Subject: Re: Capabilities issue
Posted by maratrus on Fri, 03 Sep 2010 14:50:54 GMT
View Forum Message <> Reply to Message

Hi,

thanks for providing this bugreport's link.
I'm at a loss at the moment because there is
a contradiction between vzctl utility and a
kernel from my point of view.

Look, at the kernel side the following checkings
are made

```
asmlinkage long sys_capset(cap_user_header_t header, const cap_user_data_t data)
{
<snip>
if (pid && pid != virt_pid(current) && !capable(CAP_SETPCAP))
        return -EPERM;
<snip>
}
```

As I understand, this piece if code implies that it
is possible to use "capset" system call from inside
the VE. The only thing that must be made is providing
CAP_SETCAP capability to it. The standard way to do
it is via vzctl.

On the other hand, in the bugreport provided by you we can
read

Looks like all linux kernels has init started with CAP_SETPCAP explicitly
disabled due to security implications, so that's why nobody (including vzctl)
can set it.

OK, we can't set it on boot. But if that means that it's
impossible to set this capability at all (after init is
started) it must be mirrored in vzctl's man page.

So, I would recommend you to ask this question directly in
the existing bugreport to get developer's opinion about
this situation and the ways to woraround it.