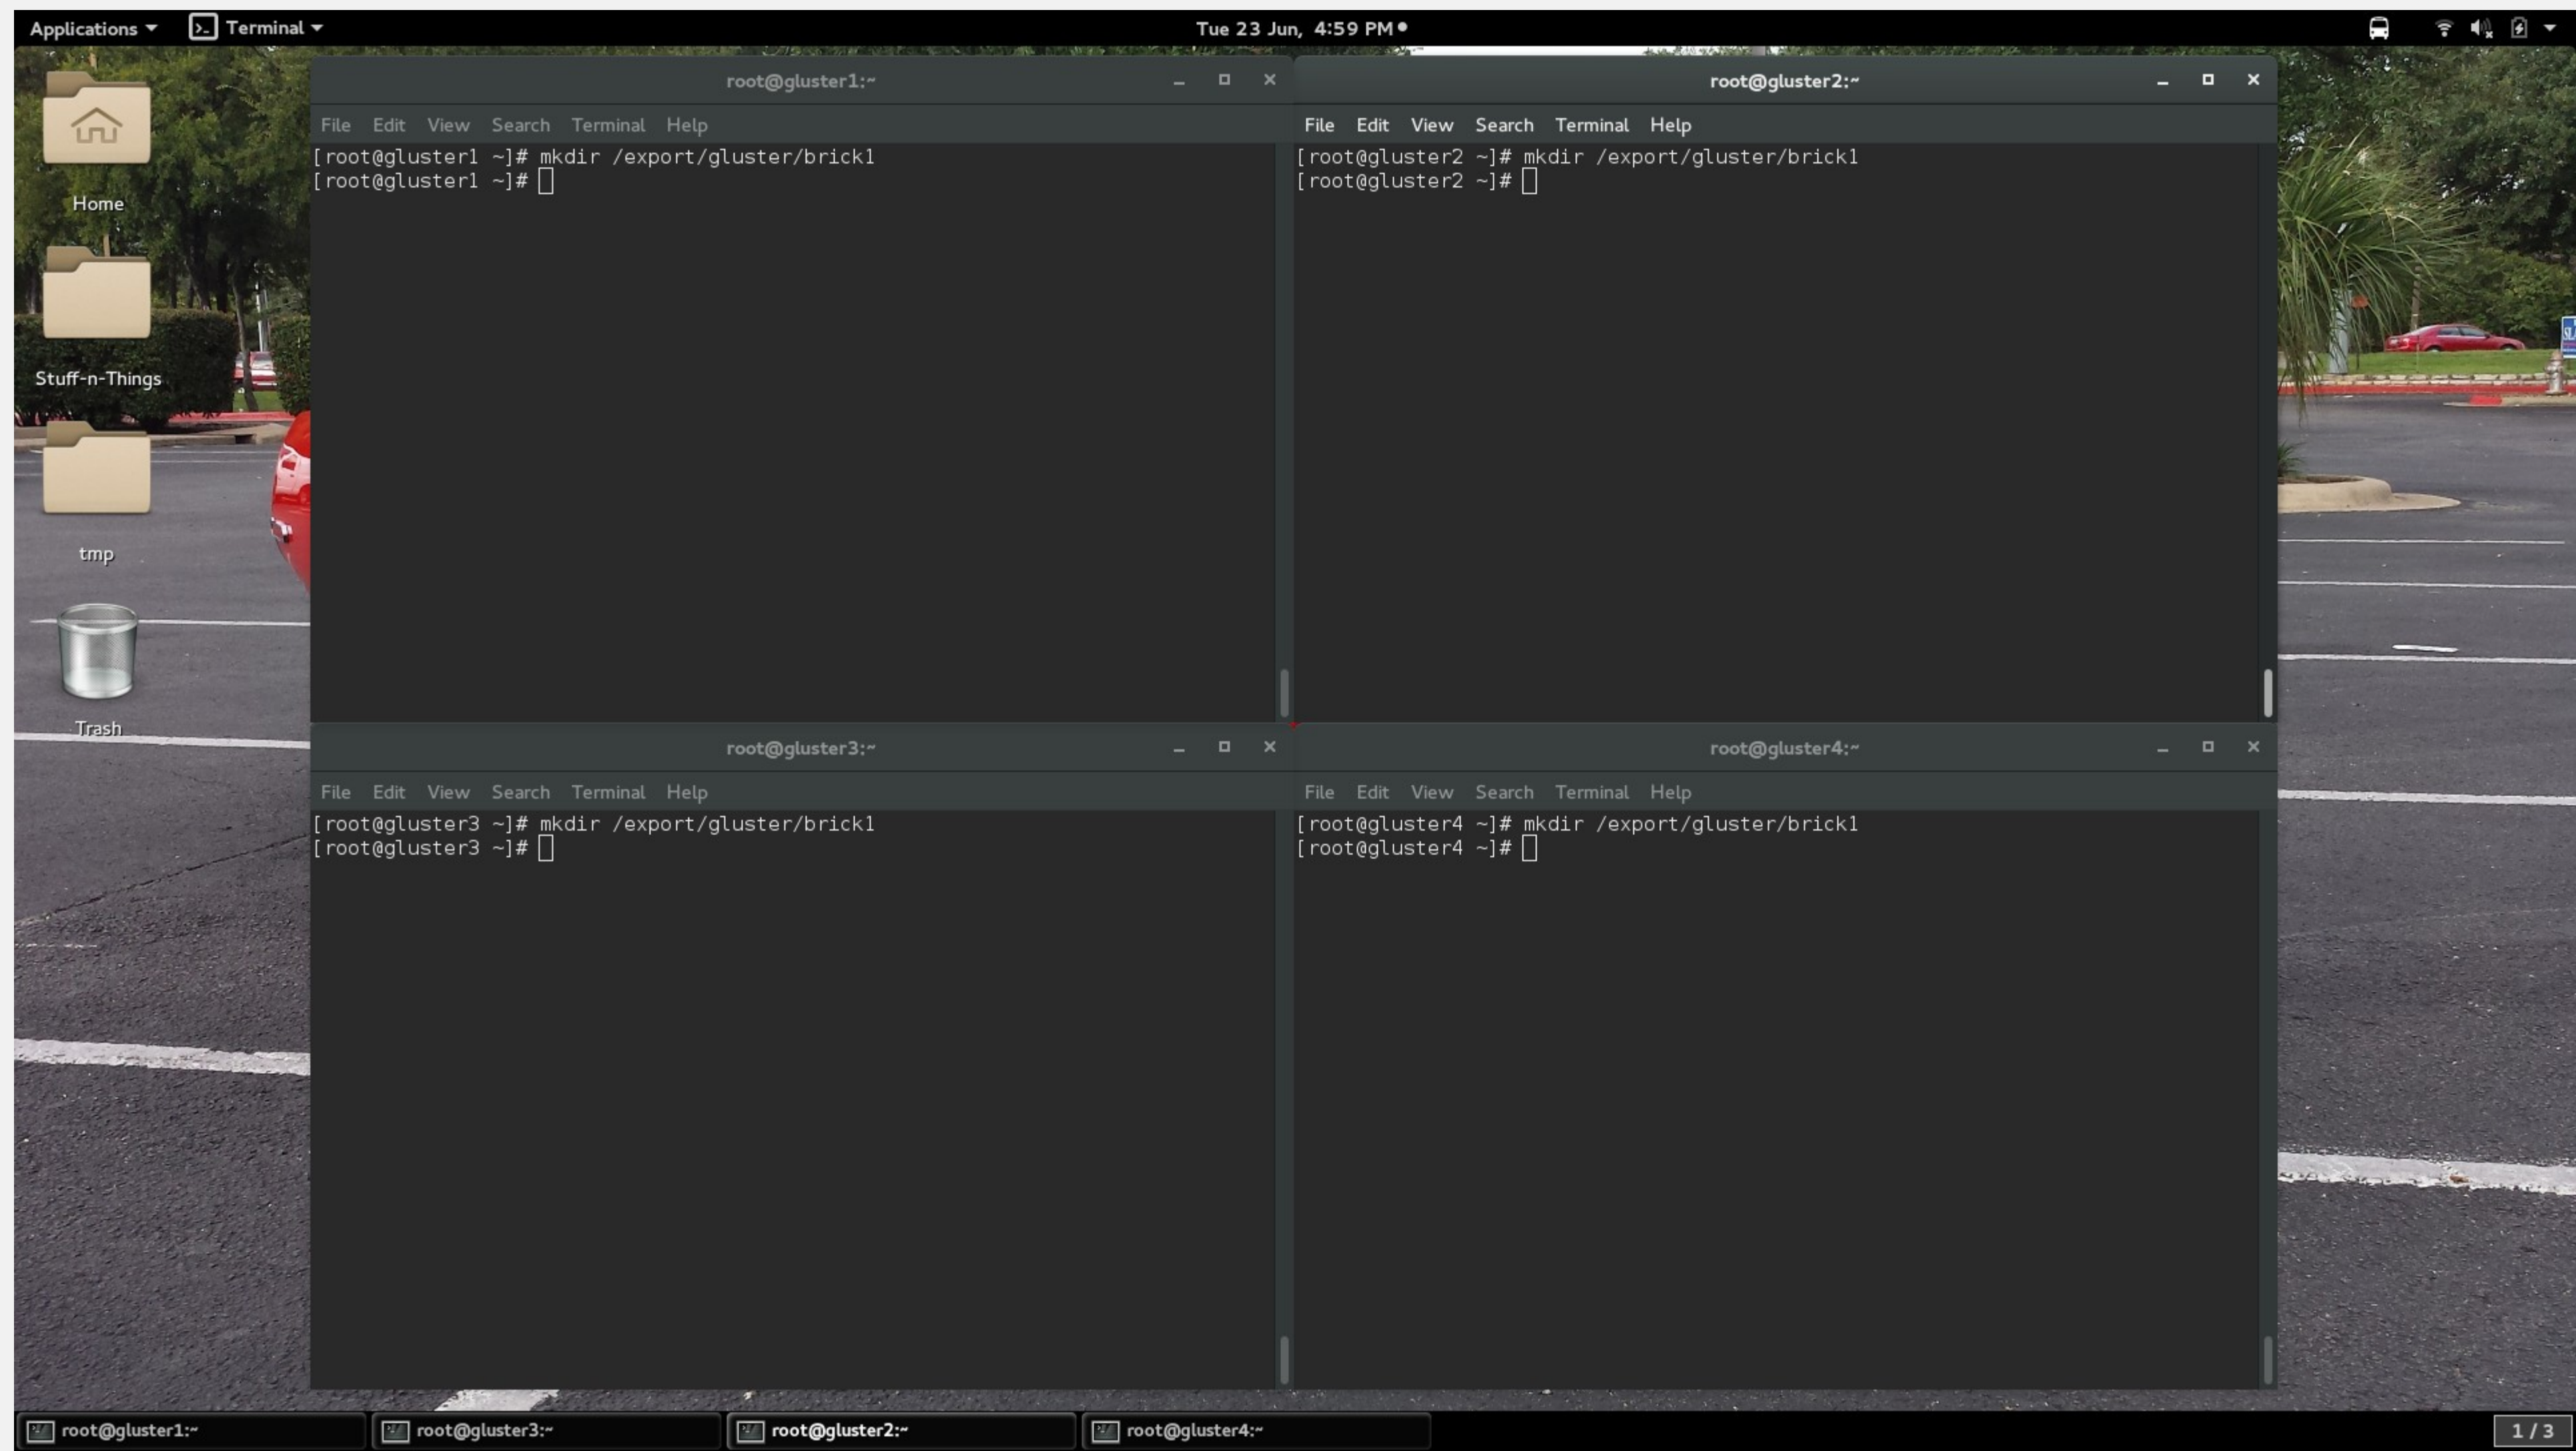

```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster peer probe gluster2.tc.redhat.com  
peer probe: success.  
[root@gluster1 ~]# gluster peer probe gluster3.tc.redhat.com  
peer probe: success.  
[root@gluster1 ~]# gluster peer probe gluster4.tc.redhat.com  
peer probe: success.  
[root@gluster1 ~]#
```

```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster peer probe gluster2.tc.redhat.com  
peer probe: success.  
[root@gluster1 ~]# gluster peer probe gluster3.tc.redhat.com  
peer probe: success.  
[root@gluster1 ~]# gluster peer probe gluster4.tc.redhat.com  
peer probe: success.  
[root@gluster1 ~]# gluster peer status  
Number of Peers: 3  
  
Hostname: gluster2.tc.redhat.com  
Uuid: a9cb18d7-1f62-45e9-9f1b-3f5fd1b4a7b6  
State: Peer in Cluster (Connected)  
  
Hostname: gluster3.tc.redhat.com  
Uuid: 03ddc0d1-f1f3-4421-8bba-de483e7c924a  
State: Peer in Cluster (Connected)  
  
Hostname: gluster4.tc.redhat.com  
Uuid: 2d995c17-c9df-47a2-8d29-7360012f4d11  
State: Peer in Cluster (Connected)  
[root@gluster1 ~]#
```

