Shared Storage Cluster

2 Knoten-Cluster
mit DRBD und GFS

„Remote-Raid 1“ dank DRBD

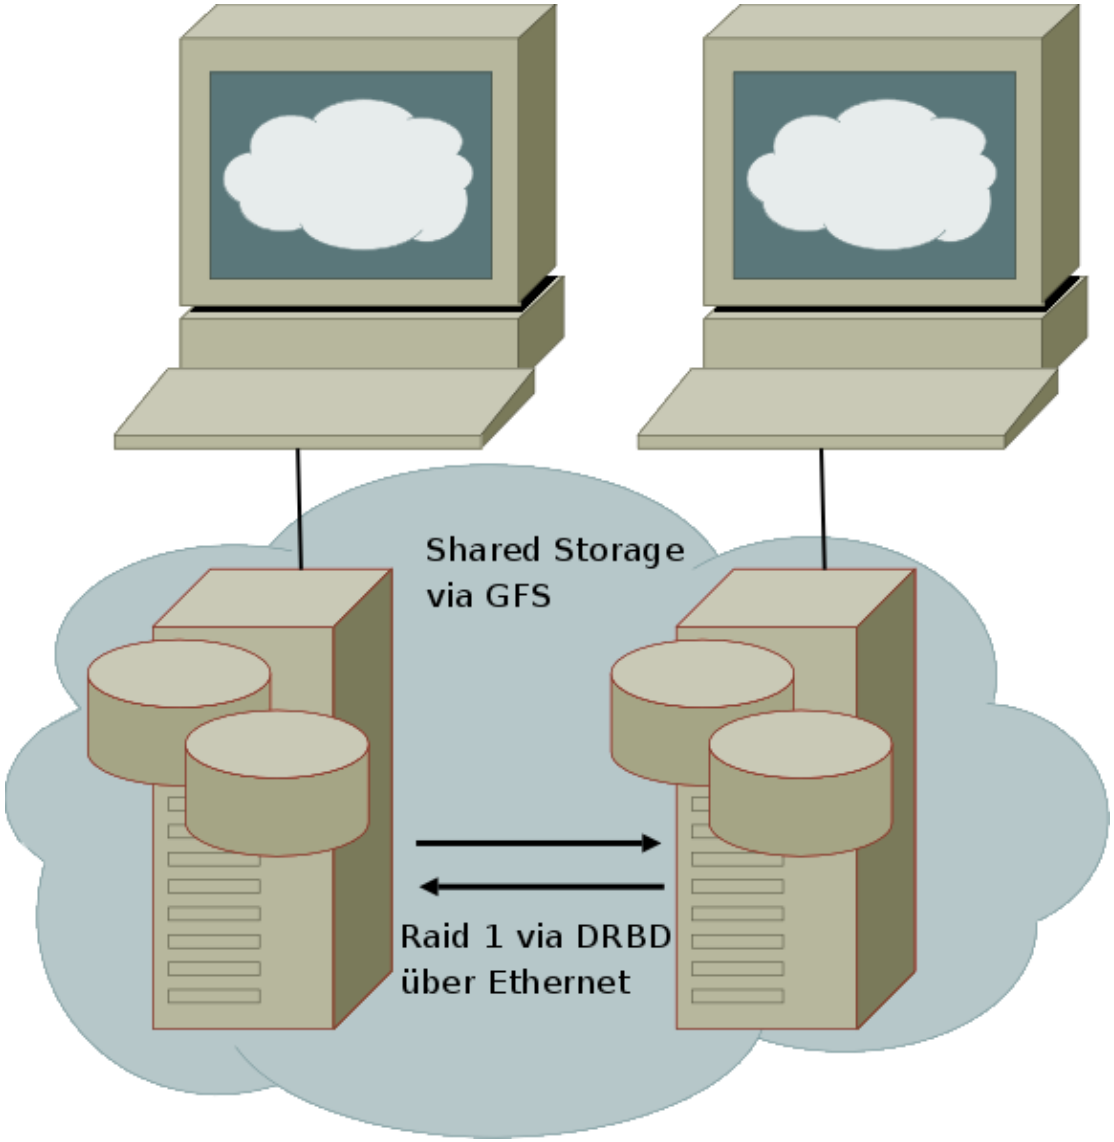

Hardware

- 2 Server
- je 2 Netzwerkinterfaces
- je 1 freie Partition

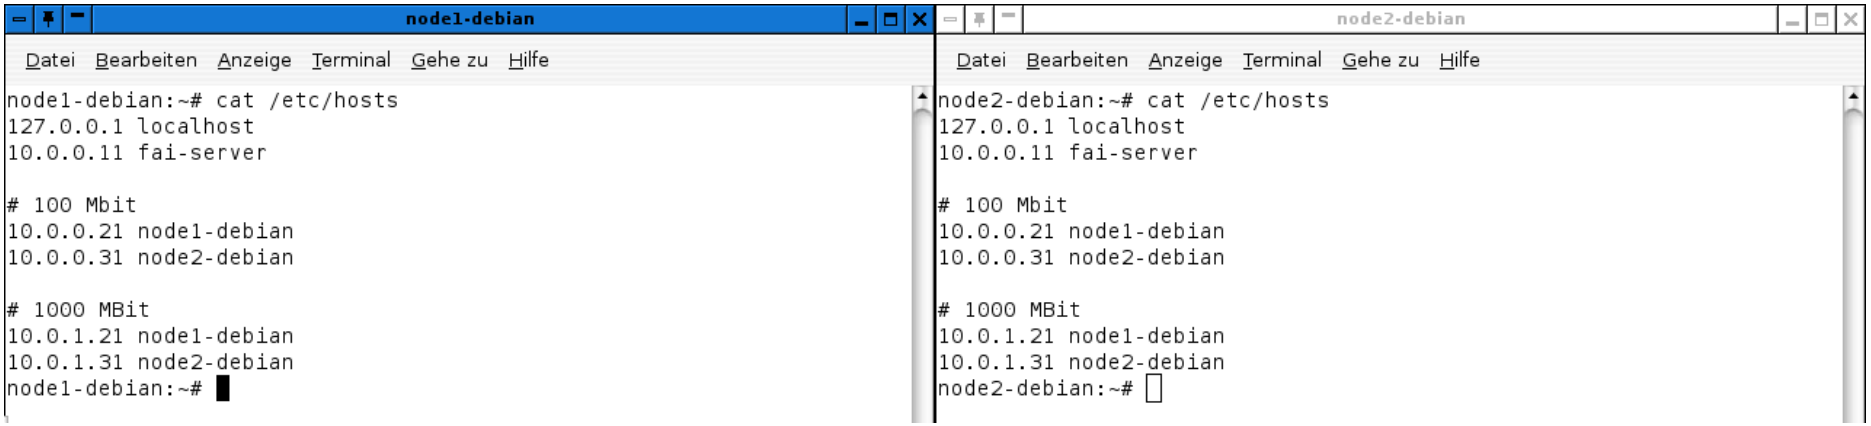

Software

- Linux (z.B. Debian/Etch)
- DRBD (>=8.0 aus backports.org)
- Redhat-Cluster (redhat-cluster-modules + Userspace-Programme)

