

add disk/partition to softRaid5

Source: <http://linux.derkeiler.com/Newsgroups/comp.os.linux.misc/2005-05/2178.html>

From: Carsten Menke (*carsten_at_menke-net.net*)

Date: 05/21/05

Date: Sat, 21 May 2005 10:15:27 +0200

Hi there,

I am trying to add a partition to a 3 disk RAID5 device by using raidreconf and after the procedure which I describe below the filesystem is in a complete mess and e2fsck crashes with a segmentation fault and up to now I do not figure out what I am doing wrong.

Currently, I am only trying out if it works before I switch over to a production system.

Basically, I do the following after creating the device and copying some garbage to the device:

```
umount /mnt/raid
fsck -f /dev/md0
raidstop /dev/md0
cp -f /etc/raidtab /etc/raidtab.old
cp -f /etc/raidtab.new /etc/raidtab
raidreconf -o /etc/raidtab.old -n /etc/raidtab -m /dev/md0
wait until the raid is synchronized
```

Please let me know if someone needs further info.

Thanks in advance.

Cheers,

Carsten

RedHat 9.0
kernel-2.6.11.7
e2fsprogs-1.37-4

```
raidtab.old
raiddev /dev/md0
raid-level 5
nr-raid-disks 3
nr-spare-disks 3
```

comp.os.linux.misc: add disk/partition to softRaid5

```
persistent-superblock 1
parity-algorithm left-symmetric
chunk-size 64
device /dev/sdc5
raid-disk 0
device /dev/sdc6
raid-disk 1
device /dev/sdc7
raid-disk 2
device /dev/sdc8
spare-disk 0
device /dev/sdc9
spare-disk 1
device /dev/sdc10
spare-disk 2
```

raidtab.new

```
raiddev /dev/md1
raid-level 5
nr-raid-disks 4
nr-spare-disks 2
persistent-superblock 1
parity-algorithm left-symmetric
chunk-size 64
device /dev/sdc5
raid-disk 0
device /dev/sdc6
raid-disk 1
device /dev/sdc7
raid-disk 2
device /dev/sdc8
raid-disk 3
device /dev/sdc9
spare-disk 0
device /dev/sdc10
spare-disk 1
```

```
#!/bin/sh
umount /mnt/raid
raidstop /dev/md0
cp -f /etc/raidtab1 /etc/raidtab
fdisk /dev/sdc <myraid.fdisk
mkraid -R /dev/md0
raidstart /dev/md0
sleep 5
cat /proc/mdstat
while grep = /proc/mdstat ; do
sleep 2
done
cat /proc/mdstat
mkfs -t ext3 /dev/md0
```

add disk/partition to softRaid5

```
mount /dev/md0 /mnt/raid
cp /var/tmp/dump/* /mnt/raid
umount /mnt/raid
fsck -f /dev/md0
raidstop /dev/md0
cp -f /etc/raidtab /etc/raidtab.old
cp -f /etc/raidtab.new /etc/raidtab
raidreconf -o /etc/raidtab.old -n /etc/raidtab -m /dev/md0
sleep 2
cat /proc/mdstat
while grep = /proc/mdstat ; do
sleep 2
done
cat /proc/mdstat
```

output:

```
./myraid.sh
```

```
umount: /mnt/raid: not mounted
```

The number of cylinders for this disk is set to 30401.

There is nothing wrong with that, but this is larger than 1024, and could in certain setups cause problems with:

- 1) software that runs at boot time (e.g., old versions of LILO)
- 2) booting and partitioning software from other OSs (e.g., DOS FDISK, OS/2 FDISK)

Command (m for help): Partition number (1-10):

Command (m for help): Partition number (1-9):

Command (m for help): Partition number (1-8):

Command (m for help): Partition number (1-7):

Command (m for help): Partition number (1-6):

Command (m for help): Partition number (1-5):

Command (m for help): Command action

l logical (5 or over)

p primary partition (1-4)

First cylinder (18-119, default 18): Using default value 18

Last cylinder or +size or +sizeM or +sizeK (18-119, default 119):

Command (m for help): Command action

l logical (5 or over)

p primary partition (1-4)