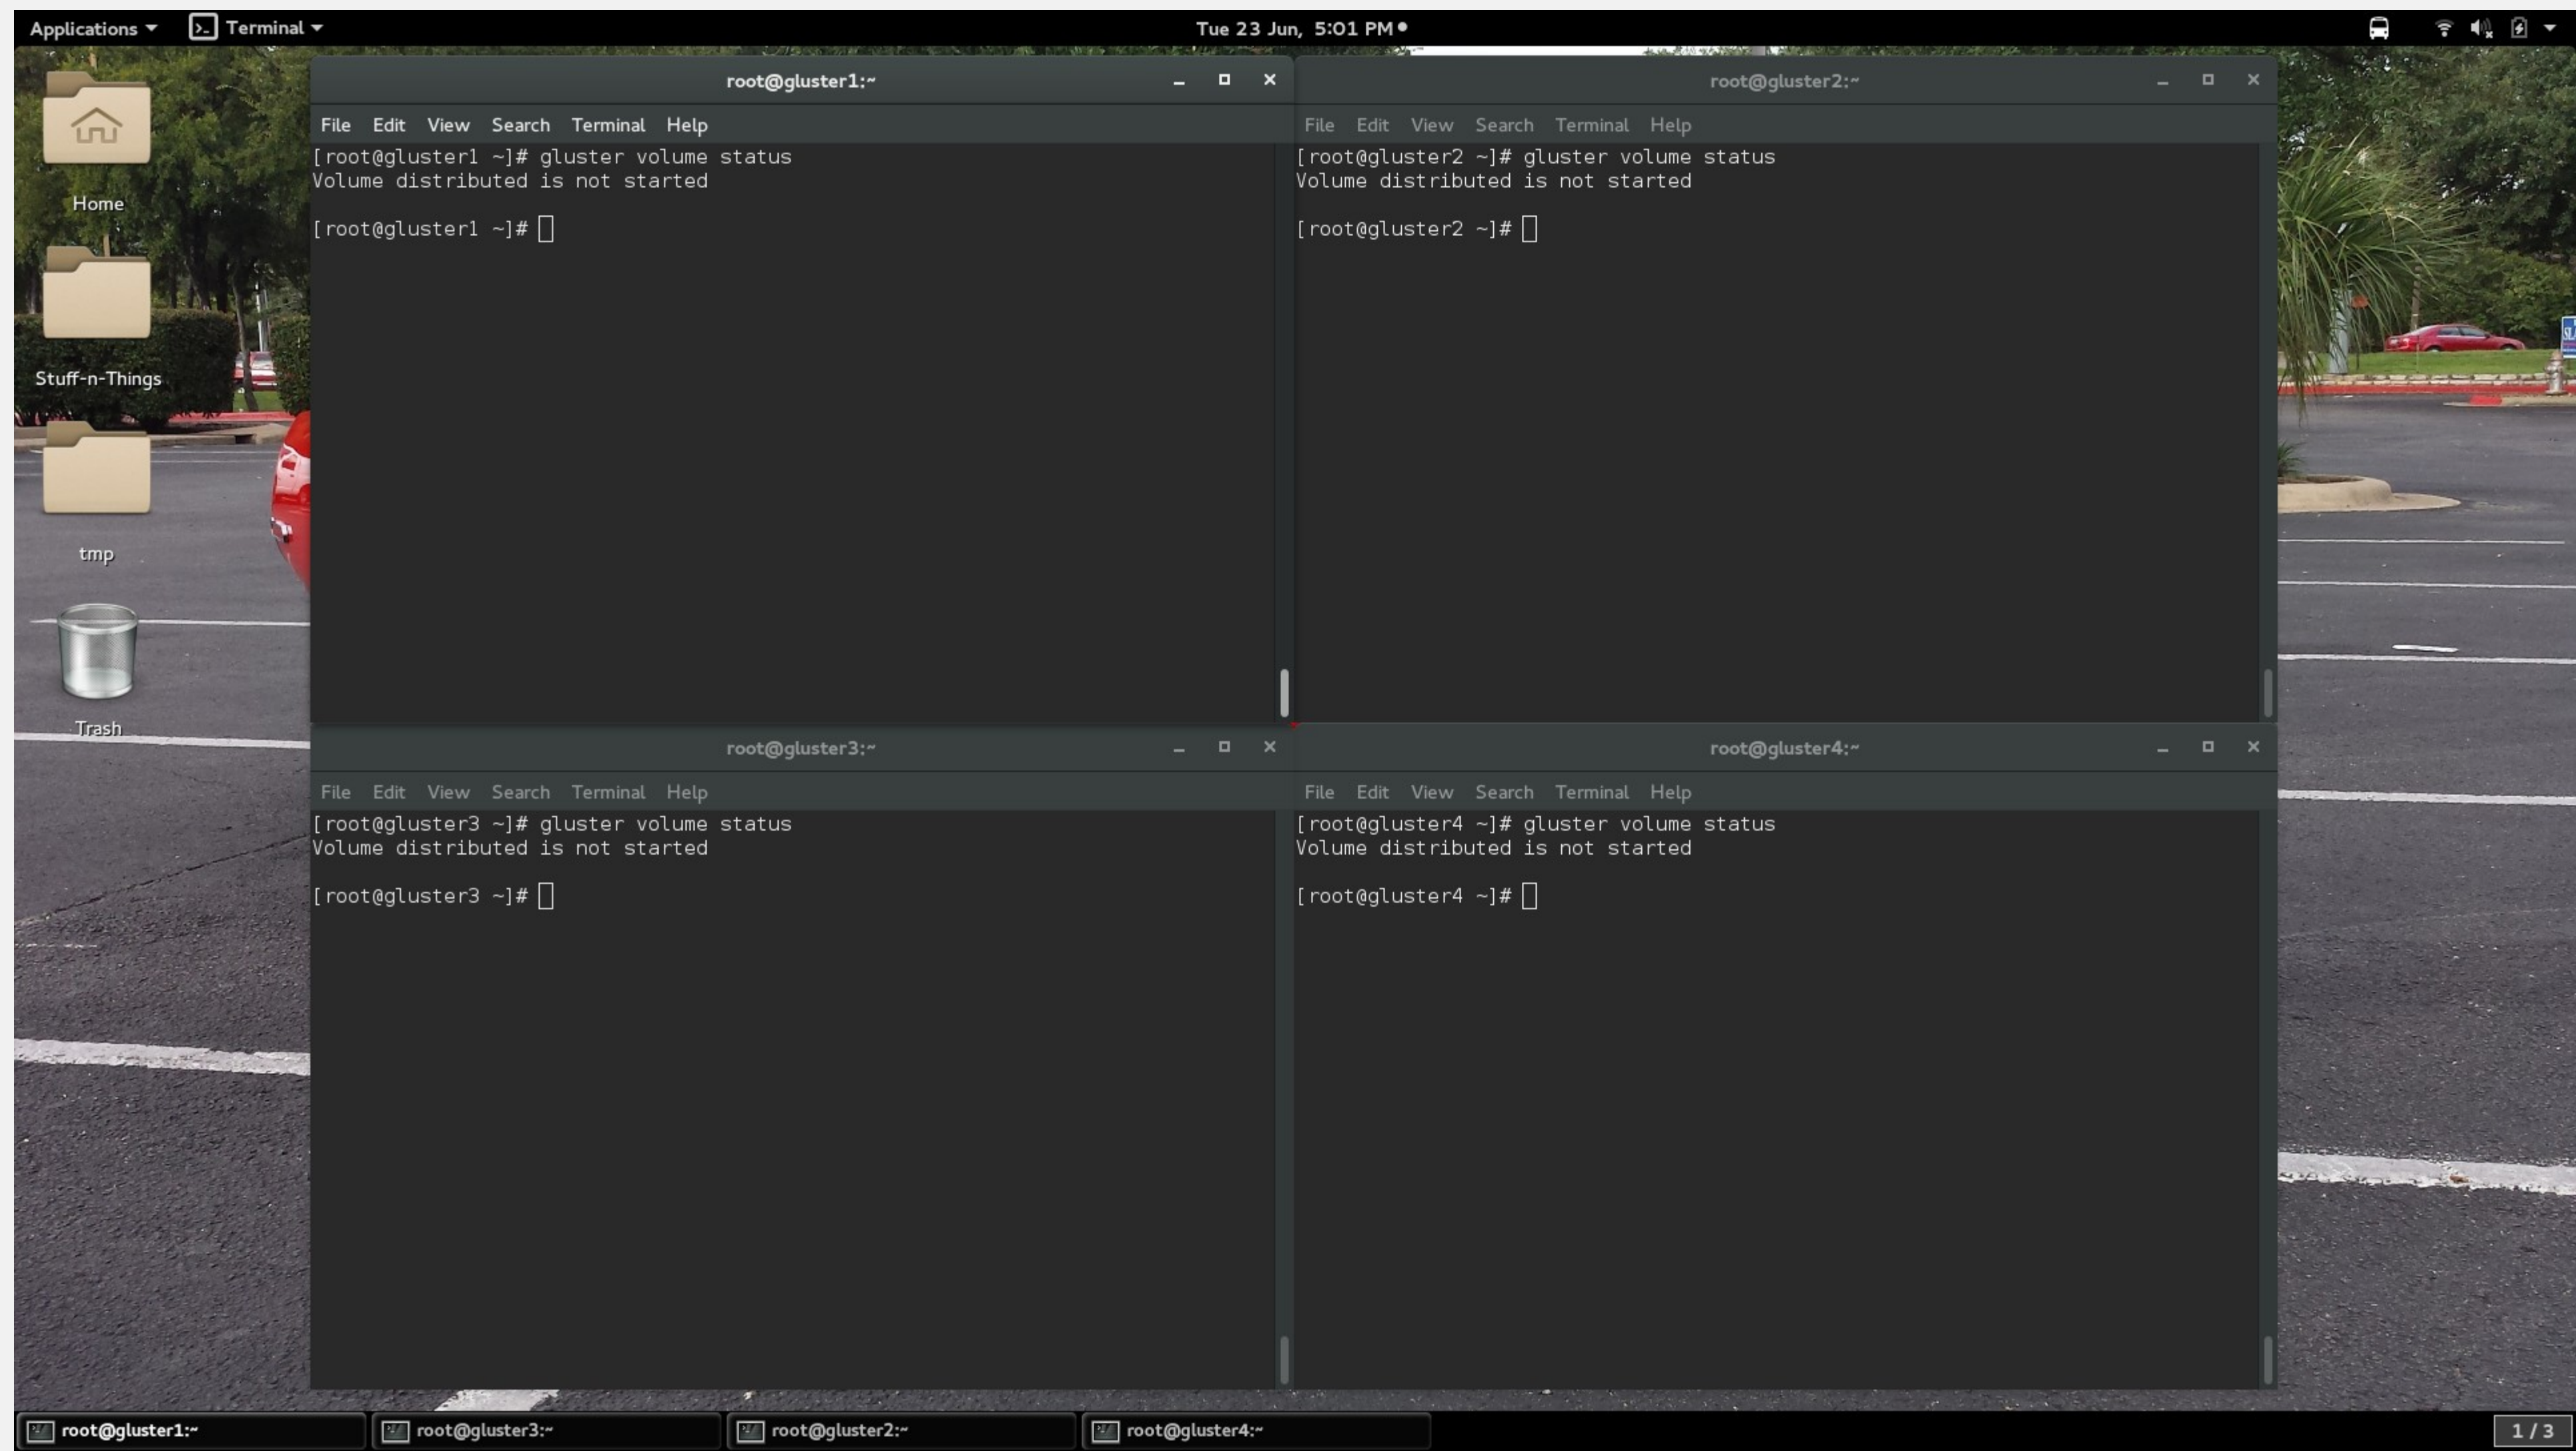
DISTRIBUTED VOLUMES


```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# mkdir /export/gluster/brick1  
[root@gluster1 ~]#
```




```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume create distributed gluster1.tc.redhat.com:/exp  
ort/gluster/brick1/ gluster2.tc.redhat.com:/export/gluster/brick1/ gluster3.tc.r  
edhat.com:/export/gluster/brick1/ gluster4.tc.redhat.com:/export/gluster/brick1/  
volume create: distributed: success: please start the volume to access data  
[root@gluster1 ~]#
```

```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume status  
Volume distributed is not started  
[root@gluster1 ~]#
```

```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume start distributed  
volume start: distributed: success  
[root@gluster1 ~]#
```



```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume start distributed  
volume start: distributed: success  
[root@gluster1 ~]# gluster volume status  
Status of volume: distributed  
Gluster process                TCP Port  RDMA Port  Online  Pid  
-----  
Brick gluster1.tc.redhat.com:/export/gluster/brick1 49152    0          Y       2584  
Brick gluster2.tc.redhat.com:/export/gluster/brick1 49152    0          Y       2442  
Brick gluster3.tc.redhat.com:/export/gluster/brick1 49152    0          Y       2439  
Brick gluster4.tc.redhat.com:/export/gluster/brick1 49152    0          Y       2435  
NFS Server on localhost          N/A      N/A        N       N/A  
NFS Server on gluster3.tc.redhat.com N/A      N/A        N       N/A  
NFS Server on gluster4.tc.redhat.com N/A      N/A        N       N/A  
NFS Server on gluster2.tc.redhat.com N/A      N/A        N       N/A  
  
Task Status of Volume distributed  
-----  
There are no active volume tasks  
  
[root@gluster1 ~]#
```