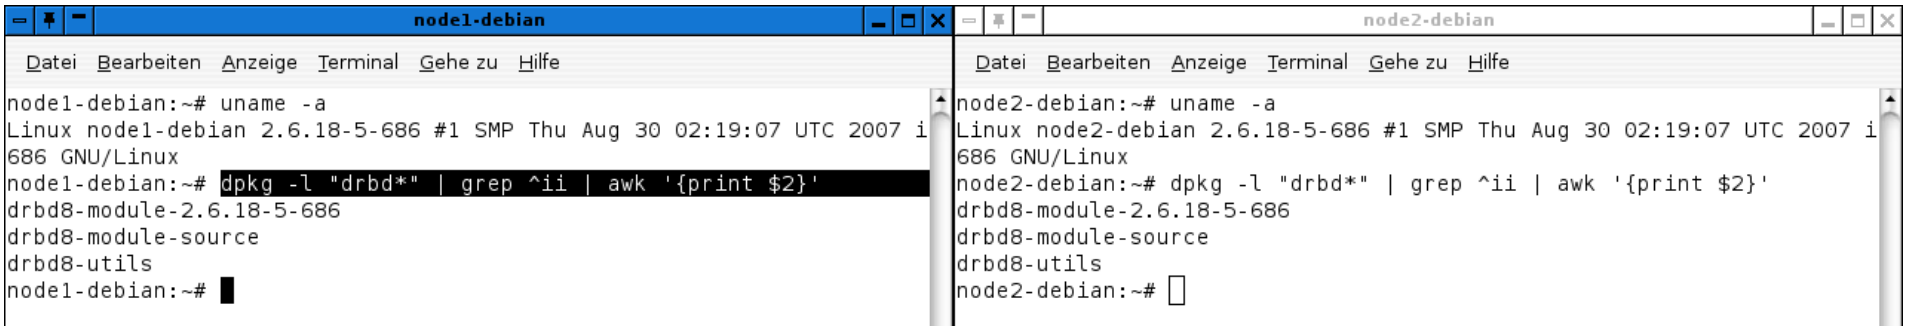

DRBD: simple /etc/drbd.conf

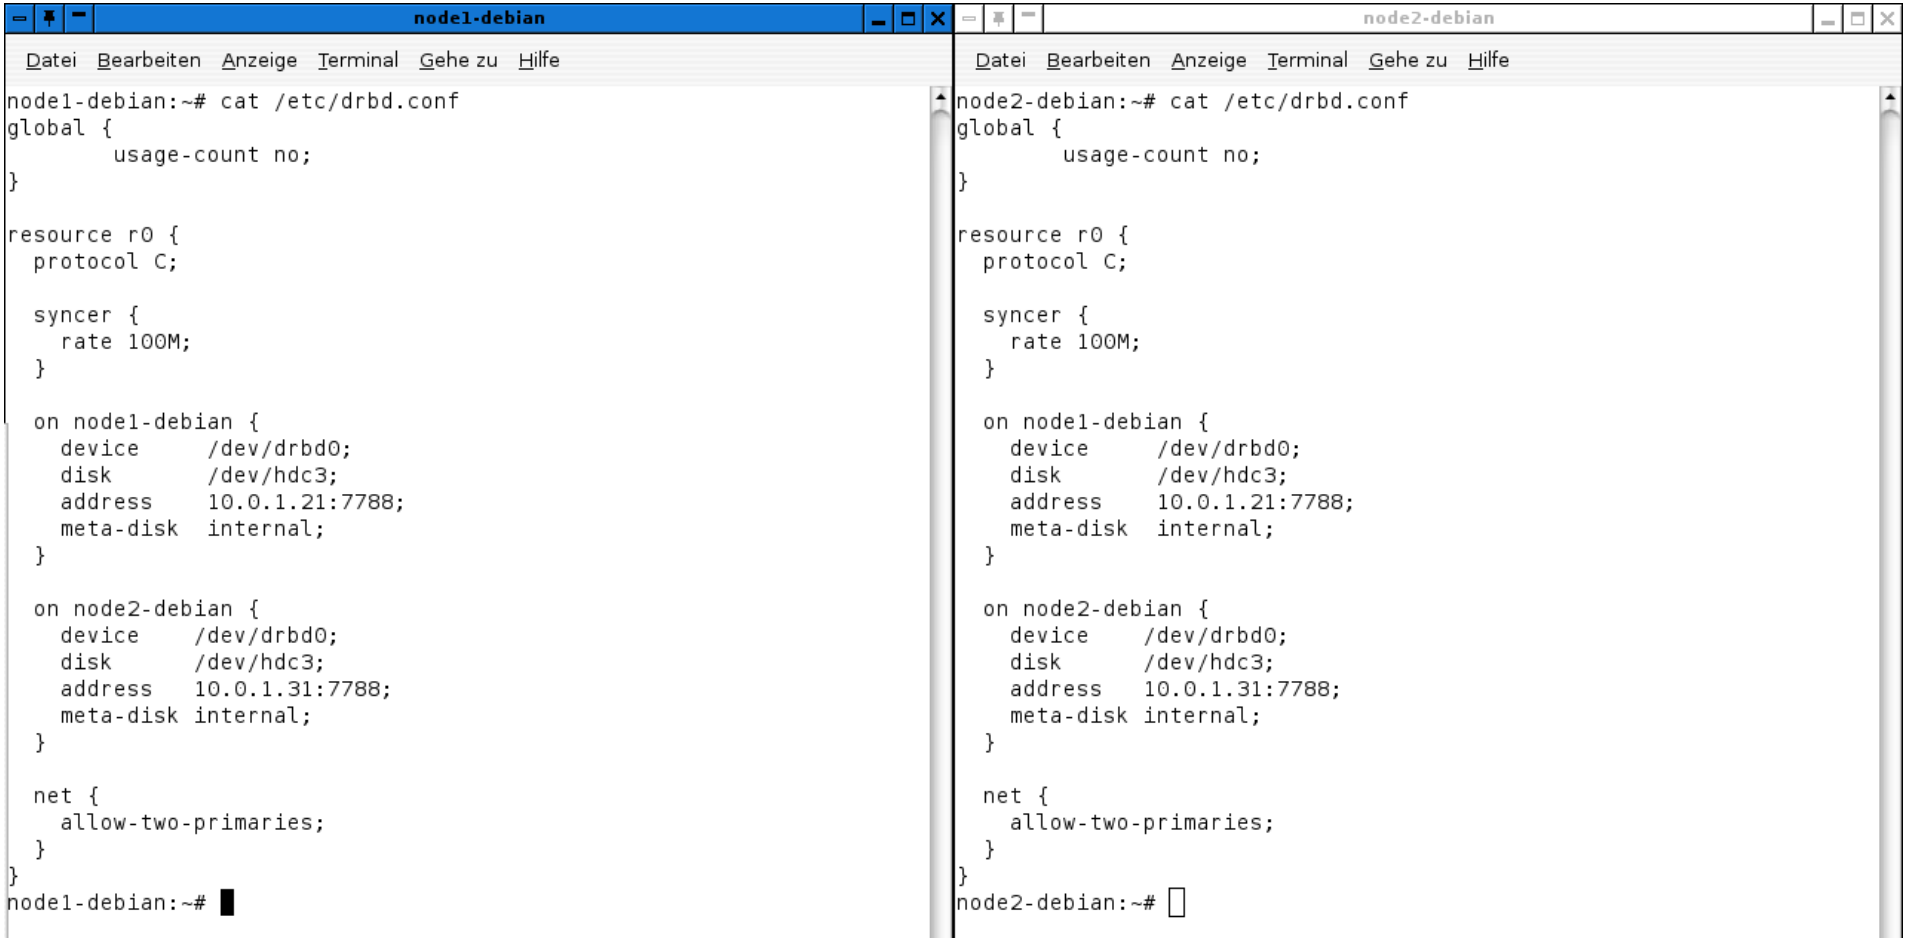

DRBD starten

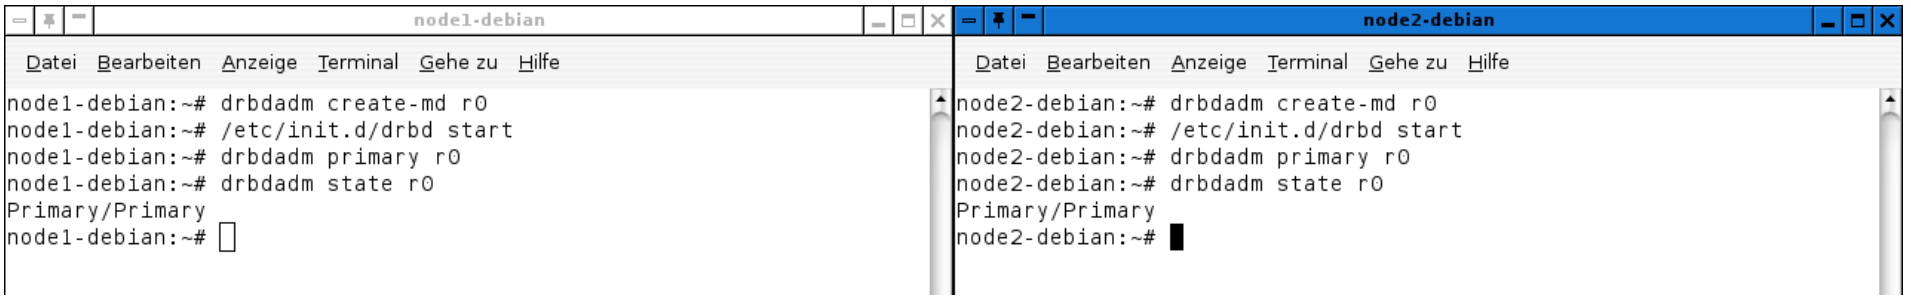

Redhat Cluster/GFS: cluster.conf

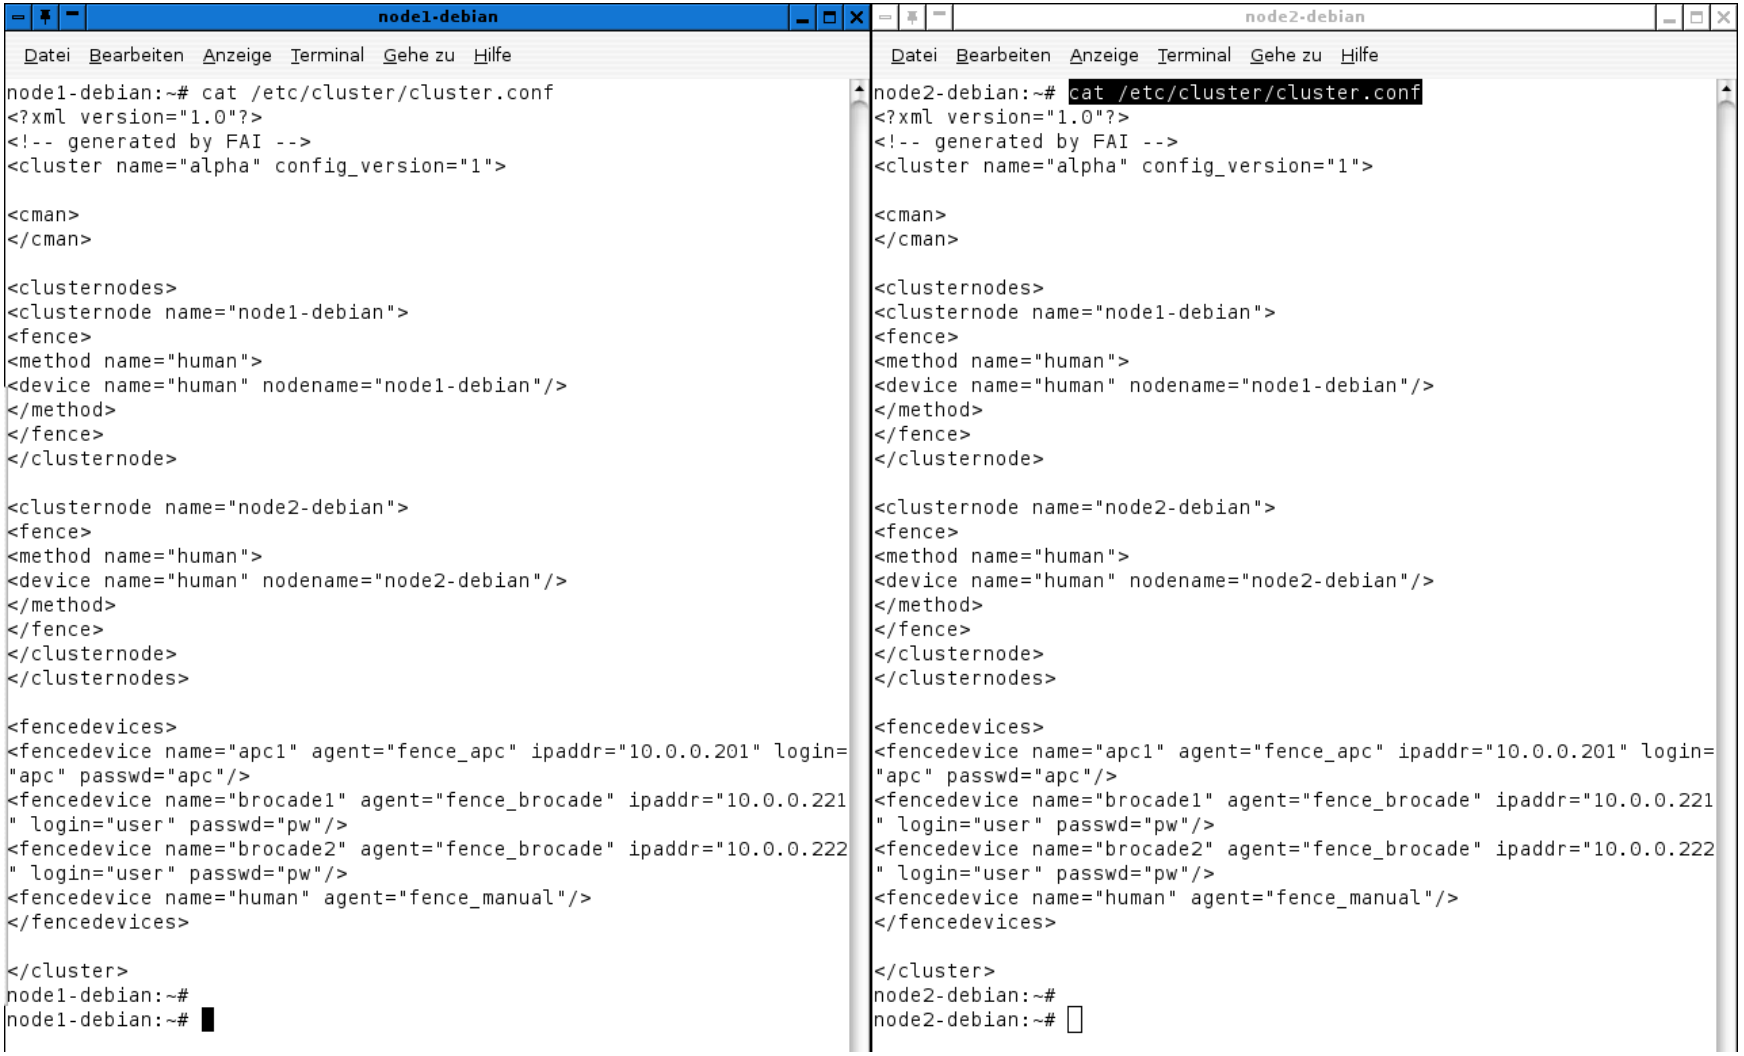

Redhat Cluster starten

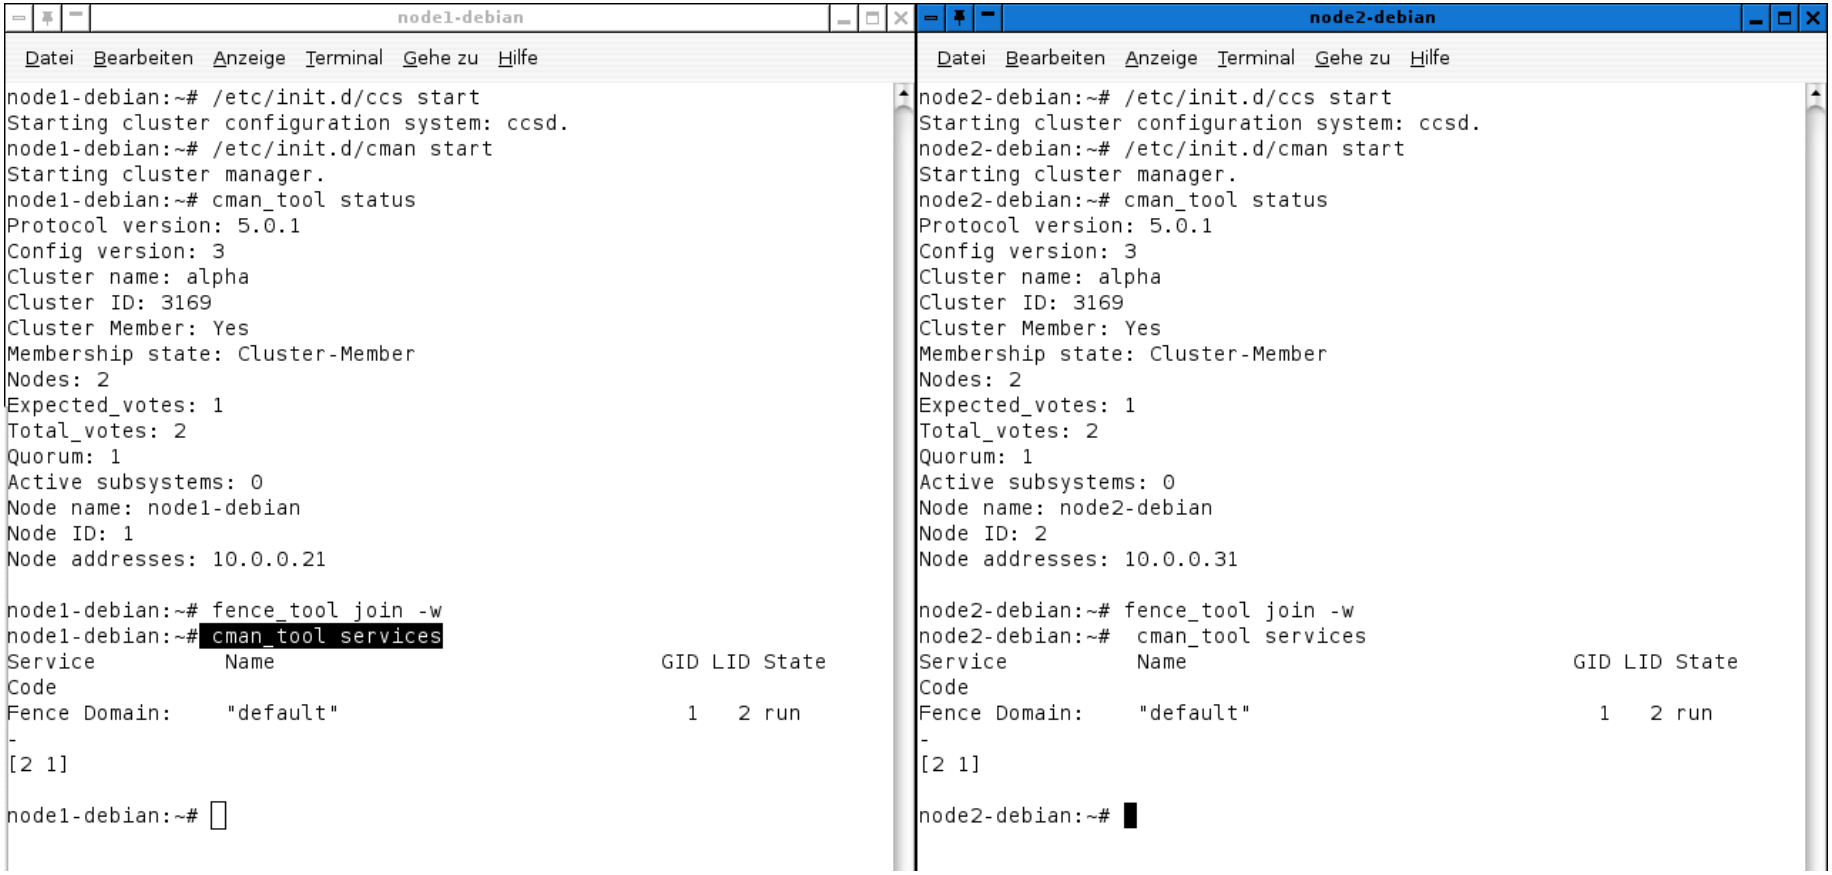

GFS: Dateisystem erstellen

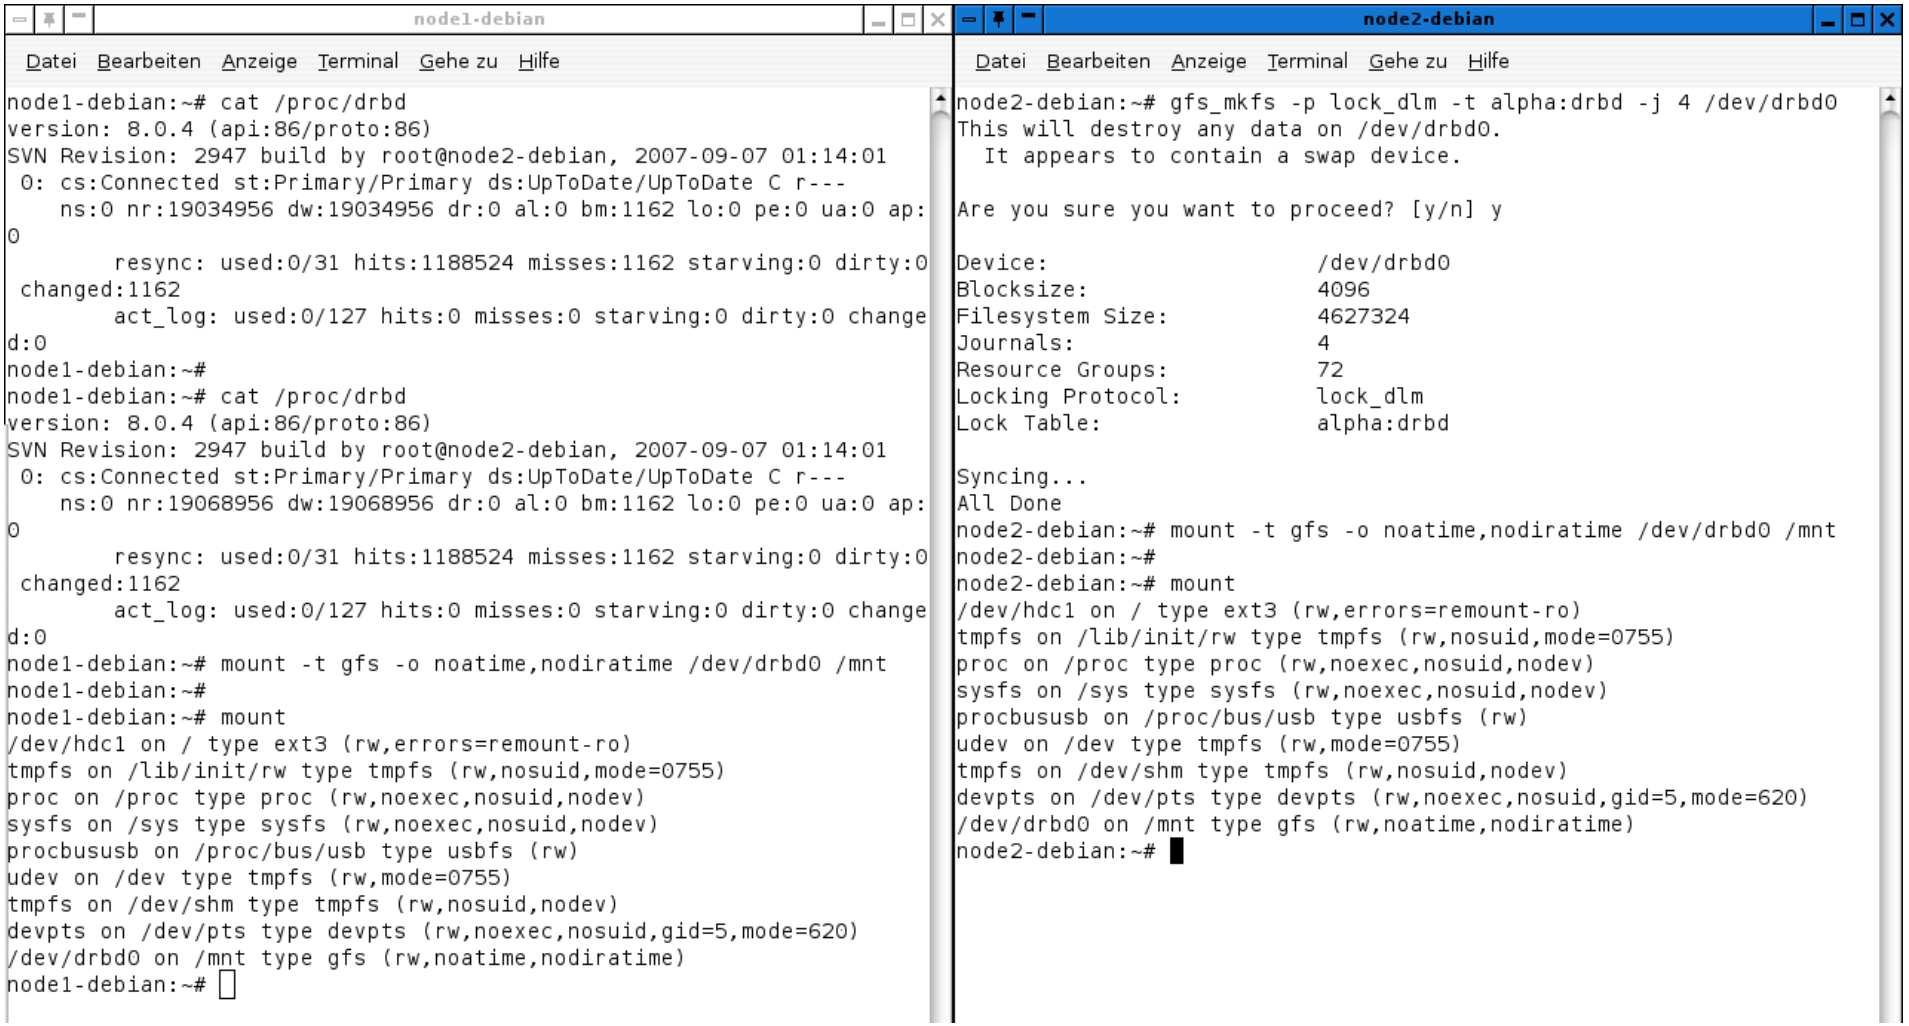

node1-debian

Datei Bearbeiten Anzeige Terminal Gehe zu Hilfe

```
node1-debian:~# cat /proc/drbd
