

How to Migrate a Single Drive Linux System to RAID-1

Written by Gene Cooper

Friday, 05 December 2008 16:50 - Last Updated Saturday, 06 December 2008 09:57

This howto uses Centos 4.7 but will be nearly identical for any recent Linux distro.

Don't make any mistakes! It's a lot easier to be careful than it is to fix problems.

Reference

http://www.centos.org/modules/newbb/viewtopic.php?topic_id=3501&forum=27

<http://togami.com/~warren/guides/remoteraidcrazies/>

<http://www.debian-administration.org/articles/238>

<http://www.linuxsa.org.au/mailling-list/2003-07/1270.html>

http://www.gnu.org/software/grub/manual/html_node/index.html

Sample Setup

Though this is a very common configuration with SATA (or SAS, SCSI or USB) drives, you may need to modify the following to meet your needs. For example, another common configuration would use */dev/hda* and */dev/hdc* instead of */dev/sda* and */dev/sdb*.

/dev/sda - original drive with system/data

/dev/sdb - new 2nd drive (preferably identical)

/dev/sdc - new hot spare (preferably identical)

Sample Partitioning of the First (Existing) Drive

/dev/sda1 - /boot

/dev/sda2 - /

/dev/sda3 - swap

Make Backup(s)

Back up any way you are comfortable with. You may want to make more than one.

Install the RAID Utilities (if not already installed)

How to Migrate a Single Drive Linux System to RAID-1

Written by Gene Cooper

Friday, 05 December 2008 16:50 - Last Updated Saturday, 06 December 2008 09:57

yum install mdadm

Let's Get Started

From now on, don't reboot until instructed to or your server may not boot up properly.

Disable the Use of Disk Labels and use Partition Device Names Instead

Edit /etc/fstab and grub.conf. Change the labels to device names in both files and add 'panic=10' to the kernel parameters in grub.conf.

vi /etc/fstab

```
/dev/sda2      /          ext3 defaults    1 1
/dev/sda1      /boot      ext3 defaults    1 2
none          /dev/pts   devpts gid=5,mode=620 0 0
none          /dev/shm   tmpfs defaults    0 0
none          /proc      proc defaults    0 0
none          /sys       sysfs defaults    0 0
/dev/sda3      swap       swap defaults    0 0
/dev/hda       /media/cdrom auto pamconsole,exec,noauto,managed 0 0
/dev/scd1      /media/cdrom1 auto pamconsole,exec,noauto,managed 0 0
/dev/scd0      /media/cdrom2 auto pamconsole,exec,noauto,managed 0 0
```

vi /etc/grub.conf

```
default=0
timeout=5
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu
title CentOS (2.6.9-78.0.5.EL)
root (hd0,0)
```

How to Migrate a Single Drive Linux System to RAID-1

Written by Gene Cooper

Friday, 05 December 2008 16:50 - Last Updated Saturday, 06 December 2008 09:57

```
kernel /vmlinuz-2.6.9-78.0.5.EL ro root=/dev/sda2 acpi=off noapic panic=10
initrd /initrd-2.6.9-78.0.5.EL.img
title CentOS-4 i386 (2.6.9-67.0.15.EL)
root (hd0,0)
kernel /vmlinuz-2.6.9-67.0.15.EL ro root=/dev/sda2 acpi=off noapic panic=10
initrd /initrd-2.6.9-67.0.15.EL.img
```

Remove the Labels from the Existing Partitions Themselves (except swap)

```
e2label /dev/sda1 ""
e2label /dev/sda2 ""
```

(/dev/sda3 is swap)

Modify Partitioning on Existing Primary Drive

Change the system types on partitions you want to mirror on the old drive to *fd* (raid autodetect) using *fdisk*. Your finished drive should resemble this output:

```
[root@firefoot root]# sfdisk -l /dev/sda
```

```
Disk /dev/sda: 8942 cylinders, 255 heads, 63 sectors/track
Units = cylinders of 8225280 bytes, blocks of 1024 bytes, counting from 0
```

```
Device Boot Start End #cyls #blocks Id System
/dev/sda1 * 0+ 242 243- 1951866 fd Linux raid autodetect
/dev/sda2 243 485 243 1951897+ fd Linux raid autodetect
/dev/sda3 486 607 122 979965 fd Linux raid autodetect
```

Now use *sfdisk* to duplicate partitions from old drive to new drive(s):

```
sfdisk -d /dev/sda | sfdisk /dev/sdb
```

How to Migrate a Single Drive Linux System to RAID-1

Written by Gene Cooper

Friday, 05 December 2008 16:50 - Last Updated Saturday, 06 December 2008 09:57

```
sfdisk -d /dev/sda | sfdisk /dev/sdc
```

Reboot now.