Applications

Terminal

Home

Stuff-n-Things

tmp

Trash

root@gluster1:~

File Edit View Search Terminal Help

[root@gluster1 ~]# gluster volume status

Status of volume: distributed

Gluster process

	TCP Port	RDMA Port	Online	Pid
Brick gluster1.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2584
Brick gluster2.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2442
Brick gluster3.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2439
Brick gluster4.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2435
NFS Server on localhost	N/A	N/A	N	N/A
NFS Server on gluster4.tc.redhat.com	N/A	N/A	N	N/A
NFS Server on gluster3.tc.redhat.com	N/A	N/A	N	N/A
NFS Server on gluster2.tc.redhat.com	N/A	N/A	N	N/A

Task Status of Volume distributed

There are no active volume tasks

[root@gluster1 ~]#

root@gluster2:~

File Edit View Search Terminal Help

[root@gluster2 ~]# gluster volume status

Status of volume: distributed

Gluster process

	TCP Port	RDMA Port	Online	Pid
Brick gluster1.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2584
Brick gluster2.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2442
Brick gluster3.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2439
Brick gluster4.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2435
NFS Server on localhost	N/A	N/A	N	N/A
NFS Server on 192.168.122.11	N/A	N/A	N	N/A
NFS Server on gluster3.tc.redhat.com	N/A	N/A	N	N/A
NFS Server on gluster4.tc.redhat.com	N/A	N/A	N	N/A

Task Status of Volume distributed

There are no active volume tasks

[root@gluster2 ~]#

root@gluster3:~

File Edit View Search Terminal Help

[root@gluster3 ~]# gluster volume status

Status of volume: distributed

Gluster process

	TCP Port	RDMA Port	Online	Pid
Brick gluster1.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2584
Brick gluster2.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2442
Brick gluster3.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2439
Brick gluster4.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2435
NFS Server on localhost	N/A	N/A	N	N/A
NFS Server on 192.168.122.11	N/A	N/A	N	N/A
NFS Server on gluster4.tc.redhat.com	N/A	N/A	N	N/A
NFS Server on gluster2.tc.redhat.com	N/A	N/A	N	N/A

Task Status of Volume distributed

There are no active volume tasks

[root@gluster3 ~]#

root@gluster4:~

File Edit View Search Terminal Help

[root@gluster4 ~]# gluster volume status

Status of volume: distributed

Gluster process

	TCP Port	RDMA Port	Online	Pid
Brick gluster1.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2584
Brick gluster2.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2442
Brick gluster3.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2439
Brick gluster4.tc.redhat.com:/export/gluster/brick1	49152	0	Y	2435
NFS Server on localhost	N/A	N/A	N	N/A
NFS Server on 192.168.122.11	N/A	N/A	N	N/A
NFS Server on gluster3.tc.redhat.com	N/A	N/A	N	N/A
NFS Server on gluster2.tc.redhat.com	N/A	N/A	N	N/A

Task Status of Volume distributed

There are no active volume tasks

[root@gluster4 ~]#

root@gluster1:~

root@gluster3:~

root@gluster2:~

root@gluster4:~

1 / 3


```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume info  
  
Volume Name: distributed  
Type: Distribute  
Volume ID: 315c0f4f-fa57-4dc7-a21e-bb45958b663b  
Status: Started  
Snap Volume: no  
Number of Bricks: 4  
Transport-type: tcp  
Bricks:  
Brick1: gluster1.tc.redhat.com:/export/gluster/brick1  
Brick2: gluster2.tc.redhat.com:/export/gluster/brick1  
Brick3: gluster3.tc.redhat.com:/export/gluster/brick1  
Brick4: gluster4.tc.redhat.com:/export/gluster/brick1  
Options Reconfigured:  
performance.readdir-ahead: on  
snap-max-hard-limit: 256  
snap-max-soft-limit: 90  
auto-delete: disable  
[root@gluster1 ~]#
```


ACCESSING DISTRIBUTED VOLUMES


```
root@t540p:~  
File Edit View Search Terminal Help  
[tcameron@t540p Desktop]$ su -  
fr3Password:  
[root@t540p ~]# mount -t glusterfs gluster1.tc.redhat.com:/distributed /mnt/gluster/  
[root@t540p ~]# df -h /mnt/gluster/  
Filesystem                                Size  Used Avail Use% Mounted on  
gluster1.tc.redhat.com:/distributed        24G   129M   24G    1% /mnt/gluster  
[root@t540p ~]#
```

