
Subject: Re: [RFC][PATCH] Devices visibility container
Posted by [Dave Hansen](#) on Thu, 27 Sep 2007 15:46:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, 2007-09-26 at 13:09 -0600, Eric W. Biederman wrote:
> Dave Hansen <haveblue@us.ibm.com> writes:
>
>> On Tue, 2007-09-25 at 07:30 -0600, Eric W. Biederman wrote:
>>> Pavel Emelyanov <xemul@openvz.org> writes:
>>>>
>>>> Oh! Can you provide us an example when after the migration some
>>>> device's major+minor pair change on the same device?
>>>>
>>>> SCSI disks on a SAN. Network accessible block devices.
>>>> All kinds of logical/virtual devices like ttys, the loop device, and
>>>> ramdisks.
>>>>
>>>> It isn't especially frequent that something cares, but fundamentally
>>>> the same issues apply.
>>>>
>>>> To be clear, this just covers cases where an application has
>>>> _internalized_ the device number, right?
>>>>
>>>> Also cases where you want to call mknod in the container.

mknod of device files only, yeah.

>> Most applications should be pretty happy with the devices having
>> persistent device names across a restart, and we can do that with udev
>> and no kernel patching.
>>
>> Yes. But the applications that do internalize stat data from files
>> aren't that uncommon. git, and backup software etc.
>>
>> There is also a fair bit of work that is needed to get sysfs
>> and the hotplug events isolated, when we start allowing mknod etc.
>>
>> Basically if I figure if we are going to deal with this we need to handle
>> the entire problem because these pieces are user visible. I don't
>> think it is a great priority.

Exactly. We have to allow mknod before any of this gets interesting in the least.

-- Dave

Containers mailing list