version: 8.0.4 (api:86/proto:86)
SVN Revision: 2947 build by root@node2-debian, 2007-09-07 01:14:01
 0: cs:Connected st:Primary/Primary ds:UpToDate/UpToDate C r---
    ns:0 nr:19034956 dw:19034956 dr:0 al:0 bm:1162 lo:0 pe:0 ua:0 ap:
0
        resync: used:0/31 hits:1188524 misses:1162 starving:0 dirty:0
 changed:1162
        act_log: used:0/127 hits:0 misses:0 starving:0 dirty:0 change
d:0
node1-debian:~#
node1-debian:~# cat /proc/drbd
version: 8.0.4 (api:86/proto:86)
SVN Revision: 2947 build by root@node2-debian, 2007-09-07 01:14:01
 0: cs:Connected st:Primary/Primary ds:UpToDate/UpToDate C r---
    ns:0 nr:19068956 dw:19068956 dr:0 al:0 bm:1162 lo:0 pe:0 ua:0 ap:
0
        resync: used:0/31 hits:1188524 misses:1162 starving:0 dirty:0
 changed:1162
        act_log: used:0/127 hits:0 misses:0 starving:0 dirty:0 change
d:0
node1-debian:~# mount -t gfs -o noatime,nodiratime /dev/drbd0 /mnt
node1-debian:~#
node1-debian:~# mount