TEST FILE CREATION

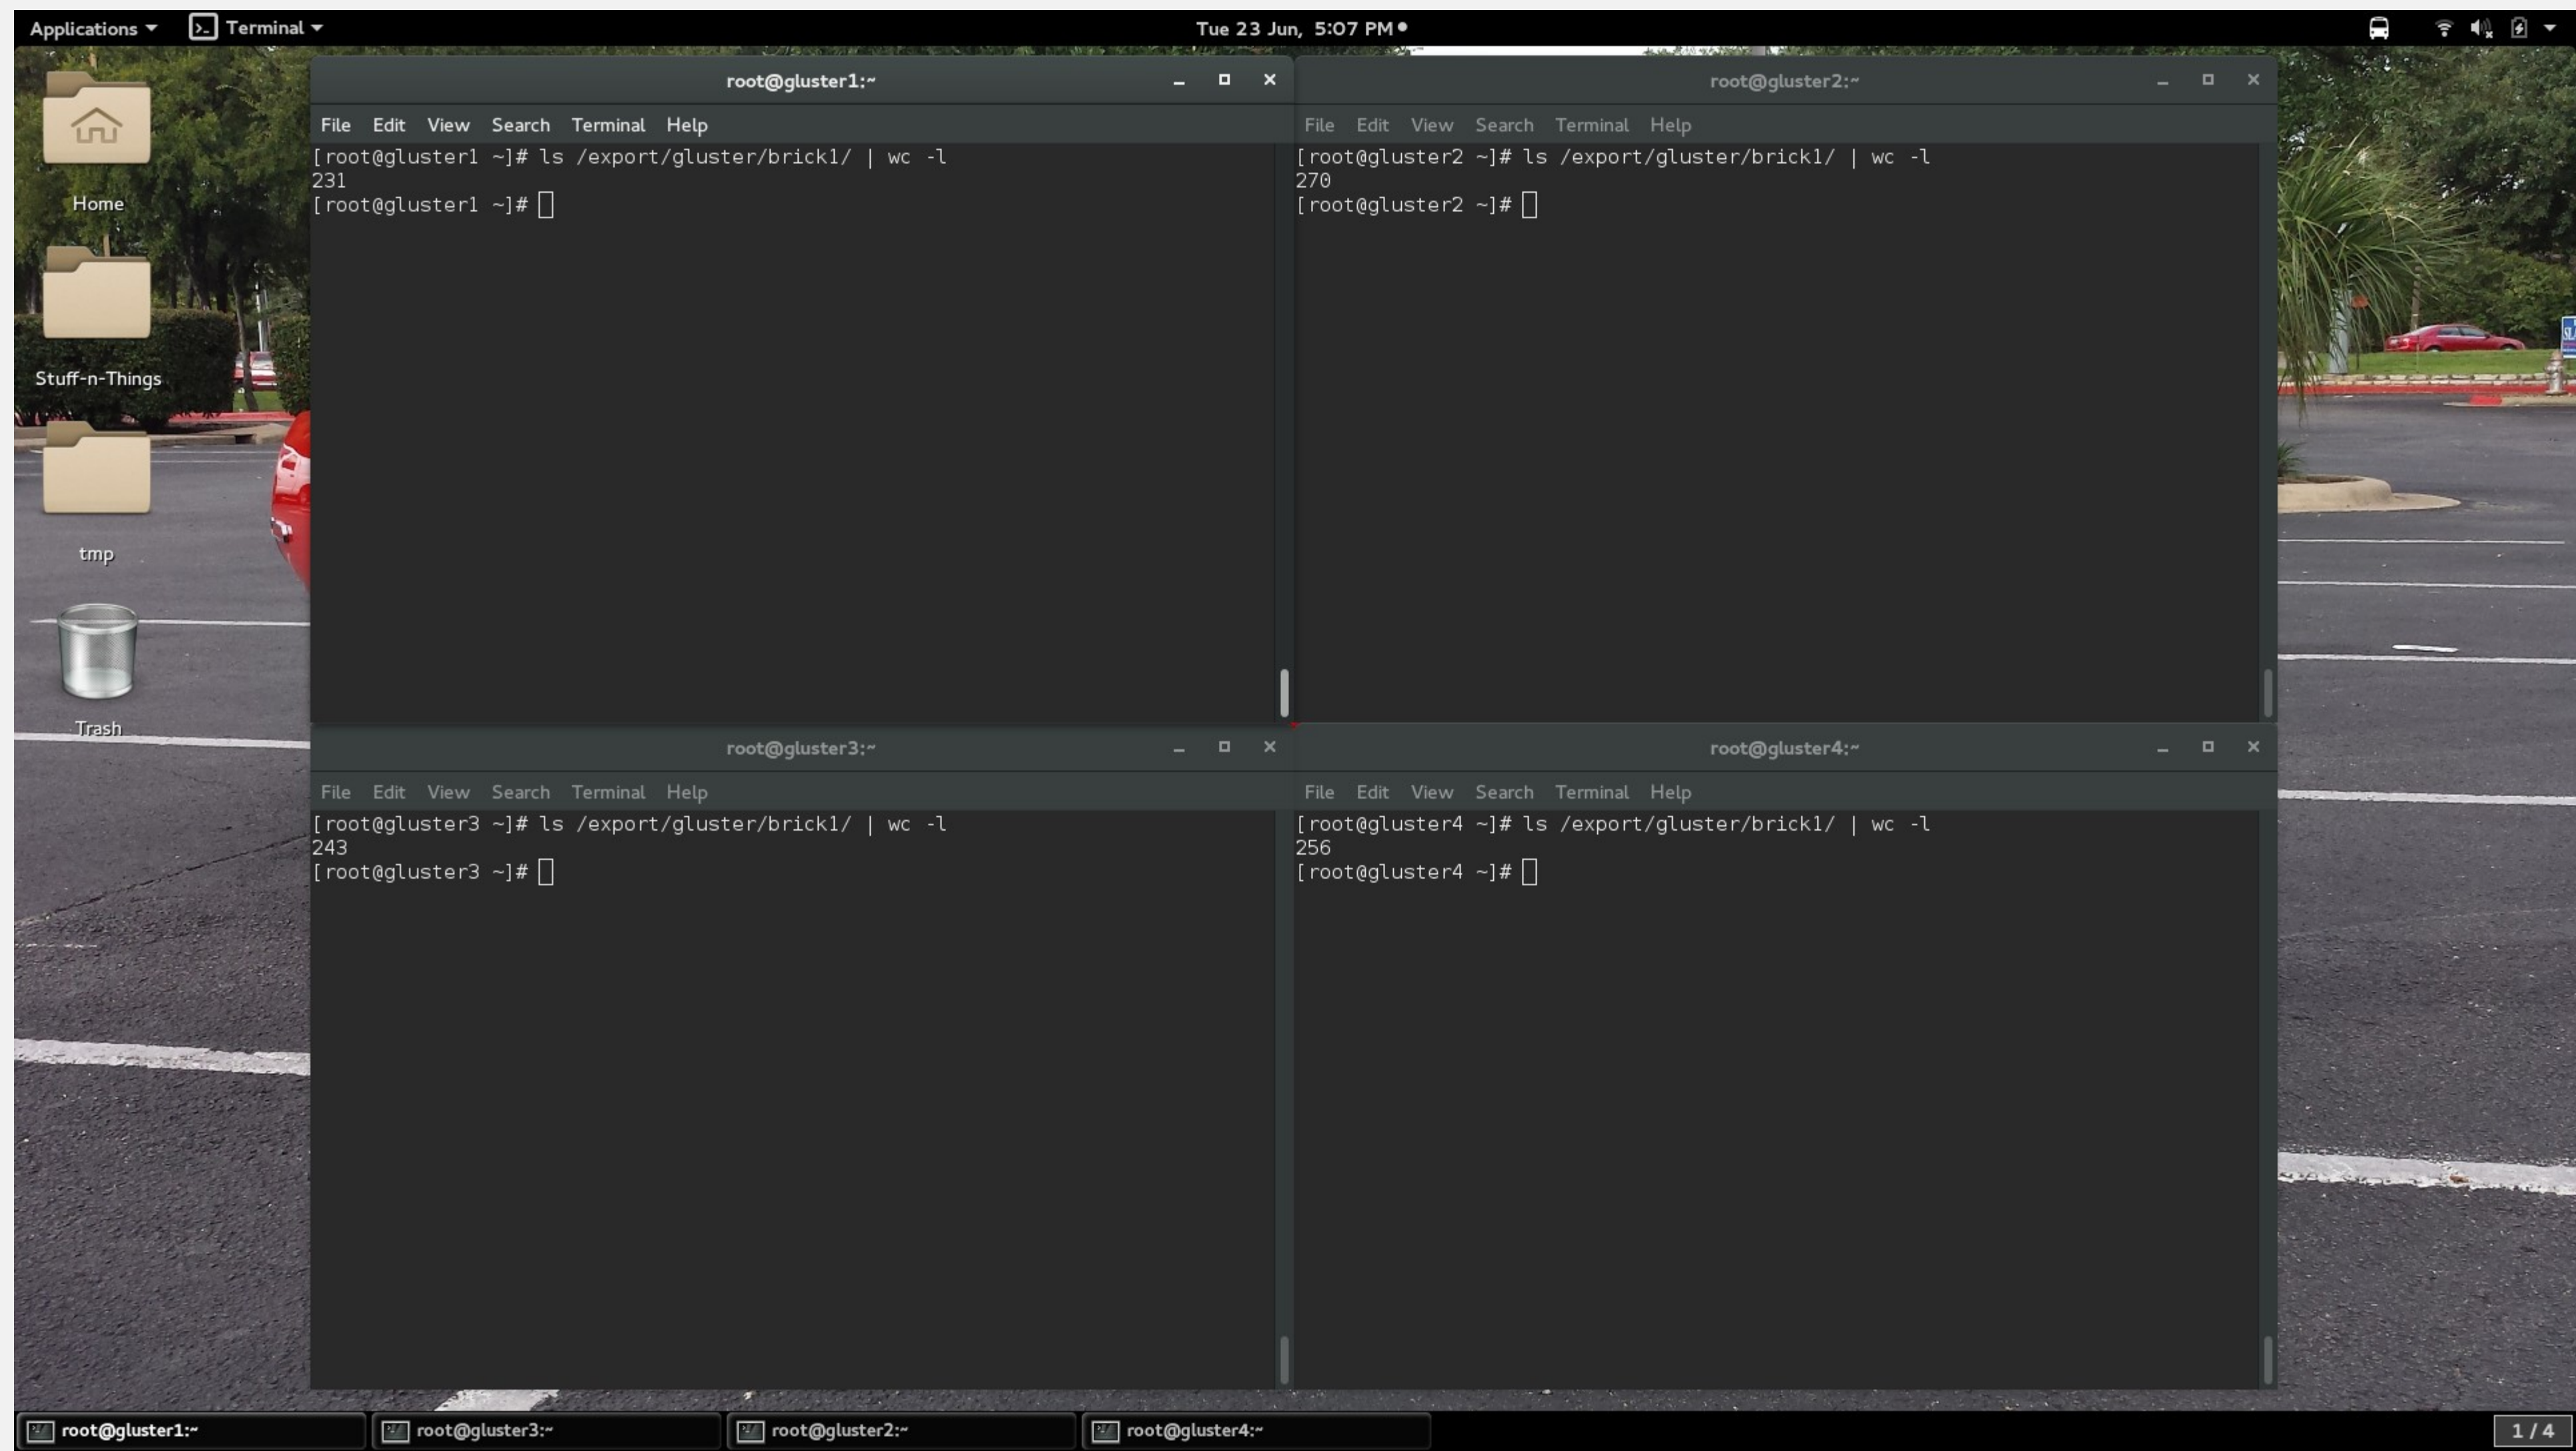
- Note that on the client, creating 1000 files indicates that there are 1000 files in one mount point


```
root@t540p:/mnt/gluster
File Edit View Search Terminal Help
[root@t540p ~]# cd /mnt/gluster/
[root@t540p gluster]# ls
[root@t540p gluster]# for i in $(seq 1 1000); do echo $i > $i; done
[root@t540p gluster]# ls | wc -l
1000
[root@t540p gluster]#
```

TEST FILE CREATION

- But when you look at the bricks, the files are **roughly** distributed with about a quarter of them on each of the 4 machines.


```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# ls /export/gluster/brick1/ | wc -l  
231  
[root@gluster1 ~]#
```




```
root@t540p:/mnt
File Edit View Search Terminal Help
[root@t540p gluster]# cd ..
[root@t540p mnt]# umount gluster
[root@t540p mnt]#
```

