By Till Brehm
Published: 2007-04-22 22:44

Recover Data From RAID1 LVM Partitions With Knoppix Linux LiveCD

Version 1.0
Author: Till Brehm <t.brehm [at] projektfarm [dot] com>
Lastedited: 04/11/2007

This tutorial describes how to rescue data from a single hard disk that was part of a LVM2 RAID1 setup like it is created by e.g the Fedora Core installer. Why is it so problematic to recover the data? Every single hard disk that formerly was a part of a LVM RAID1 setup contains all data that was stored in the RAID, but the hard disk cannot simply be mounted. First, a RAID setup must be configured for the partition(s) and then LVM must be set up to use this (these) RAID partition(s) before you will be able to mount it. I will use the Knoppix Linux LiveCD to do the data recovery.

Prerequisites

I used a **Knoppix** 5.1 LiveCD for this tutorial. Download the CD ISO image from **here** and burn it on CD, then connect the hard disk which contains the RAID partition(s) to the IDE / ATA controller of your mainboard, put the Knoppix CD in your CD drive and boot from the CD.

The hard disk I used is an IDE drive that is attached to the first IDE controller (hda). In my case, the hard disk contained only one partition.

Restoring The Raid

After Knoppix has booted, open a shell and execute the command:

```
sudo su
```

to become the root user.

As I don't have the *mdadm.conf* file from the original configuration, I create it with this command:

```
mdadm --examine --scan /dev/hda1 >> /etc/mdadm/mdadm.conf
```

The result should be similar to this one:

```
DEVICE partitions
    CREATE owner=root group=disk mode=0660 auto=yes metadata=1
    MAILADDR root
ARRAY /dev/md0 level=raid1 num-devices=2 UUID=a28090aa:6893be8b:c4024dfc:29cdb07a
```

Edit the file and add *devices=/dev/hda1,missing* at the end of the line that describes the RAID array.

```
vi /etc/mdadm/mdadm.conf
```

Finally the file looks like this:

```
DEVICE partitions
    CREATE owner=root group=disk mode=0660 auto=yes metadata=1
    MAILADDR root
ARRAY /dev/md0 level=raid1 num-devices=2 UUID=a28090aa:6893be8b:c4024dfc:29cdb07a devices=/dev/hda1,missing
```

The string */dev/hda1* is the hardware device and *missing* means that the second disk in this RAID array is not present at the moment.

Edit the file */etc/default/mdadm*:

```
vi /etc/default/mdadm
```

and change the line:

```
AUTOSTART=false
```

to:

```
AUTOSTART=true
```

Now we can start our RAID setup:

```
/etc/init.d/mdadm start

/etc/init.d/mdadm-raid start
```

To check if our RAID device is ok, run the command:

```
cat /proc/mdstat
```

The output should look like this:

```
Personalities : [linear] [multipath] [raid0] [raid1] [raid6] [raid5] [raid4] [ra
                                                    id10]
  md0 : active raid1 hda1[1]
  293049600 blocks [2/1] [_U]

unused devices: <none>
```

Recovering The LVM Setup

The LVM configuration file cannot be created by an easy command like the `mdadm.conf`, but LVM stores one or more copy(s) of the configuration file content at the beginning of the partition. I use the command `dd` to extract the first part of the partition and write it to a text file:

```
dd if=/dev/md0 bs=512 count=255 skip=1 of=/tmp/md0.txt
```

Open the file with a text editor:

```
vi /tmp/md0.txt
```

You will find some binary data first and then a configuration file part like this:

```
VolGroup00 {
 id = "evRkPK-aCjV-HiHY-oaaD-SwUO-zN7A-LyRhoj"
 seqno = 2
 status = ["RESIZEABLE", "READ", "WRITE"]
 extent_size = 65536  # 32 Megabytes
 max_lv = 0
 max_pv = 0

 physical_volumes {

  pv0 {
   id = "uMJ8uM-sfTJ-La9j-oIuy-W3NX-ObiT-n464Rv"
   device = "/dev/md0" # Hint only

   status = ["ALLOCATABLE"]
   pe_start = 384
   pe_count = 8943 # 279,469 Gigabytes
  }
 }

 logical_volumes {

  LogVol00 {
   id = "ohesOX-VRSi-CsnK-PUoI-GjUE-0nT7-ltxWoy"
   status = ["READ", "WRITE", "VISIBLE"]
   segment_count = 1

   segment1 {
    start_extent = 0
    extent_count = 8942 # 279,438 Gigabytes

    type = "striped"
    stripe_count = 1 # linear
```

```
    stripes = [
     "pv0", 0
    ]
   }
  }
 }
}
```

Create the file */etc/lvm/backup/VolGroup00*:

```
vi /etc/lvm/backup/VolGroup00
```

and insert the configuration data so the file looks similar to the above example.

Now we can start LVM:

```
/etc/init.d/lvm start
```

Read in the volume:

```
vgscan
```

```
  Reading all physical volumes.  This may take a while...
    Found volume group "VolGroup00" using metadata type lvm2
```

```
pvscan
```

```
 PV /dev/md0   VG VolGroup00   lvm2 [279,47 GB / 32,00 MB free]
    Total: 1 [279,47 GB] / in use: 1 [279,47 GB] / in no VG: 0 [0   ]
```

and activate the volume:

```
vgchange VolGroup00 -a y
```

```
 1 logical volume(s) in volume group "VolGroup00" now active
```

Now we are able to mount the partition to */mnt/data*:

```
mkdir /mnt/data

mount /dev/VolGroup00/LogVol00 /mnt/data/
```

If you recover data from a hard disk with filenames in UTF-8 format, it might be necessary to convert them to your current non-UTF-8 locale. In my case, the RAID hard disk is from a Fedora Core system with UTF-8 encoded filenames. My target locale is ISO-8859-1. In this case, the Perl script *convmv* helps to convert the filenames to the target locale. **Installation Of convmv**

```
cd /tmp

wget http://j3e.de/linux/convmv/convmv-1.10.tar.gz

tar xvfz convmv-1.10.tar.gz

cd convmv-1.10

cp convmv /usr/bin/convmv
```

To convert all filenames in */mnt/data* to the ISO-8859-1 locale, run this command:

```
convmv -f UTF-8 -t ISO-8859-1 -r --notest /mnt/data/*
```

If you want to test the conversion first, use:

```
convmv -f UTF-8 -t ISO-8859-1 -r /mnt/data/*
```

Links

- **http://j3e.de/linux/convmv/**
- **http://www.linuxjournal.com/article/8874**