

Enable software RAID on XenServer 6.2

Bevore we start:

I shall not use this on production servers. I shall not use this on production servers. I shall not use this on production servers. I shall not use this on production servers.

[Daaa daadidaada daa dadada daaa] (The Simpson theme rolling ...)

Umm... where was I? Right, enabling a software RAID on XenServer 6.2. Possible? Yes. Recommended? No. Field of use: Poor techies like unable to afford a hardware RAID controller for a simple test box ;-)

So, there's plenty of docs out there, telling you the steps on how to enable software RAID using mdadm on a XenServer. Some of the How-To's I came along:

<http://sysadmin.compxtreme.ro/xen-6-0-2-software-raid-installation-procedure>
<http://blog.codeaddict.org/2011/12/04/xenserver-6-0-software-raid1-installation>
<http://www.bauer-power.net/2011/06/easy-software-raid-1-with-xenserver-56.html>

So, I have no point in replicating the steps other's have already written up, so go ahead and check out those HowTo's carefully.

My thoughts, when I was reading up this stuff was like: "Oh my god, this ought to be scripted"

Luckily, some of the guys out there started already off with some scripts, however the ones I came along only covered parts of the procedure, namely steps 1 to 9 (referring to the <http://sysadmin.compxtreme.ro> how-to). So I then essentially took the script provided over at <http://sysadmin.compxtreme.ro> and took it one step further.

It now has no very basic interactive mode: By scanning for available hard drives, it does not rely on hard coded values (which struck me on my server). Thus it does not only detect the system root device (where XenServer bootet from), but also tries to find hard drives that meet the exact same specification (ie. same size).

Well, yes, mdadm allows to build software RAID from different, non-identical hard drives. But from my own expertise, building software RAID off different hard drives will sooner or later bite you!

I prefer to build software RAID always on identical hard drives, which makes things much more predictable. Thus, my incarnation of the setup script does pay attention to the disks being equal in size. Period.

Oh well, I also added the snippets that will finish off your RAID setup, namely steps 10 to 14 as well.

My motivation for doing this enhancement was simple: If you're like me, re-installing things over and over again, because of breaking things just to see how they work, well, then you want something, that really helps you get back up again in no time. I love scripts :-)

So, here's the direct download link: <http://phunsites.net/scripts/x6raidsetup.sh.gz>

You may also run it directly off your XenServer like this:

```
wget -qO- http://phunsites.net/scripts/x6raidsetup.sh.gz | gunzip | bash
```

The script will scan your disk devices and propose to build a software RAID `_if_` and disk identical to your boot drive is found.

During the process, it will create a script file at `/root/x6raidsetup2.sh`, which you can call up after the first reboot in order to finish the setup procedure.

Here's an example output from the first-stage setup script:

```
Trying to figure out root device ...
system's root device: sda
Please wait, locating disk drives ...
got sda (976773168 blocks)
got sdb (976773168 blocks)
device sdc ignored because it's size is not equal to the system's root device
device sdd ignored because it's size is not equal to the system's root device
Disk is 'sdb' is the only suitable disk and has been auto-selected along with root device 'sda'
!! Last Chance !! Do you want to continue? [Y, n] y
clearing the partition table on 'sdb' ...
installing gpt partition table onto 'sdb' ...
creating partition /dev/sdb1 ...
creating partition /dev/sdb2 ...
creating partition /dev/sdb3 ...
sleeping for 10 seconds ...
creating RAID device: /dev/md0 ...
creating bitmaps for /dev/md0 ...
creating RAID device: /dev/md1 ...
creating bitmaps for /dev/md1 ...
creating RAID device: /dev/md2 ...
creating bitmaps for /dev/md2 ...
Formatting the root disk and mount it to /mnt ...
Copying the root filesystem to the RAID array (this will take a while) ...
Updating root filesystem in /mnt/etc/fstab to /dev/md0 ...
Installing the boot loader to 'sdb' ...
Updating initrd ...
Updating boot loader configuration ...
Preparing second stage setup script ...
```

There we are, everything should be configured right now.
Now, please check all output above carefully for any errors.
Resolve any errors, then, if you feel everything is right, perform a reboot.
Enter BIOS setup and set it to boot from the second harddrive 'sdb'.
Once your system has rebooted, you may run this command to finish:

```
/root/x6raidsetup2.sh
```

So, as seen from the script output, you should be switching the device boot order in your BIOS in order to boot off the new mdadm RAID. Then then second-stage script at /root/x6raidsetup2.sh will finish off, which will look like this:

```
Adding original system root device 'sda' to the software RAID ...
Monitoring RAID rebuild progres (press CTRL-C anytime to stop monitoring)...
md1 : active raid1 sda2[0] sdb2[1]
      4193216 blocks [2/2] [UU]
      bitmap: 1/128 pages [4KB], 16KB chunk

md2 : active raid1 sdb3[1]
      479996864 blocks [2/1] [_U]
      bitmap: 0/229 pages [0KB], 1024KB chunk
```

```
md0 : active raid1 sda1[0] sdb1[1]
      4193216 blocks [2/2] [UU]
      bitmap: 0/128 pages [0KB], 16KB chunk
```

unused devices:

Now you should have your mdadm RAID1 up and running in virtually now time (took me 5 minutes). Enjoy!