

---

Subject: Re: [PATCH 06/11] sysfs: Implement sysfs tagged directory support.  
Posted by [Daniel Lezcano](#) on Thu, 03 Jul 2008 12:55:38 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

Eric W. Biederman wrote:

> Daniel Lezcano <dlezcano@fr.ibm.com> writes:

>  
>> The kobject events are sent through a netlink message which is not currently per  
>> network namespace. Shouldn't be useful to have a way to retrieve from the  
>> kobject the network namespace or the uevent socket associated with it ? IMHO  
>> having idr in the kobject + netns pointer associated may help to handle the  
>> sysfs isolation and makes the uevent per namespace trivial, no ?

>  
> Grumble. I have been conveniently been forgetting about that socket.  
> Similarly we have the user mode helpers to deal with.

>  
> For this conversation there is a simple answer. All of that is in the  
> kobject layer, and works even when you compile sysfs out of your kernel.  
> Therefore it is a separate problem. And sysfs idr tags have nothing  
> to do with it.

Ah Ok, I am not really familiar with kobject/sysfs so I thought there was a proposition to store the id in the kobject instead of using the tag callbacks, so I figured, perhaps, the idr could have been used in the kobject layer and the sysfs being built upon that.

> It is most definitely something we need to come back to. I bet there  
> are some interesting interactions when you have multiple network devices  
> with the same name generating events.

Yes as mentionned Benjamin, we have the eth0 in the init\_net which is shut down when a network namespace with a netdev with the same name exists. There is a udev rule which ifdown eth0 :)

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

<https://lists.linux-foundation.org/mailman/listinfo/containers>

---