/dev/hdc1 on / type ext3 (rw,errors=remount-ro)
tmpfs on /lib/init/rw type tmpfs (rw,nosuid,mode=0755)
proc on /proc type proc (rw,noexec,nosuid,nodev)
sysfs on /sys type sysfs (rw,noexec,nosuid,nodev)
procbususb on /proc/bus/usb type usbfs (rw)
udev on /dev type tmpfs (rw,mode=0755)
tmpfs on /dev/shm type tmpfs (rw,nosuid,nodev)
devpts on /dev/pts type devpts (rw,noexec,nosuid,gid=5,mode=620)
/dev/drbd0 on /mnt type gfs (rw,noatime,nodiratime)
node1-debian:~#
```

node2-debian

Datei Bearbeiten Anzeige Terminal Gehe zu Hilfe

```
node2-debian:~# gfs_mkfs -p lock_dlm -t alpha:drbd -j 4 /dev/drbd0
This will destroy any data on /dev/drbd0.
  It appears to contain a swap device.

Are you sure you want to proceed? [y/n] y

Device:                    /dev/drbd0
Blocksize:                 4096
Filesystem Size:           4627324
Journals:                  4
Resource Groups:           72
Locking Protocol:          lock_dlm
Lock Table:                alpha:drbd

Syncing...
All Done
node2-debian:~# mount -t gfs -o noatime,nodiratime /dev/drbd0 /mnt
node2-debian:~#
node2-debian:~# mount
/dev/hdc1 on / type ext3 (rw,errors=remount-ro)
tmpfs on /lib/init/rw type tmpfs (rw,nosuid,mode=0755)
proc on /proc type proc (rw,noexec,nosuid,nodev)
sysfs on /sys type sysfs (rw,noexec,nosuid,nodev)
procbususb on /proc/bus/usb type usbfs (rw)
udev on /dev type tmpfs (rw,mode=0755)
tmpfs on /dev/shm type tmpfs (rw,nosuid,nodev)