Create and Prepare the RAID Array

Once rebooted, turn off swap so you can safely create the swap array:

```
swapoff -a
```

Now use mdadm to create the raid arrays. We mark the first drive (sda) as "missing" so the process doesn't wipe out our existing data. Create the raid volumes md0, md1, md2, etc.:

```
# /boot partition
mdadm --create /dev/md0 --level 1 --raid-devices=2 missing /dev/sdb1
# / or root partition
mdadm --create /dev/md1 --level 1 --raid-devices=2 missing /dev/sdb2
# swap partition
mdadm --create /dev/md2 --level 1 --raid-devices=2 missing /dev/sdb3
```

Now that the volumes are ready, initialize the swap and create filesystems for the raid devices:

```
mkfs.ext3 /dev/md0
mkfs.ext3 /dev/md1
mkswap /dev/md2
```

Recreate the initrd image with raid1 support. Adjust for your kernel:

```
cd ~
```

How to Migrate a Single Drive Linux System to RAID-1

Written by Gene Cooper

Friday, 05 December 2008 16:50 - Last Updated Saturday, 06 December 2008 09:57

```
rpm -q kernel # display the installed kernel package
uname -r # display the running kernel
mkinitrd -v --preload=raid1 /root/initrd-2.6.9-78.0.5.EL.img 2.6.9-78.0.5.EL
# the new image should be slightly larger
ll initrd-2.6.9-78.0.5.EL.img /boot/initrd-2.6.9-78.0.5.EL.img
# backup the original
mv /boot/initrd-2.6.9-78.0.5.EL.img /boot/initrd-2.6.9-78.0.5.EL.img.orig
# install the new initrd image
mv initrd-2.6.9-78.0.5.EL.img /boot/
```

Set up the grub bootloader to boot the new initrd image with a fallback to the original in case of a kernel panic. Be sure to change the 'default' line:

```
vi /etc/grub.conf
```

```
default=saved # important
timeout=5
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu
title CentOS (2.6.9-78.0.5.EL)
root (hd0,0)
kernel /vmlinuz-2.6.9-78.0.5.EL ro root=/dev/sda2 acpi=off noapic panic=10
initrd /initrd-2.6.9-78.0.5.EL.img.orig
title CentOS RAID (2.6.9-78.0.5.EL)
root (hd0,0)
kernel /vmlinuz-2.6.9-78.0.5.EL ro root=/dev/sda2 acpi=off noapic panic=10
initrd /initrd-2.6.9-78.0.5.EL.img
```

Configure grub so the second item is the 'saved' default, but only for a single reboot:

```
echo "savedefault --default=1 --once" | grub --batch
```

Double-check everything and reboot.

How to Migrate a Single Drive Linux System to RAID-1

Written by Gene Cooper

Friday, 05 December 2008 16:50 - Last Updated Saturday, 06 December 2008 09:57

If it boots up OK, check the RAID status. Here we see all is well:

```
cat /proc/mdstat
```

```
[root@thirdlane ~]# mdstat
-bash: mdstat: command not found
[root@thirdlane ~]# cat /proc/mdstat
Personalities : [raid1]
md1 : active raid1 sdb2[1]
243312384 blocks [2/1] [_U]

md2 : active raid1 sdb3[1]
779072 blocks [2/1] [_U]

md0 : active raid1 sdb1[1]
104320 blocks [2/1] [_U]
```

Verify the System Will Boot from the New RAID Array

Now that we know it will boot up with the new initrd image, we can modify GRUB so the second (new) stanza points to the second drive (hd1,0) with root on the RAID partition (/dev/md1):

```
title CentOS RAID (2.6.9-78.0.5.EL)
root (hd1,0)
kernel /vmlinuz-2.6.9-78.0.5.EL ro root=/dev/md1 acpi=off noapic panic=10
initrd /initrd-2.6.9-78.0.5.EL.img
```