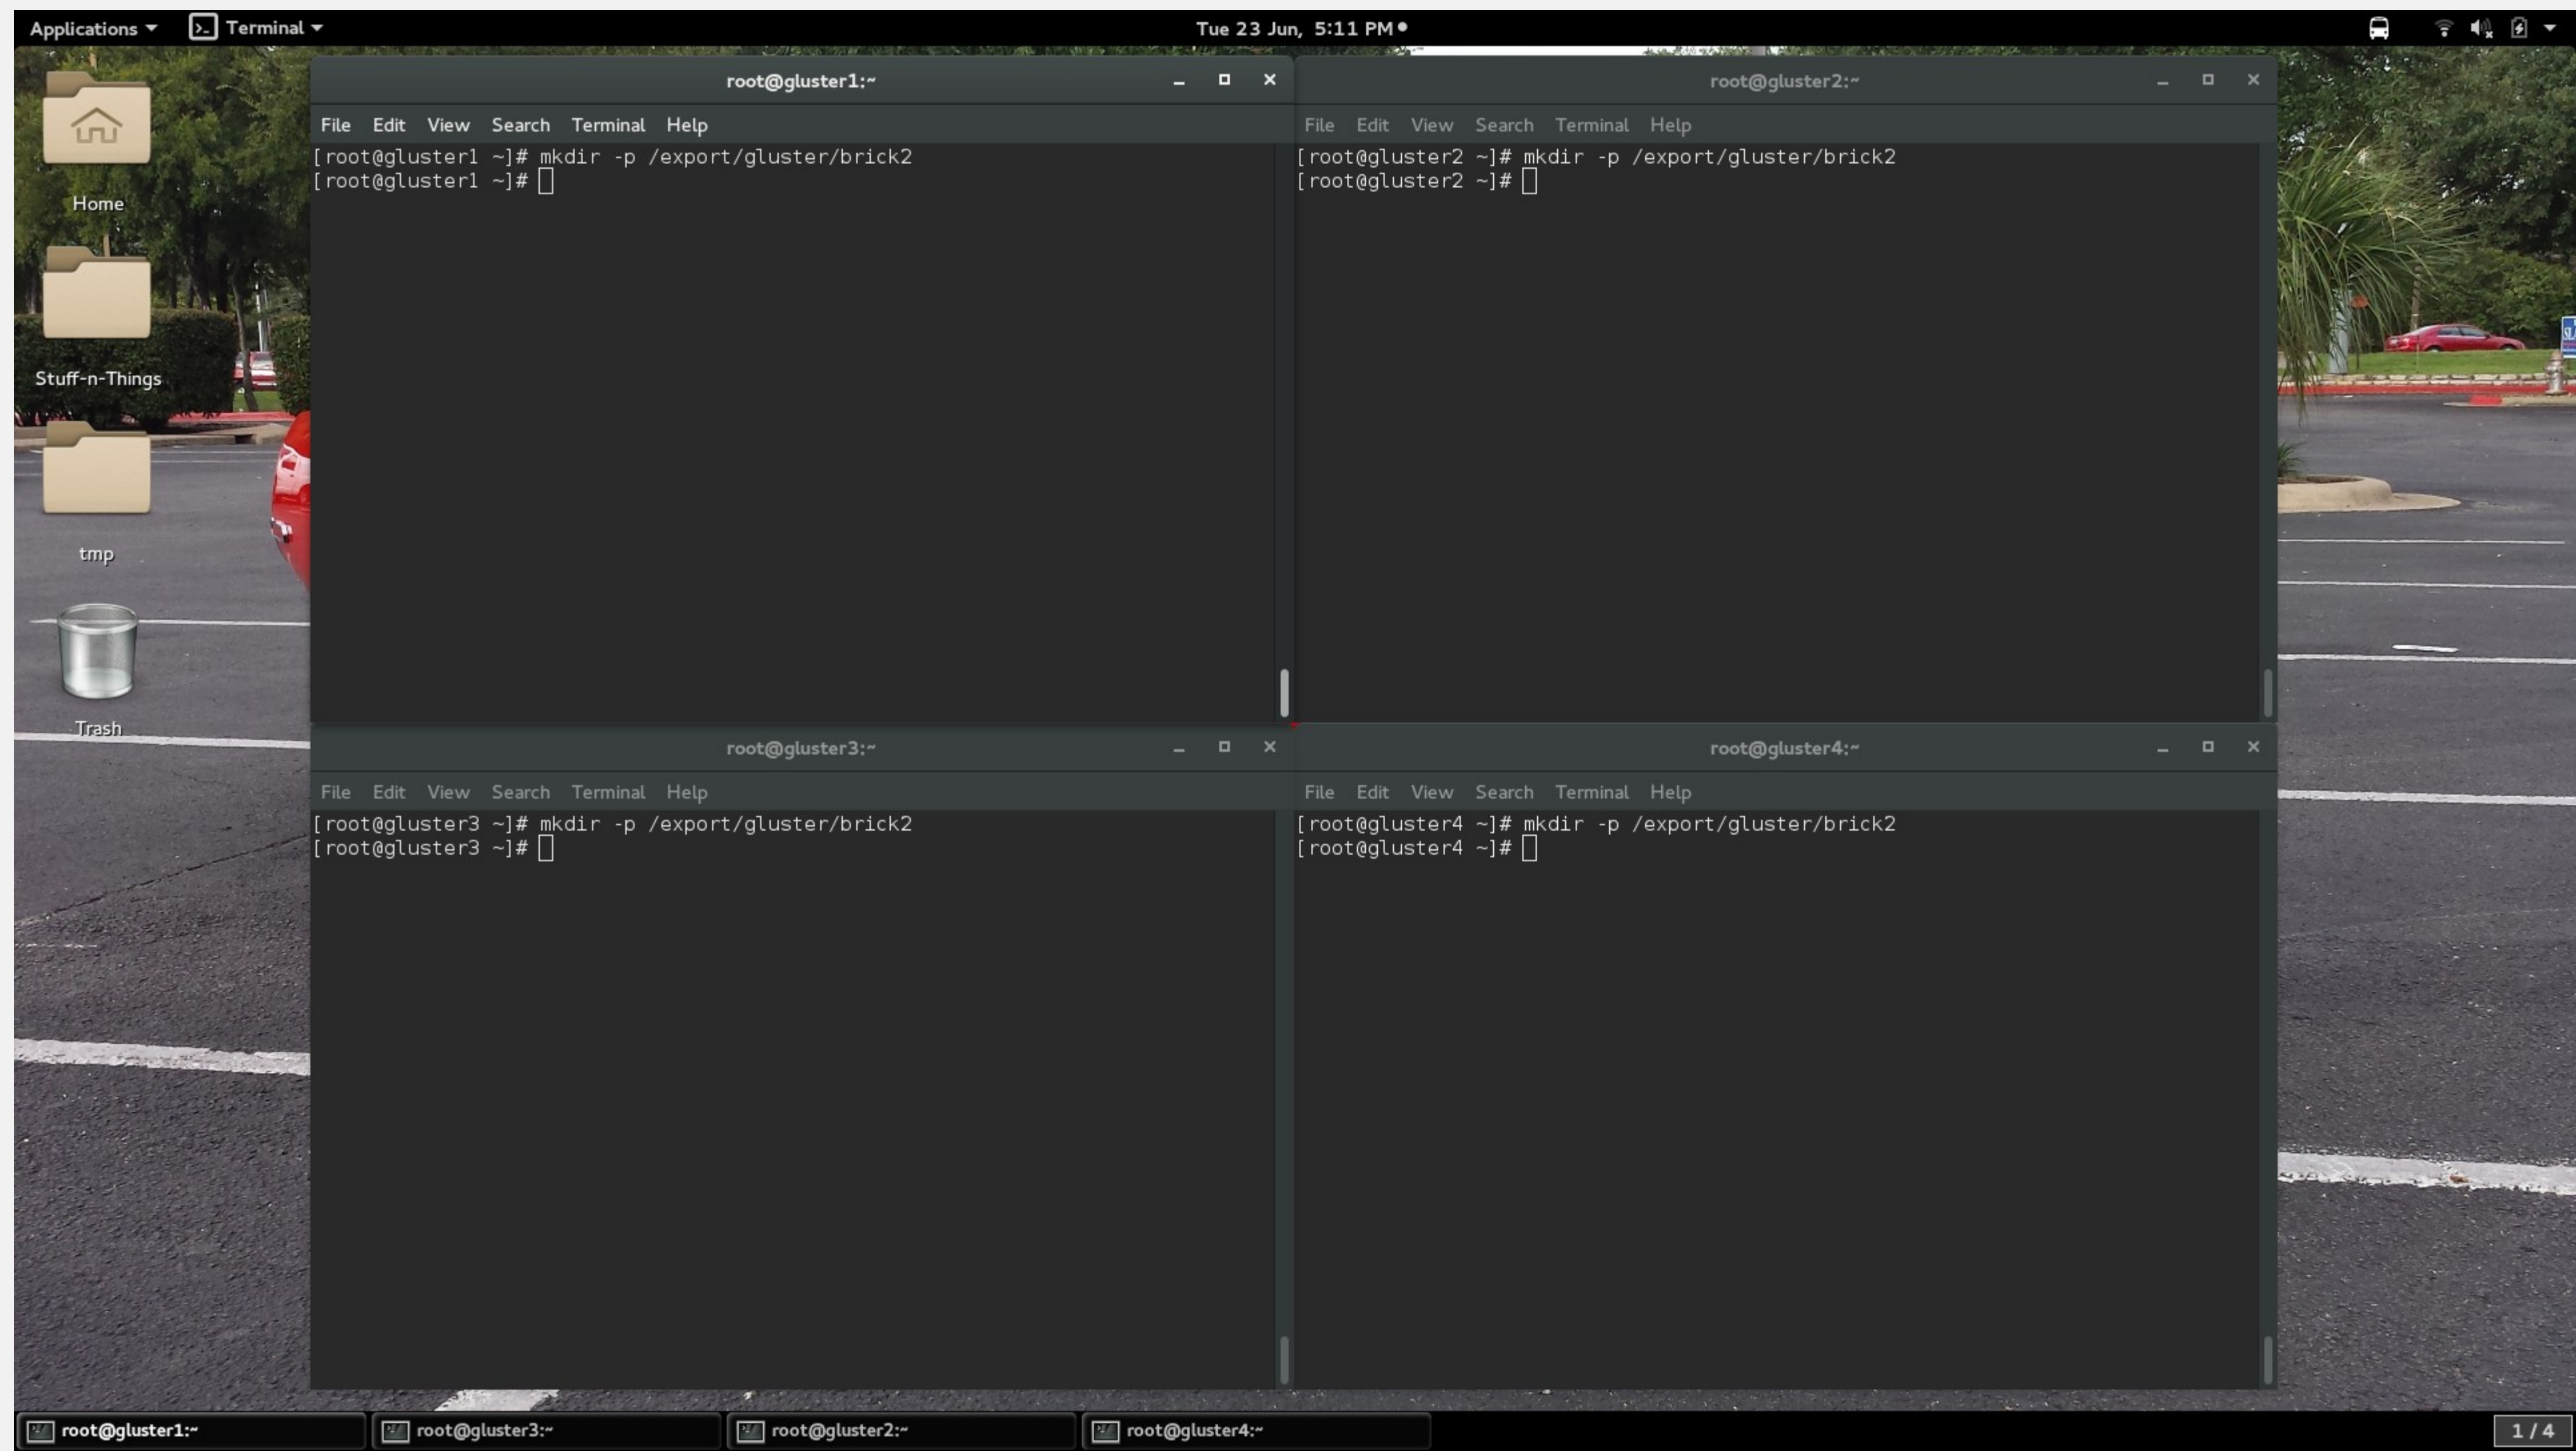
```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume stop distributed  
Stopping volume will make its data inaccessible. Do you want to continue? (y/n)  
y  
volume stop: distributed: success  
[root@gluster1 ~]#
```



```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume delete distributed  
Deleting volume will erase all information about the volume. Do you want to continue? (y/n) y  
volume delete: distributed: success  
[root@gluster1 ~]#
```


