Subject: Re: [RFC][PATCH] Devices visibility container
Posted by ebiederm on Mon, 24 Sep 2007 16:47:50 GMT

View Forum Message <> Reply to Message

Pavel Emelyanov <xemul@openvz.org> writes:

> Hi.
>
> At KS we have pointed out the need in some container, that allows
> to limit the visibility of some devices to task within it. I.e.
> allow for /dev/null, /dev/zero etc, but disable (by default) some
> IDE devices or SCSI discs and so on.

NAK

We do not want a control group subsystem for this.

For the short term we can just drop CAP_SYS_MKNOD.

For the long term we need a device namespace for application migration, so they can continue to use devices with the same major+minor number pair after the migration event. Things like ensuring a call to stat on a given file before and after the migration return the exact same information sounds compelling. So I don't think this is even strictly limited to virtual devices anymore. How many applications are there out there that memorize the stat data on a file and so they can detect if it has changed?

If we need something between those two it may make sense to enhance the LSM or perhaps introduce an alternate set security hooks. Still if we are going to need a full device namespace that may be a little much.

Eric

____
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers