Subject: Re: [PATCHSET 3/4] sysfs: divorce sysfs from kobject and driver model
Posted by ebiederm on Tue, 16 Oct 2007 23:54:06 GMT
View Forum Message <> Reply to Message

sukadev@us.ibm.com writes:

> | 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 ?

No. It is worse you need to access a filesystem and probably
a block device that is available on both machine A and machine B.
With care we can introduce appropriate namespaces and namespace semantics
so we can make the names be what we need.

For a classic tricky case think what it would require to migrate
a git archive with checked out files and not need to say
"git-update-index --refresh" before you work with the files.

I used names like disk-1 and disk-2 instead of UUIDs because it
was easier for me to type and think about. You do need some
kind of absolute disk or filesystem identity you can refer back to.

Eric

_____

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