REPLICATED VOLUMES


```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# mkdir -p /export/gluster/brick2  
[root@gluster1 ~]#
```




```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume create replicated replica 2 gluster1.tc.redhat  
.com:/export/gluster/brick2/ gluster2.tc.redhat.com:/export/gluster/brick2/  
volume create: replicated: success: please start the volume to access data  
[root@gluster1 ~]#
```

```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume start replicated  
volume start: replicated: success  
[root@gluster1 ~]#
```



```
root@gluster1:~
File Edit View Search Terminal Help
[root@gluster1 ~]# gluster volume status
Volume distributed is not started

Status of volume: replicated
Gluster process                TCP Port  RDMA Port  Online  Pid
-----
Brick gluster1.tc.redhat.com:/export/gluster/brick2 49153      0           Y       2824
Brick gluster2.tc.redhat.com:/export/gluster/brick2 49153      0           Y       2615
NFS Server on localhost N/A        N/A         N       N/A
Self-heal Daemon on localhost N/A        N/A         Y       2846
NFS Server on gluster4.tc.redhat.com N/A        N/A         N       N/A
Self-heal Daemon on gluster4.tc.redhat.com N/A        N/A         Y       2584
NFS Server on gluster3.tc.redhat.com N/A        N/A         N       N/A
Self-heal Daemon on gluster3.tc.redhat.com N/A        N/A         Y       2588
NFS Server on gluster2.tc.redhat.com N/A        N/A         N       N/A
Self-heal Daemon on gluster2.tc.redhat.com N/A        N/A         Y       2638

Task Status of Volume replicated
-----
There are no active volume tasks

[root@gluster1 ~]#
```

```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume info replicated  
  
Volume Name: replicated  
Type: Replicate  
Volume ID: f9418ba4-2d43-4f87-8d89-dbab82df0f2c  
Status: Started  
Snap Volume: no  
Number of Bricks: 1 x 2 = 2  
Transport-type: tcp  
Bricks:  
Brick1: gluster1.tc.redhat.com:/export/gluster/brick2  
Brick2: gluster2.tc.redhat.com:/export/gluster/brick2  
Options Reconfigured:  
performance.readdir-ahead: on  
auto-delete: disable  
snap-max-soft-limit: 90  
snap-max-hard-limit: 256  
[root@gluster1 ~]#
```


ACCESSING REPLICATED VOLUMES


```
root@t540p:~  
File Edit View Search Terminal Help  
[root@t540p ~]# mount -t glusterfs gluster1.tc.redhat.com:replicated /mnt/gluster/  
r/  
[root@t540p ~]# df -h /mnt/gluster/  
Filesystem                                Size  Used Avail Use% Mounted on  
gluster1.tc.redhat.com:replicated 5.9G   34M  5.9G   1% /mnt/gluster  
[root@t540p ~]#
```


ACCESSING REPLICATED VOLUMES

- This time, creating 1000 files from the client creates 1000 files on each of the replicas.

```
root@t540p:/mnt/gluster
File Edit View Search Terminal Help
[root@t540p ~]# cd /mnt/gluster/
[root@t540p gluster]# for i in $(seq 1 1000); do echo $i > $i; done
[root@t540p gluster]# ls | wc -l
1000
[root@t540p gluster]#
```



```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# ls /export/gluster/brick2/ | wc -l  
1000  
[root@gluster1 ~]#
```

```
root@t540p:/mnt
File Edit View Search Terminal Help
[root@t540p gluster]# cd ..
[root@t540p mnt]# umount gluster/
[root@t540p mnt]#
```

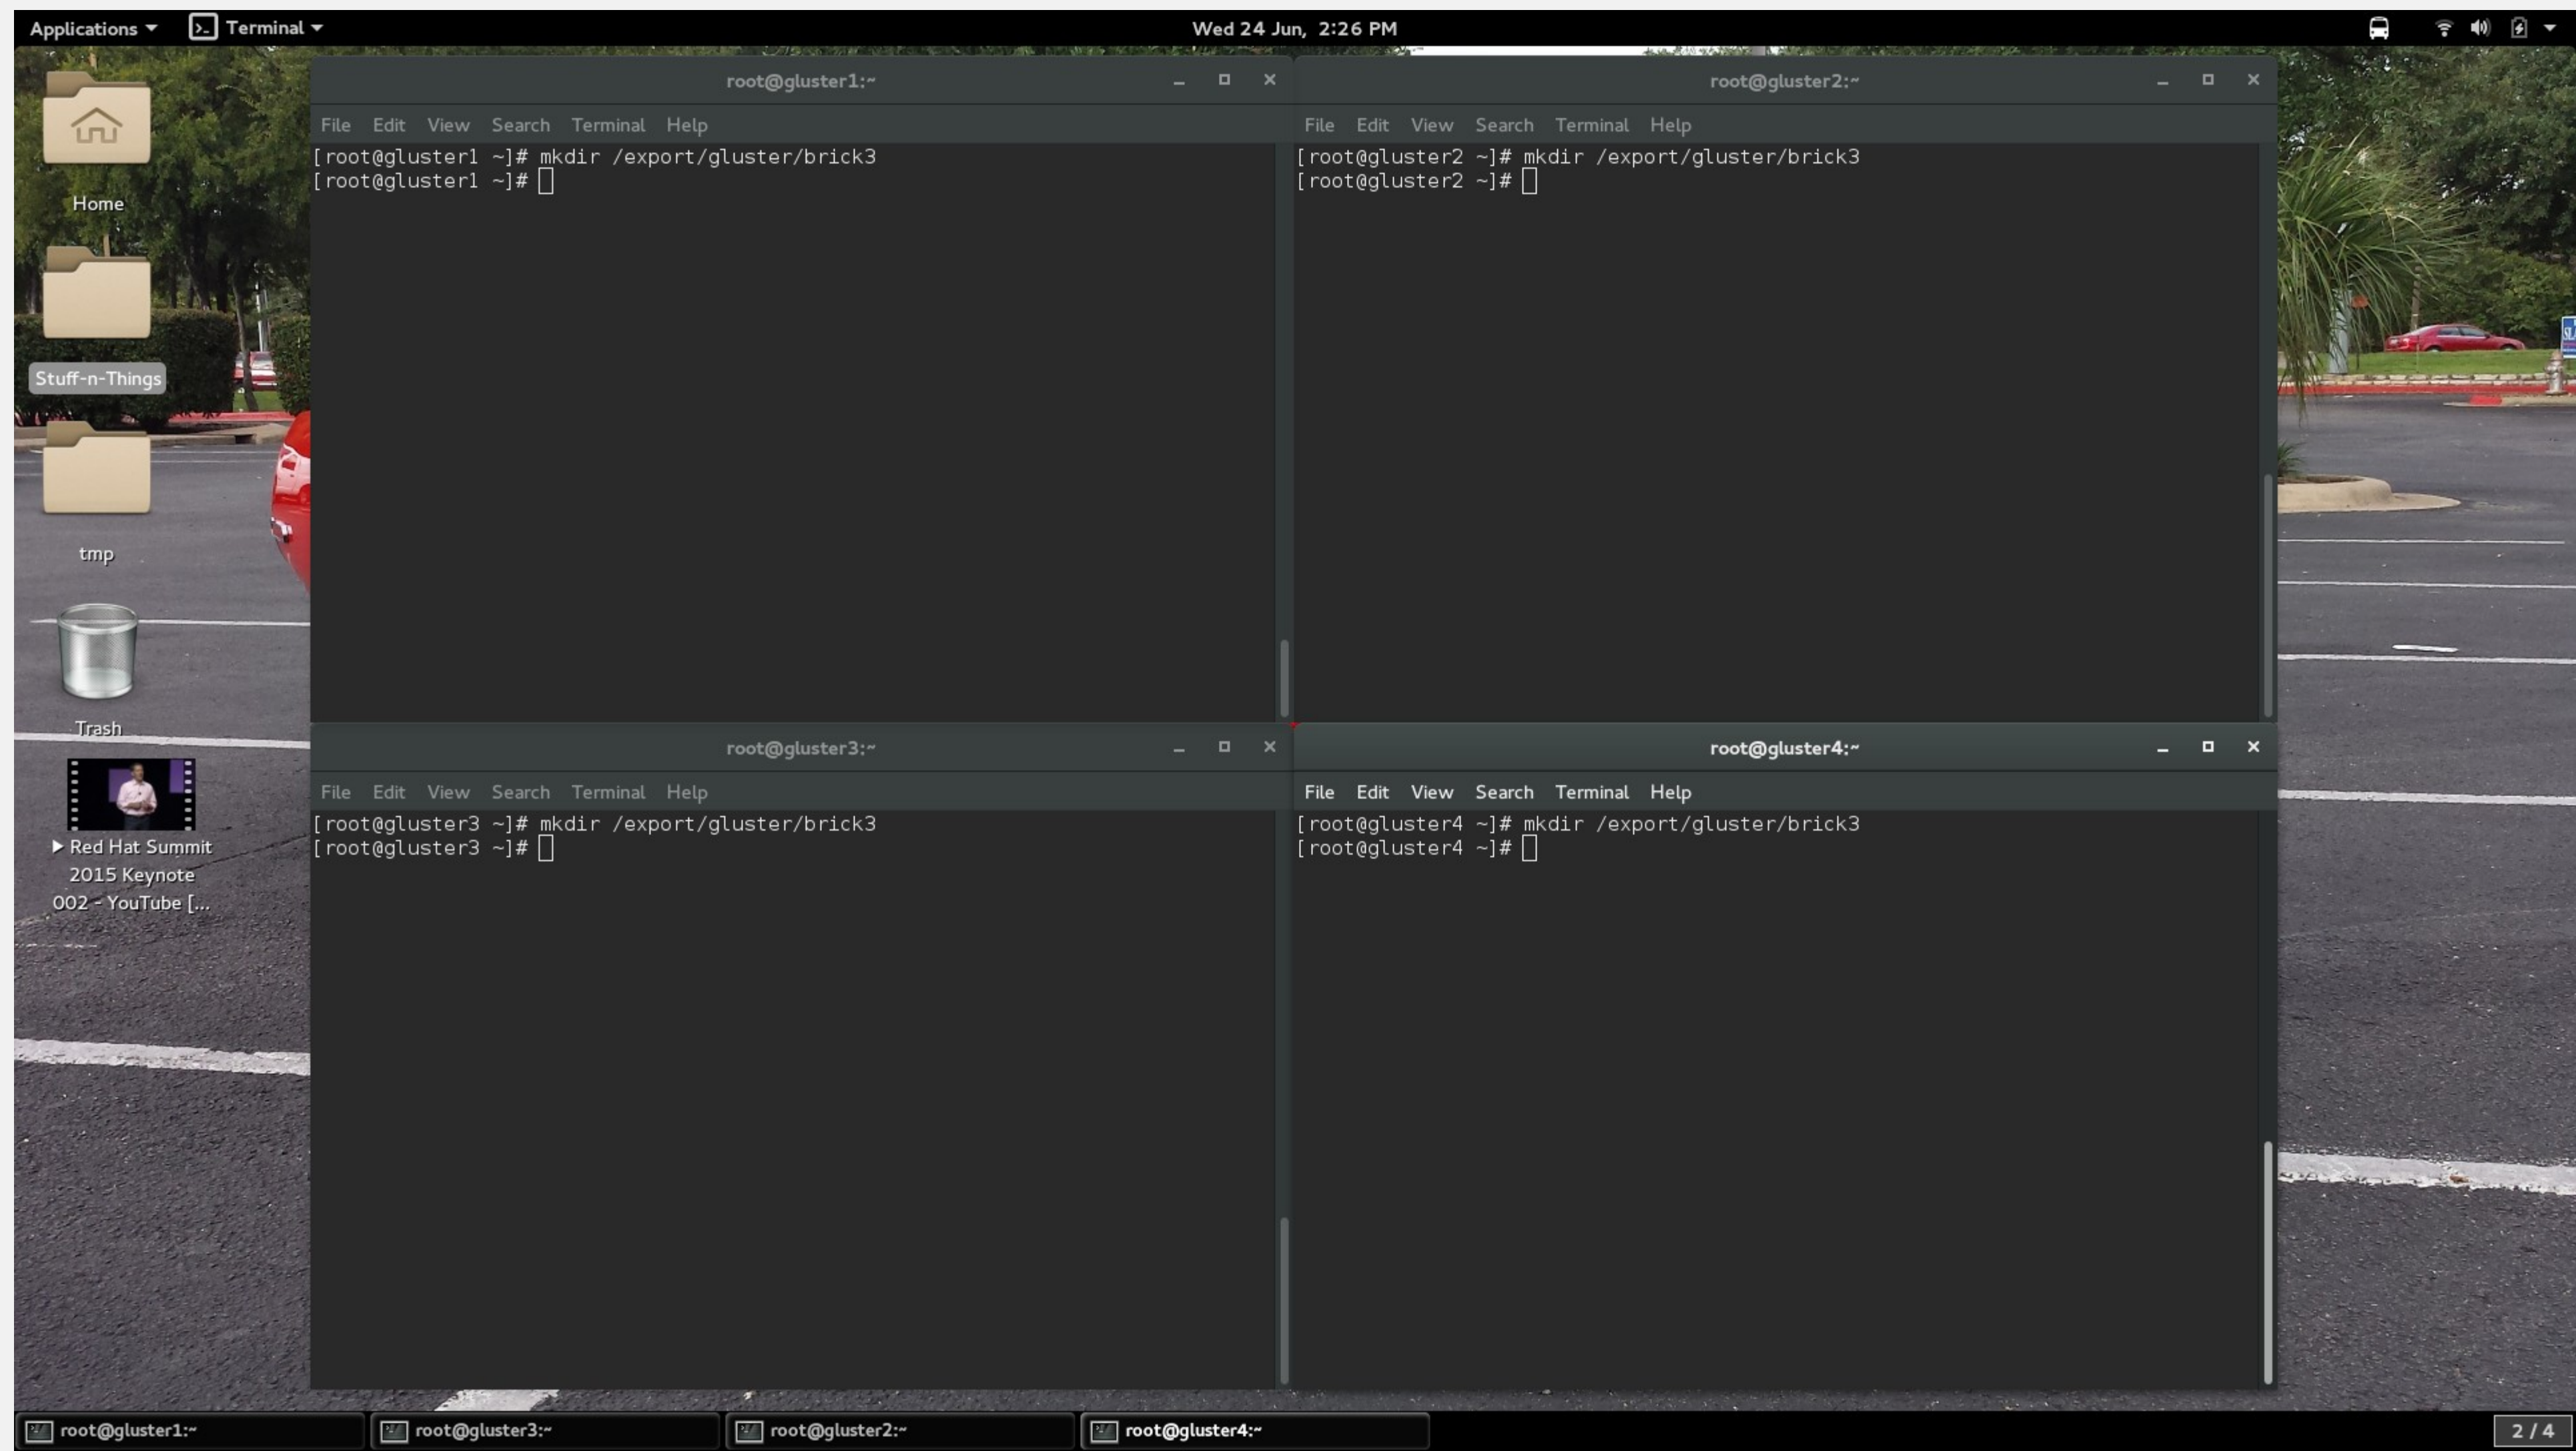


```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume stop replicated  
Stopping volume will make its data inaccessible. Do you want to continue? (y/n)  
y  
volume stop: replicated: success  
[root@gluster1 ~]#
```

```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume delete replicated  
Deleting volume will erase all information about the volume. Do you want to continue? (y/n) y  
volume delete: replicated: success  
[root@gluster1 ~]#
```