First cylinder (35-119, default 35): Using default value 35

Last cylinder or +size or +sizeM or +sizeK (35-119, default 119):

Command (m for help): Command action

l logical (5 or over)

p primary partition (1-4)

First cylinder (52-119, default 52): Using default value 52

Last cylinder or +size or +sizeM or +sizeK (52-119, default 119):

Command (m for help): Command action

l logical (5 or over)

p primary partition (1-4)

First cylinder (69-119, default 69): Using default value 69

comp.os.linux.misc: add disk/partition to softRaid5

Last cylinder or +size or +sizeM or +sizeK (69–119, default 119):

Command (m for help): Command action

l logical (5 or over)

p primary partition (1–4)

First cylinder (86–119, default 86): Using default value 86

Last cylinder or +size or +sizeM or +sizeK (86–119, default 119):

Command (m for help): Command action

l logical (5 or over)

p primary partition (1–4)

First cylinder (103–119, default 103): Using default value 103

Last cylinder or +size or +sizeM or +sizeK (103–119, default 119):

Command (m for help): The partition table has been altered!

Calling ioctl() to re-read partition table.

Syncing disks.

DESTROYING the contents of /dev/md0 in 5 seconds, Ctrl-C if unsure!

handling MD device /dev/md0

analyzing super-block

disk 0: /dev/sdc5, 136521kB, raid superblock at 136448kB

disk 1: /dev/sdc6, 136521kB, raid superblock at 136448kB

disk 2: /dev/sdc7, 136521kB, raid superblock at 136448kB

disk 3: /dev/sdc8, 136521kB, raid superblock at 136448kB

disk 4: /dev/sdc9, 136521kB, raid superblock at 136448kB

disk 5: /dev/sdc10, 136521kB, raid superblock at 136448kB

/dev/md0: Invalid argument

Personalities : [raid5]

md0 : active raid5 sdc10[3] sdc9[4] sdc8[5] sdc7[2] sdc6[1] sdc5[0]

272896 blocks level 5, 64k chunk, algorithm 2 [3/3] [UUU]

[=====>.....] resync = 29.8% (41728/136448)

finish=0.1min speed=8345K/sec

unused devices: <none>

[=====>.....] resync = 29.8% (41856/136448)

finish=0.1min speed=8371K/sec

[=====>.....] resync = 41.7% (58280/136448)

finish=0.1min speed=8325K/sec

[=====>.....] resync = 53.7% (74408/136448)

finish=0.1min speed=8267K/sec

[=====>.....] resync = 64.9% (90084/136448)

finish=0.0min speed=8189K/sec

[=====>.....] resync = 76.8% (105608/136448)

finish=0.0min speed=8123K/sec

[=====>.....] resync = 89.5% (123644/136448)

finish=0.0min speed=8242K/sec

Personalities : [raid5]

md0 : active raid5 sdc10[3] sdc9[4] sdc8[5] sdc7[2] sdc6[1] sdc5[0]

272896 blocks level 5, 64k chunk, algorithm 2 [3/3] [UUU]

unused devices: <none>

mke2fs 1.37 (21-Mar-2005)

Filesystem label=

OS type: Linux

add disk/partition to softRaid5

comp.os.linux.misc: add disk/partition to softRaid5

Block size=1024 (log=0)
Fragment size=1024 (log=0)
68272 inodes, 272896 blocks
13644 blocks (5.00%) reserved for the super user
First data block=1
Maximum filesystem blocks=67633152
34 block groups
8192 blocks per group, 8192 fragments per group
2008 inodes per group
Superblock backups stored on blocks:
8193, 24577, 40961, 57345, 73729, 204801, 221185