Copy the Original Disk Information to the New Array

Now mount the new raid volumes and copy the partitions. Be careful to match your md devices with your filesystem layout. This will take some time:

```
mount /dev/md1 /mnt
cp -dpRx / /mnt
```

How to Migrate a Single Drive Linux System to RAID-1

Written by Gene Cooper

Friday, 05 December 2008 16:50 - Last Updated Saturday, 06 December 2008 09:57

```
mount /dev/md0 /mnt/boot
cp -dpRx /boot /mnt
```

Boot from the New RAID Array

Modify the new /etc/fstab to match the RAID device names:

```
vim /mnt/etc/fstab
```

```
# /etc/fstab
#
/dev/md1      /          ext3 defaults    1 1
/dev/md0     /boot     ext3 defaults    1 2
none        /dev/pts  devpts gid=5,mode=620 0 0
none        /dev/shm  tmpfs defaults     0 0
none        /proc     proc  defaults     0 0
none        /sys      sysfs defaults     0 0
/dev/md2     swap      swap  defaults     0 0
/dev/hda     /media/cdrom auto pamconsole,exec,noauto,managed 0 0
/dev/scd1    /media/cdrom1 auto pamconsole,exec,noauto,managed 0 0
/dev/scd0    /media/cdrom2 auto pamconsole,exec,noauto,managed 0 0
```

Once again, configure grub so the second item is the 'saved' default, but only for a single reboot:

```
echo "savedefault --default=1 --once" | grub --batch
```

Double-check everything and reboot.

Now we verify we booted to the RAID array (on the second hard disk) then add your first disk to the array. This will destroy all data on the first (original) disk, as well as the boot sector where grub is installed. Your system will not boot properly if it reboots during this procedure!

How to Migrate a Single Drive Linux System to RAID-1

Written by Gene Cooper

Friday, 05 December 2008 16:50 - Last Updated Saturday, 06 December 2008 09:57

Verify the system booted up using the RAID:

```
[root@pbx ~]# mount
/dev/md1 on / type ext3 (rw)
none on /proc type proc (rw)
none on /sys type sysfs (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
usbfs on /proc/bus/usb type usbfs (rw)
/dev/md0 on /boot type ext3 (rw)
none on /dev/shm type tmpfs (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
```

Show current RAID status:

```
[root@pbx ~]# cat /proc/mdstat
Personalities : [raid1]
md1 : active raid1 sdb2[1]
243312384 blocks [2/1] [_U]

md2 : active raid1 sdb3[1]
779072 blocks [2/1] [_U]

md0 : active raid1 sdb1[1]
104320 blocks [2/1] [_U]
```

At this point, you have all of your original data on the new drive and you have verified it will boot to the new array, so we can safely add the original drive to the raid volume (smallest to largest shown here):

```
mdadm --add /dev/md2 /dev/sda3
mdadm --add /dev/md0 /dev/sda1
mdadm --add /dev/md1 /dev/sda2
```


How to Migrate a Single Drive Linux System to RAID-1

Written by Gene Cooper

Friday, 05 December 2008 16:50 - Last Updated Saturday, 06 December 2008 09:57

Check `/proc/mdstat` for status. When everything is done, all the devices should show `[UU]`.

Don't reboot until it's done syncing the array and it's likely to take a while.

Finalize the Bootloader (GRUB) Configuration

Edit `/boot/grub/grub.conf`. Here we delete the old non-RAID stanza (`#0`), modify the RAID stanza and add an additional entry to boot a recovery mode in case the first drive fails. Don't forget to change the 'default' line. Here is our finalized `grub.conf`:

```
# grub.conf
#
# Boot automatically after 10 secs.
timeout=10
# By default, boot the first entry.
default=0
# Fallback to the second entry.
fallback=1
splashimage=(hd0,0)/grub/splash.xpm.gz
# For booting with disk 0 kernel
title CentOS RAID (2.6.9-78.0.5.EL)
root (hd0,0)
kernel /vmlinuz-2.6.9-78.0.5.EL ro root=/dev/md1 acpi=off noapic panic=10
initrd /initrd-2.6.9-78.0.5.EL.img
# For booting with disk 1 kernel if disk 0 is damaged
title CentOS RAID (RECOVERY) (2.6.9-78.0.5.EL)
root (hd1,0)
kernel /vmlinuz-2.6.9-78.0.5.EL ro root=/dev/md1 acpi=off noapic panic=10
initrd /initrd-2.6.9-78.0.5.EL.img
```