DISTRIBUTED + REPLICATED VOLUMES


```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# mkdir /export/gluster/brick3  
[root@gluster1 ~]#
```

```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume create dist-rep replica 2 gluster1.tc.redhat.c  
om:/export/gluster/brick3/ gluster2.tc.redhat.com:/export/gluster/brick3/ gluste  
r3.tc.redhat.com:/export/gluster/brick3/ gluster4.tc.redhat.com:/export/gluster/  
brick3/  
volume create: dist-rep: success: please start the volume to access data  
[root@gluster1 ~]#
```



```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume start dist-rep  
volume start: dist-rep: success  
[root@gluster1 ~]#
```

```
root@gluster1:~
File Edit View Search Terminal Help
[root@gluster1 ~]# gluster volume status dist-rep
Status of volume: dist-rep
Gluster process                                TCP Port  RDMA Port  Online  Pid
-----
Brick gluster1.tc.redhat.com:/export/gluster/brick3 49154      0           Y       3628
Brick gluster2.tc.redhat.com:/export/gluster/brick3 49154      0           Y       3504
Brick gluster3.tc.redhat.com:/export/gluster/brick3 49153      0           Y       3488
Brick gluster4.tc.redhat.com:/export/gluster/brick3 49153      0           Y       3488
NFS Server on localhost                          2049      0           Y       3642
Self-heal Daemon on localhost                    N/A       N/A         Y       3651
NFS Server on gluster2.tc.redhat.com              2049      0           Y       3519
Self-heal Daemon on gluster2.tc.redhat.com        N/A       N/A         Y       3526
NFS Server on gluster4.tc.redhat.com              2049      0           Y       3502
Self-heal Daemon on gluster4.tc.redhat.com        N/A       N/A         Y       3510
NFS Server on gluster3.tc.redhat.com              2049      0           Y       3503
Self-heal Daemon on gluster3.tc.redhat.com        N/A       N/A         Y       3510

Task Status of Volume dist-rep
-----
There are no active volume tasks

[root@gluster1 ~]#
```



```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume info dist-rep  
  
Volume Name: dist-rep  
Type: Distributed-Replicate  
Volume ID: 816828a4-2b89-4d10-b4de-ad1166649e0c  
Status: Started  
Snap Volume: no  
Number of Bricks: 2 x 2 = 4  
Transport-type: tcp  
Bricks:  
Brick1: gluster1.tc.redhat.com:/export/gluster/brick3  
Brick2: gluster2.tc.redhat.com:/export/gluster/brick3  
Brick3: gluster3.tc.redhat.com:/export/gluster/brick3  
Brick4: gluster4.tc.redhat.com:/export/gluster/brick3  
Options Reconfigured:  
performance.readdir-ahead: on  
auto-delete: disable  
snap-max-soft-limit: 90  
snap-max-hard-limit: 256  
[root@gluster1 ~]#
```


