

Re: [opensuse] boot after installation problem

Source: <http://linux.derkeiler.com/Mailing-Lists/SuSE/2008-02/msg00663.html>

- *From:* Ian Marlier <ian.marlier@xxxxxxxxxxxxxxxxxxxxxx>
 - *Date:* Thu, 07 Feb 2008 01:31:02 -0500
-

On 2/6/08 11:58 PM, "Fred A. Miller" <fmiller@xxxxxxxxxxxxxx> wrote:

Lew Wolfgang wrote:

Hi Folks,

I'm having a problem installing 10.3 x86_64 on a box with many raid disk controllers and disks. It has two disks mounted in the back that are supposed to be the system disks, configured as a raid1 mirror. Alas, this disk shows up after all the other disks, in this particular case as /dev/sde instead of /dev/sda.

A full install goes well until the first boot. The boot fails saying it can't find /dev/sde3 (the root partition). It eventually drops into a limited sh shell running out of ram.

This box worked ok when the system disk appeared as /dev/sdc. I wonder if there is some bug about booting from disks too far removed from /dev/sda?

I've got a call in to the manufacturer to see if there's a way to have the system disks show up as /dev/sda, but no word from them yet.

Anyone seen this behavior? The controllers are 3-Ware, the disks SATA, if that matters.

I don't have an answer for you nor could I find anything except that apparently the company has been bought out by a non-Linux friend company (AMCC) and there are performance complaints now showing up. Here's a reference: <http://lxxer.com/module/forums/t/26605/>

Re: [opensuse] boot after installation problem

I'm running several dozen servers with 3ware cards, all of them running opensuse 10.1 to opensuse 10.3. The kernel modules for those cards have been updated consistently for their latest hardware, backwards compatibility for new cards has always worked, and they've made a point to keep all of their management tools (3dm2) up to date on Linux as well as Windows and other platforms.

As far as vendors in the Linux world go, they're about as close to the top of the heap as anyone. Whatever Sander Marechal's complaints, they don't have any relation to my real-world experience.

That said: I've had problems like yours in the past.

At root, it has to do with something seemingly ridiculous: alphabetical order. When the installation system is probing for hardware, it does so by loading each kernel module and seeing whether any new hardware appears. (Or something else, the effect of which is identical to this.) Because the 3ware modules are loaded very early — 3w-9xxx.so is, after all, very early in alphabetical order — 3ware disks almost always show up as /dev/sda.

However, this order does not hold when the initrd for the installed system is created. There's an explicit order that's created in that case.

Within the initrd filesystem, at the root level, there's a file called "init". It's basically a shell script, and is created dynamically when the initrd is created at the end of the installation process. It includes a section that loads the various modules. Somehow the 3w_9xxx kernel module always ends up being loaded after others, especially the modules for onboard SATA controllers.

The only fix that I've found is to rebuild the initrd. Fortunately, it's pretty easy:

- Boot from a rescue disk;
- chroot to the installed system;
- edit /etc/sysconfig/kernel; for the parameter "INITRD_MODULES", make sure that "3w_9xxx" is listed before any other scsi/sata/ata/etc modules.
- do "mkinitrd"

I do almost all of my installs via autoyast, and have had this problem off and on with builds from suse 9.1 through 10.3. At one time, I actually had a chroot-script (run after the installation of the system but before the first reboot, chrooted within the installed system) that would do this. For opensuse 10.2 I found that I didn't need it, but I'm about to start to set up the 10.3 autobuild for these servers, so who knows...

In my case, my machines almost all have Silicon Image controllers onboard, in addition to the 3ware cards. So the chroot script was something basically like this (I'll see if I can dig the exact script out of the subversion repository where I keep my autoyast configs tomorrow):

Re: [opensuse] boot after installation problem

```
#!/bin/sh
## Alter the INITRD_MODULES line appropriately
sed -i s/sata_sil 3w_9xxx/3w_9xxx sata_sil/g" /etc/sysconfig/kernel
## Rebuild initrd
mkinitrd
exit 0
```

FWIW, my personal opinion is that the bug lies in initrd/mkinitrd/the initrd creation process, and not the card; that is, the process of creating the initrd should be smarter about disk detection order, and force a load of the module needed for the disk on which the bootloader is installed before any others.

You might be able to duplicate the effect of this rebuild by specifying "insmod=3w_9xxx" on the kernel command line, in the grub boot menu, but I won't swear to that.

I know, I know, it's absurd. But those 3ware cards are reliable, very fast, and 3ware/AMCC's service has been outstanding (I've gotten advance-replace cards shipped same day, at no cost, via American Airlines Cargo when FedEx just wasn't fast enough.) I've entirely given up on the idea of using Promise or Adaptec RAID controller cards at this point because I've been so impressed with them.

On the downside, documentation could be a lot better (I suspect that a search of my name and the term "3ware" on both this list and the suse-autoinstall list would turn up a few hits), but, well...once you've played with the cards a bit...

– Ian

—

To unsubscribe, e-mail: opensuse+unsubscribe@xxxxxxxxxxxxx
For additional commands, e-mail: opensuse+help@xxxxxxxxxxxxx