Writing inode tables: done
Creating journal (8192 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 23 mounts or 180 days, whichever comes first. Use tune2fs -c or -i to override.
cp: omitting directory `/root/linux-2.6.11.7'
cp: omitting directory `/root/Mail'
fsck 1.37 (21-Mar-2005)
e2fsck 1.37 (21-Mar-2005)
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
/dev/md0: 41/68272 files (4.9% non-contiguous), 20085/272896 blocks
Working with device /dev/md0
Parsing /etc/raidtab.old
Parsing /etc/raidtab.new
Size of old array: 1638252 blocks, Size of new array: 1638252 blocks
Old raid-disk 0 has 2132 chunks, 136448 blocks
Old raid-disk 1 has 2132 chunks, 136448 blocks
Old raid-disk 2 has 2132 chunks, 136448 blocks
Old raid-disk 3 has 2132 chunks, 136448 blocks
Old raid-disk 4 has 2132 chunks, 136448 blocks
Old raid-disk 5 has 2132 chunks, 136448 blocks
New raid-disk 0 has 2132 chunks, 136448 blocks
New raid-disk 1 has 2132 chunks, 136448 blocks
New raid-disk 2 has 2132 chunks, 136448 blocks
New raid-disk 3 has 2132 chunks, 136448 blocks
New raid-disk 4 has 2132 chunks, 136448 blocks
New raid-disk 5 has 2132 chunks, 136448 blocks
Using 64 Kbyte blocks to move from 64 Kbyte chunks to 64 Kbyte chunks.
Detected 1036464 KB of physical memory in system
A maximum of 3454 outstanding requests is allowed

I will convert your old device /dev/md0 of 10660 blocks
to a new device /dev/md0 of same size
using a block-size of 64 KB

add disk/partition to softRaid5

comp.os.linux.misc: add disk/partition to softRaid5

```
Is this what you want? (yes/no): yes
Converting 10660 block device to 10660 block device
Allocated free block map for 6 disks
6 unique disks detected.
Working (/) [00010660/00010660]
#####]
Source drained, flushing sink.
Reconfiguration succeeded, will update superblocks...
Updating superblocks...
handling MD device /dev/md0
analyzing super-block
disk 0: /dev/sdc5, 136521kB, raid superblock at 136448kB
disk 1: /dev/sdc6, 136521kB, raid superblock at 136448kB
disk 2: /dev/sdc7, 136521kB, raid superblock at 136448kB
disk 3: /dev/sdc8, 136521kB, raid superblock at 136448kB
disk 4: /dev/sdc9, 136521kB, raid superblock at 136448kB
disk 5: /dev/sdc10, 136521kB, raid superblock at 136448kB
Array is updated with kernel.
Disks re-inserted in array... Hold on while starting the array...
Maximum friend-freeing depth: 1
Total wishes hooked: 10660
Maximum wishes hooked: 3454
Total gifts hooked: 10660
Maximum gifts hooked: 2110
Congratulations, your array has been reconfigured,
and no errors seem to have occurred.
Personalities : [raid5]
md0 : active raid5 sdc10[4] sdc9[5] sdc8[3] sdc7[2] sdc6[1] sdc5[0]
      409344 blocks level 5, 64k chunk, algorithm 2 [4/4] [UUUU]
      [==>.....] resync = 10.4% (14784/136448)
finish=0.2min speed=7392K/sec
unused devices: <none>
      [==>.....] resync = 11.1% (15628/136448)
finish=0.2min speed=7814K/sec
      [====>.....] resync = 20.8% (28876/136448)
finish=0.2min speed=7219K/sec
      [=====>.....] resync = 29.8% (41136/136448)
finish=0.2min speed=6856K/sec
      [=====>.....] resync = 38.8% (53488/136448)
finish=0.2min speed=6686K/sec
      [=====>.....] resync = 47.7% (65784/136448)
finish=0.1min speed=6578K/sec
      [=====>.....] resync = 56.7% (78552/136448)
finish=0.1min speed=6546K/sec
      [=====>.....] resync = 65.6% (91124/136448)
finish=0.1min speed=6508K/sec
      [=====>.....] resync = 75.3% (103764/136448)
finish=0.0min speed=6485K/sec
      [=====>.....] resync = 85.8% (118448/136448)
finish=0.0min speed=6580K/sec
      [=====>.....] resync = 94.7% (130816/136448)
```

add disk/partition to softRaid5

comp.os.linux.misc: add disk/partition to softRaid5

finish=0.0min speed=6540K/sec

Personalities : [raid5]

md0 : active raid5 sdc10[4] sdc9[5] sdc8[3] sdc7[2] sdc6[1] sdc5[0]
409344 blocks level 5, 64k chunk, algorithm 2 [4/4] [UUUU]

unused devices: <none>