ACCESSING DISTRIBUTED + REPLICATED VOLUMES

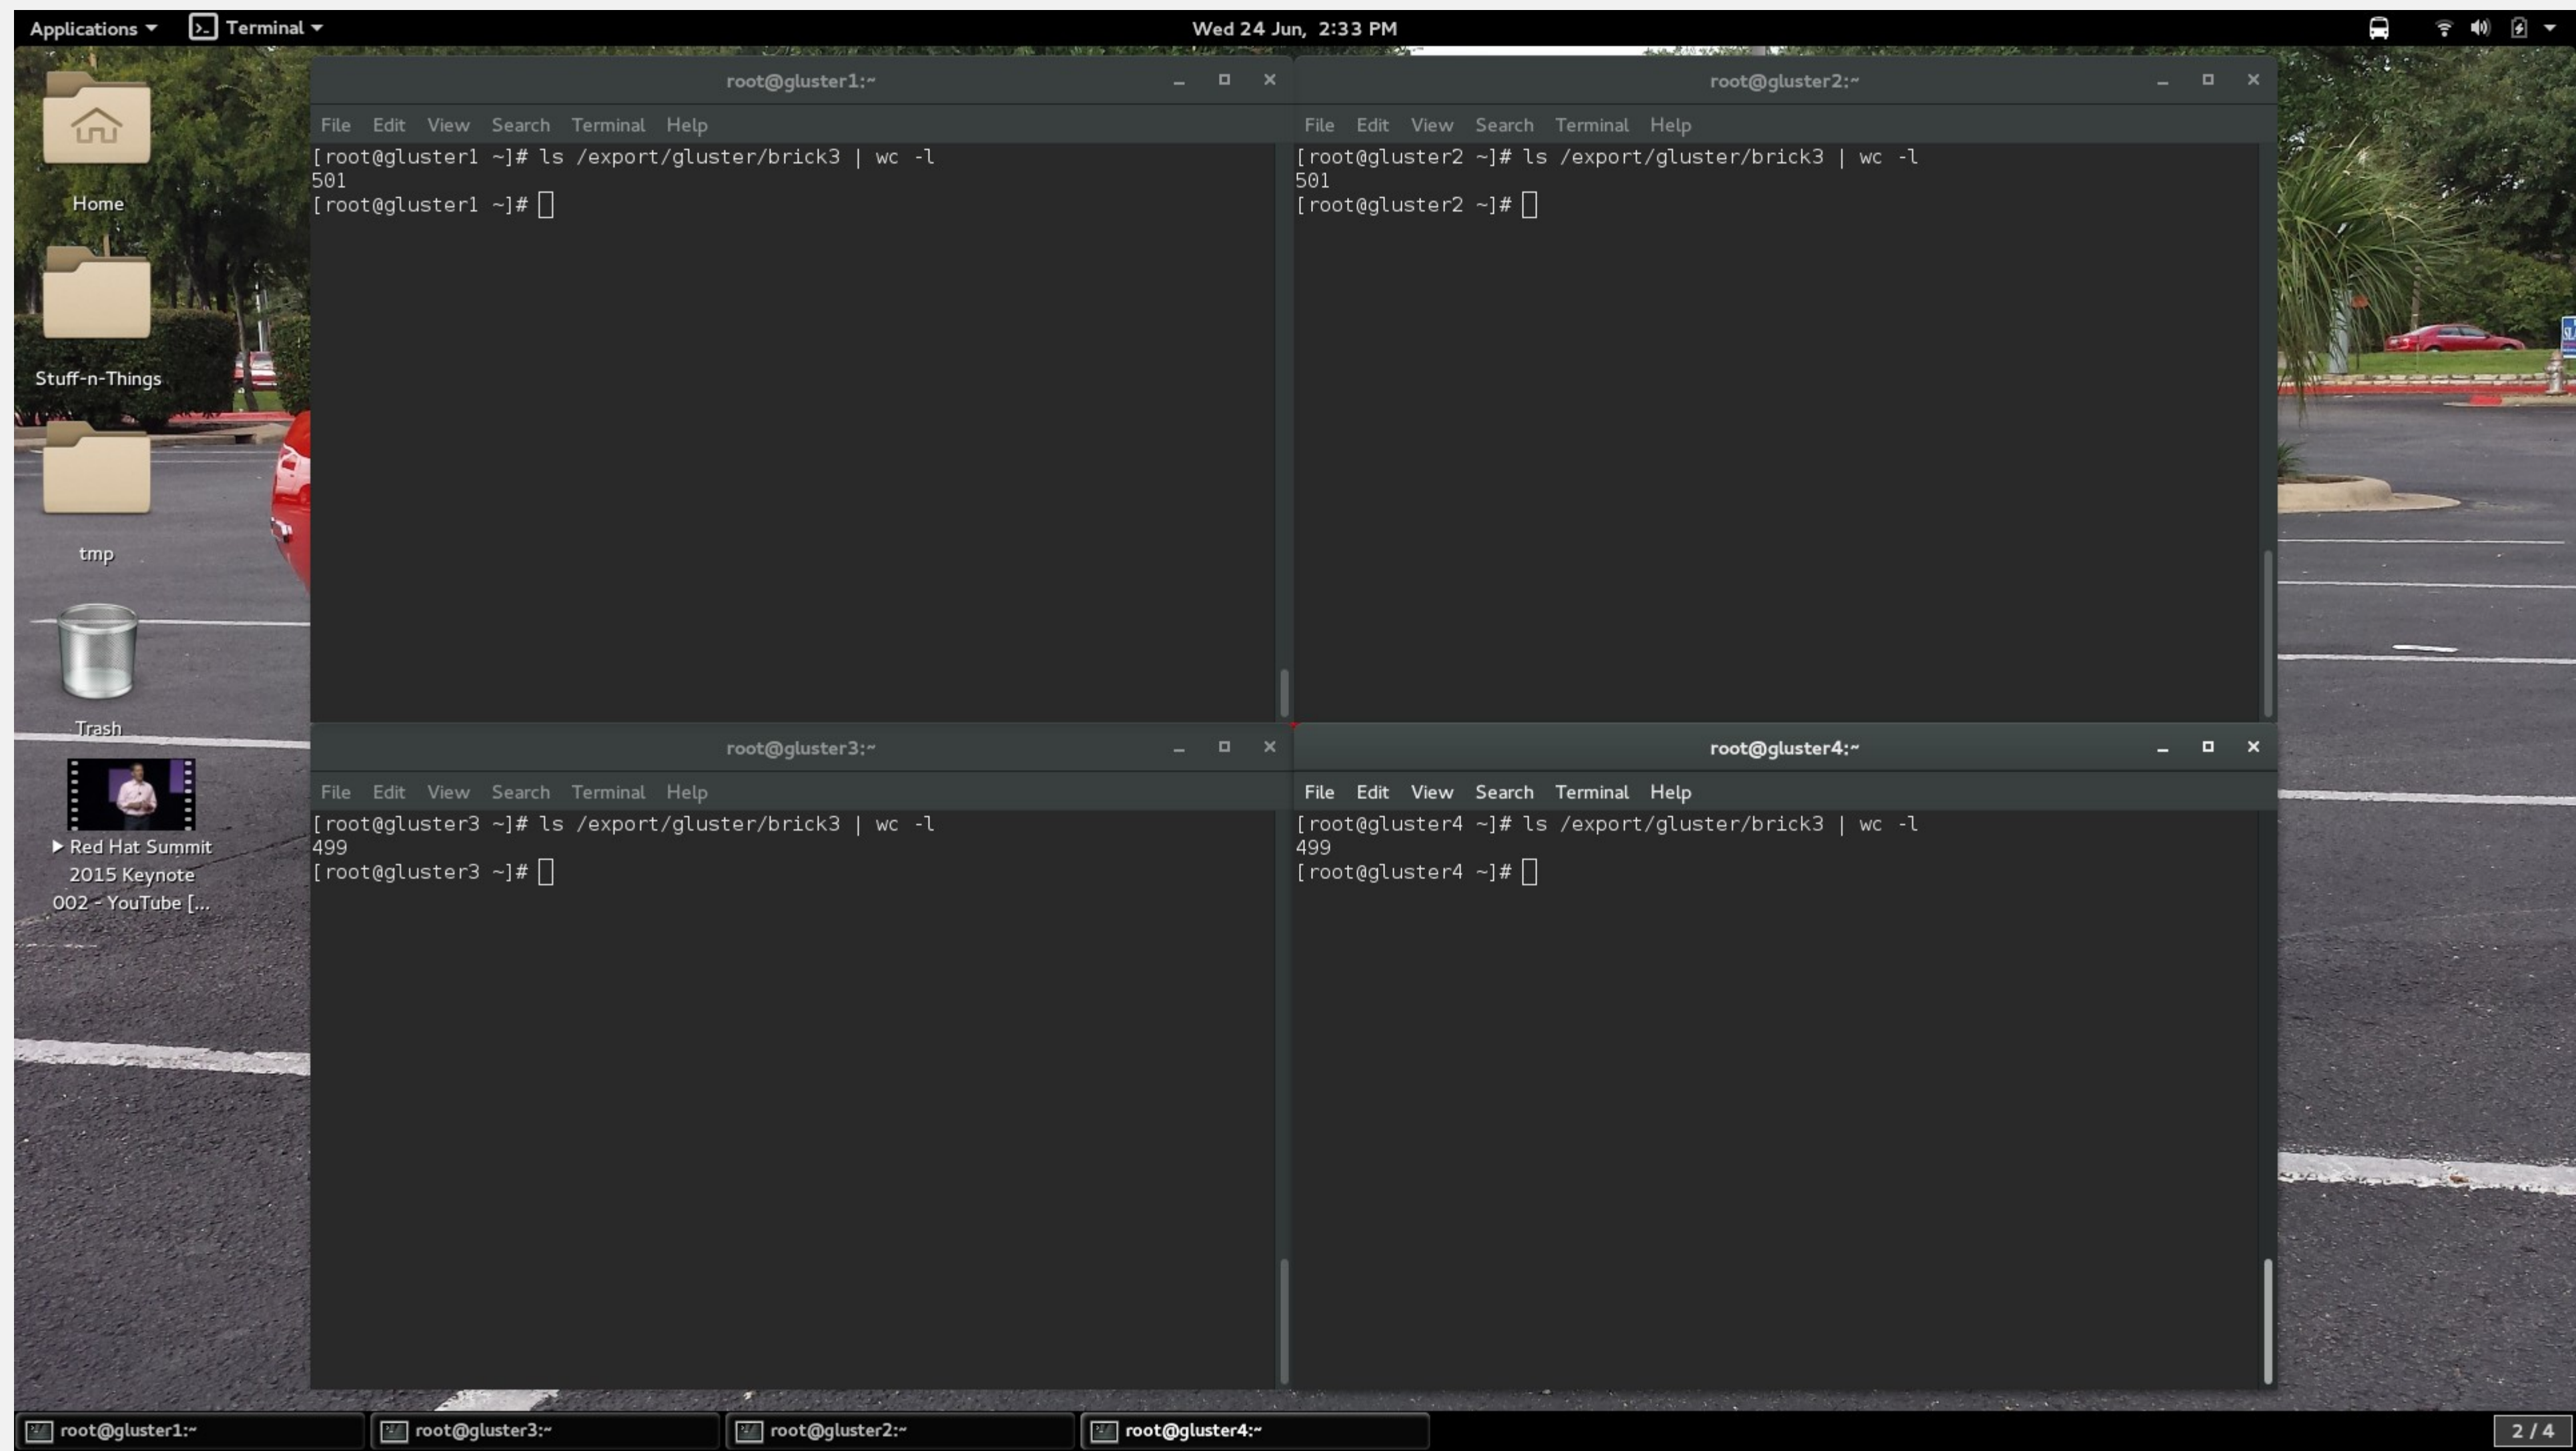

```
root@host189:~  
File Edit View Search Terminal Help  
[root@host189 ~]# mount -t glusterfs gluster1:dist-rep /mnt/gluster/  
[root@host189 ~]# df -h /mnt/gluster/  
Filesystem      Size  Used Avail Use% Mounted on  
gluster1:dist-rep 12G   65M   12G   1% /mnt/gluster  
[root@host189 ~]#
```

ACCESSING DISTRIBUTED+REPLICATED VOLUMES

- This time, since it's a distributed set of replicated (in this case, two way, or mirrored) volumes, creating 1000 files on the clients results in approximately half of the files getting created on each of the distributed nodes, then it is replicated. So 500-ish files on each node.


```
root@host189:/mnt/gluster
File Edit View Search Terminal Help
[root@host189 gluster]# for i in $(seq 1 1000); do echo $i > $i; done
[root@host189 gluster]# ls | wc -l
1000
[root@host189 gluster]#
```

```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# ls /export/gluster/brick3 | wc -l  
501  
[root@gluster1 ~]#
```

```
root@gluster1:~  
File Edit View Search Terminal Help  
[root@gluster1 ~]# gluster volume stop dist-rep  
Stopping volume will make its data inaccessible. Do you want to continue? (y/n)  
y  
volume stop: dist-rep: success  
[root@gluster1 ~]# gluster volume delete dist-rep  
Deleting volume will erase all information about the volume. Do you want to continue? (y/n) y  
volume delete: dist-rep: success  
[root@gluster1 ~]#
```


CONNECTING VIA NFS

CONNECTING VIA NFS

- On the client, install the utilities necessary for NFS


```
root@host189:~  
File Edit View Search Terminal Help  
[root@host189 ~]# yum install nfs-utils
```

```
root@host189:~  
File Edit View Search Terminal Help  
Cleanup      : keyutils-libs-1.4-4.el6.x86_64 9/9  
Verifying    : keyutils-libs-1.4-5.el6.x86_64 1/9  
Verifying    : libevent-1.4.13-4.el6.x86_64 2/9  
Verifying    : rpcbind-0.2.0-11.el6.x86_64 3/9  
Verifying    : keyutils-1.4-5.el6.x86_64 4/9  
Verifying    : 1:nfs-utils-1.2.3-54.el6.x86_64 5/9  
Verifying    : nfs-utils-lib-1.1.5-9.el6_6.x86_64 6/9  
Verifying    : libtirpc-0.2.1-10.el6.x86_64 7/9  
Verifying    : libgssglue-0.1-11.el6.x86_64 8/9  
Verifying    : keyutils-libs-1.4-4.el6.x86_64 9/9  
  
Installed:  
  nfs-utils.x86_64 1:1.2.3-54.el6  
  
Dependency Installed:  
  keyutils.x86_64 0:1.4-5.el6      libevent.x86_64 0:1.4.13-4.el6  
  libgssglue.x86_64 0:0.1-11.el6  libtirpc.x86_64 0:0.2.1-10.el6  
  nfs-utils-lib.x86_64 0:1.1.5-9.el6_6  rpcbind.x86_64 0:0.2.0-11.el6  
  
Dependency Updated:  
  keyutils-libs.x86_64 0:1.4-5.el6  
  
Complete!  
[root@host189 ~]#
```



```
root@host189:~  
File Edit View Search Terminal Help  
[root@host189 ~]# chkconfig rpcbind on  
[root@host189 ~]# service rpcbind start  
Starting rpcbind: [ OK ]  
[root@host189 ~]# chkconfig nfslock on  
[root@host189 ~]# service nfslock start  
Starting NFS statd: [ OK ]  
[root@host189 ~]#
```

```
root@host189:~  
File Edit View Search Terminal Help  
[root@host189 ~]# showmount -e gluster1.tc.redhat.com  
Export list for gluster1.tc.redhat.com:  
/dist-rep *  
[root@host189 ~]#
```



```
root@host189:~  
File Edit View Search Terminal Help  
[root@host189 ~]# mkdir /mnt/nfs && mount -t nfs gluster1.tc.redhat.com:/dist-rep /mnt/nfs  
[root@host189 ~]# df -h /mnt/nfs  
Filesystem                                Size  Used Avail Use% Mounted on  
gluster1.tc.redhat.com:/dist-rep         12G   65M   12G    1% /mnt/nfs  
[root@host189 ~]#
```


QUESTIONS?

THANKS! DO THE SURVEY!

RED HAT **SUMMIT**

LEARN. NETWORK.
EXPERIENCE OPEN SOURCE.