Subject: Re: [PATCHSET 3/4] sysfs: divorce sysfs from kobject and driver model
Posted by Sukadev Bhattiprolu on Tue, 16 Oct 2007 22:18:16 GMT
View Forum Message <> Reply to Message

Eric W. Biederman [ebiederm@xmission.com] wrote:
| Greg KH <greg@kroah.com> writes:
|
| > On Fri, Oct 05, 2007 at 06:12:41AM -0600, Eric W. Biederman wrote:
| >> Greg KH <greg@kroah.com> writes:
| >> >
| >> >> Also fun is that the dev file implementation needs to be able to
| >> >> report different major:minor numbers based on which mount of
| >> >> sysfs we are dealing with.
| >> >
| >> > Um, no, that's not going to happen. /dev/sda will _always_ have the
| >> > same major:minor number, as defined by the LSB spec. You can not break
| >> > that at all. So while you might not want to show all mounts
| >> > /sys/devices/block/sda/ the ones that you do, will all have the LSB
| >> > defined major:minor number assigned to it.
| >>
| >> Hmm. If that is in the LSB it must come from
| >> Documentation/devices.txt
| >
| > Yes, that is the requirement.
| >
| >> I'm not after changing the user visible major/minor assignments.
| >
| > Oh, I misunderstood what you wrote above then.
|
| My above sentence is slightly misleading. That should have been.
| I am not after changing the device name to major:minor assignments
| as specified in Documentation/devices.txt.
|
| So within a single device namespace everything is normal and as it
| always has been. Weirdness only ensues when you look across device
| namespaces.
|
| >> Let me see if a concrete example will help. Suppose I have
| >> have a SAN with two disks: disk-1 and disk-2. I have
| >> two machines A and B. On machine A I get the mapping:
| >> sda -> disk-1, sdb ->disk-2. On machine B I wind up with
| >> a different probe order so I get the mapping: sda -> disk-2
| >> sdb ->disk-1.
| >
| > Ok.
| >
| >> To be very clear by sda I mean the block device with major 8 and
| >> minor 0, and by sdb I mean the block device with major 8 and minor

| >> 16.
| >
| > Ok.
| >
| >> So I decide I want an environment on machine B that looks just
| >> like the environment on machine A, so I can bring transfer over
| >> a running program or whatever. So I run around looking at UUID
| >> labels and what not and I discover that the machine B knows disk-1 as
| >> sdb and that machine A knows disk-1 as sda. So I want to say:
| >> /sys/devices/block/sdb show up in this other device namespace as
| >> /sys/devices/block/sda.
|
| >
| > Ah, but if you do that then the "other" device namespace would have
| > /sys/devices/block/sda/dev be 8:16, right?
|
| No. The "other" device namespace I would construct on machine B to
| look just like the device namespace that existed on machine A.
| Making /sys/devices/block/sda would still be 8:0.
|
| So to be very clear on machine B when talking about disk-1 I would have.
| initial device namespace:
| /sys/devices/block/sdb
| /sys/devices/block/sdb/dev 8:16
|
| "other" device namespace:
| /sys/devices/block/sda
| /sys/devices/block/sda/dev 8:0
|
| Similarly on machine B when talking about disk-2 I would have.
| initial device namespace:
| /sys/devices/block/sda
| /sys/devices/block/sda/dev 8:0
|
| "other" device namespace:
| /sys/devices/block/sdb
| /sys/devices/block/sdb/dev 8:16
|
| So between the two devices namespaces on machine B the two disks
| would exchange their user visible identities.

So an application that would migrate from machine A to B has to
use virtual names (like "disk-1" and "disk-2") to access the disk
right ?

|
| Eric
| _____

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

_____

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