Update the GRUB Bootloader

Here we install the updated GRUB bootloader code into the MBR of all drives in our array. Part of the goal is that the machine will boot regardless of the status of the first drive. There is more than one way to do this, but it's a pain if it doesn't work.

How to Migrate a Single Drive Linux System to RAID-1

Written by Gene Cooper

Friday, 05 December 2008 16:50 - Last Updated Saturday, 06 December 2008 09:57

grub

```
find /grub/stage1 # optional - this will show you where grub setup files are
root (hd0,0)      # use the setup files on the first drive
setup (hd0)      # install grub on the first drive
setup (hd1)      # install grub on the second drive
quit
```

The final test is a reboot, though none is required here. If the machine doesn't boot properly (it might just print 'GRUB' to the console), you can boot from the distribution (Centos, in this case) CD into 'rescue mode' (by entering 'linux rescue' at the CD boot prompt, in this case). Once in rescue mode, you can mount, then chroot into the installation on disk ('chroot /mnt/sysimage' on Centos) to fix GRUB.

Hot Spare

Lets add one more disk to the RAID as a hot spare using the mdadm --add command.

```
sfdisk -d /dev/sda | sfdisk /dev/sdc # if not already done above
mdadm --add /dev/md0 /dev/sdc1
mdadm --add /dev/md1 /dev/sdc2
mdadm --add /dev/md2 /dev/sdc3
```

Verify the new disk partitions are added. Here we see sda is disk 0, sdb is disk 1 and sdc is disk 2 and that the RAID 1 array is fully mirrored ([UU]):

```
[root@thirdlane ~]# cat /proc/mdstat
Personalities : [raid1]
md1 : active raid1 sdc2[2] sdb2[1] sda2[0]
243312384 blocks [2/2] [UU]

md2 : active raid1 sdc3[2] sdb3[1] sda3[0]
779072 blocks [2/2] [UU]

md0 : active raid1 sdc1[2] sdb1[1] sda1[0]
```

How to Migrate a Single Drive Linux System to RAID-1

Written by Gene Cooper

Friday, 05 December 2008 16:50 - Last Updated Saturday, 06 December 2008 09:57

```
104320 blocks [2/2] [UU]
```

Install the bootloader on the hot spare drive:

```
grub
```

```
find /grub/stage1 # this will show you where grub setup files are
root (hd0,0) # use the setup files on the first drive
setup (hd2) # install grub on the second drive
quit
```

Monitor the RAID Array

```
vi /usr/local/bin/linux-raid-status.sh
```

```
#!/bin/bash
#
# linux-raid-status.sh
#
# For Linux Software RAID
#
# Run from cron:
#
# Check RAID status every 10 minutes
#0,10,20,30,40,50 * * * * root /usr/local/bin/linux-raid-status.sh &> /dev/null
```

```
ADMIN="admin@yourdomain.com"
```

```
HOSTNAME=`/bin/hostname`
```

```
if egrep "[._*]" /proc/mdstat > /dev/null
```

```
then
```

```
logger -p daemon.error "mdcheck: Failure of one or more software RAID devices"
```

```
echo "Failure of one or more software RAID devices on ${HOSTNAME}" |
```

```
/bin/mail -s "Software RAID device failure on ${HOSTNAME}" ${ADMIN}
```

How to Migrate a Single Drive Linux System to RAID-1

Written by Gene Cooper

Friday, 05 December 2008 16:50 - Last Updated Saturday, 06 December 2008 09:57

fi

Make the new script executable:

```
chmod +x /usr/local/bin/linux-raid-status.sh
```

Add the crontab entry like this:

```
cat >> /etc/crontab /dev/null  
EOF
```