devpts on /dev/pts type devpts (rw,noexec,nosuid,gid=5,mode=620)
/dev/drbd0 on /mnt type gfs (rw,noatime,nodiratime)
node2-debian:~#
```

Clusterdateisystem benutzen

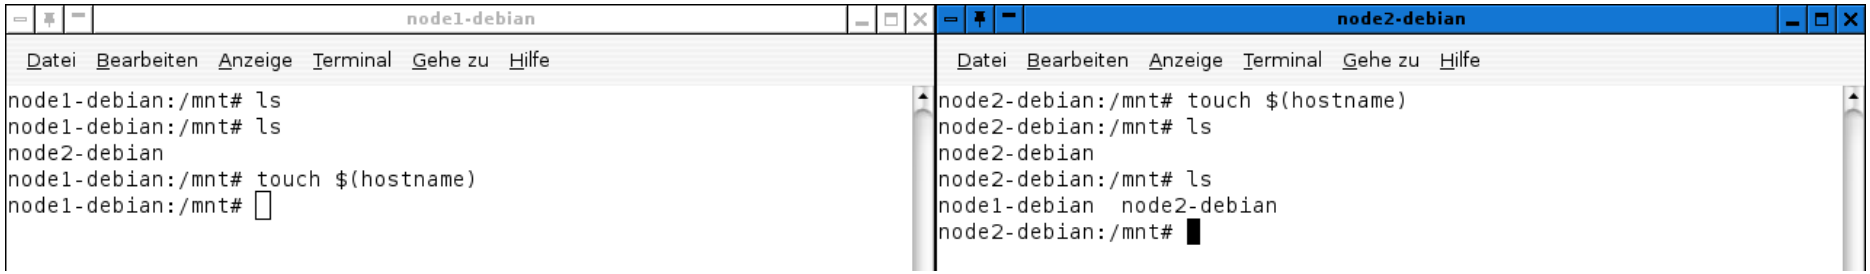

Links

- Debian: http://www.debian.org/
- Redhat-Cluster: http://sources.redhat.com/cluster/
- DRBD: http://www.drbd.org/

Autor

Michael Mende IT Consulting

Herderstraße 29

22085 Hamburg

Tel: 040 / 413 46 412

Email: kontakt@failover.de

Web: http://www.failover.de