

---

Subject: Re: [PATCH 2/4] sysfs: Implement sysfs managed shadow directory support.

Posted by [dev](#) on Mon, 30 Jul 2007 13:57:09 GMT

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

---

Tejun Heo wrote:

> Hello, Kirill.

>

> Kirill Korotaev wrote:

>

>> Imho then OpenVZ approach with multiple sysfs trees is better.

>> it allows to use cached dentries with multiple sysfs mounts

>> each having different view.

>> It also allows to hide hw-related entries and events from the containers

>> and has quite little modifications in the code.

>

>

> I thought something like supermount plus some twists or fuse based sysfs

> proxy would fit better. Dunno whether or how uevent and polling stuff

> can work that way tho. Note that sysfs no longer keeps dentries and

> inodes pinned. It might make the shared dentry stuff harder.

We simply don't share sysfs dentries/inodes between containers.

It's not that frequently used time critical fs to be super-optimized... :)

I don't like the idea with fuse, since sysfs exports kernel-related stuff, so doing it via user-space would be pain.

Thanks,

Kirill

---

Containers mailing list

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

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

---

---

Subject: Re: [PATCH 2/4] sysfs: Implement sysfs managed shadow directory support.

Posted by [Tejun Heo](#) on Mon, 30 Jul 2007 14:04:51 GMT

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

---

Kirill Korotaev wrote:

> Tejun Heo wrote:

>> I thought something like supermount plus some twists or fuse based sysfs

>> proxy would fit better. Dunno whether or how uevent and polling stuff

>> can work that way tho. Note that sysfs no longer keeps dentries and

>> inodes pinned. It might make the shared dentry stuff harder.

>

> We simply don't share sysfs dentries/inodes between containers.

> It's not that frequently used time critical fs to be super-optimized... :)

OIC, dentries and inodes are not shared. Good then. Agreed that sysfs doesn't need to be super-optimized as long as big machines aren't penalized too much (both memory and cpu cycle wise).

> I don't like the idea with fuse, since sysfs exports kernel-related stuff,  
> so doing it via user-space would be pain.

Yeah, it would be cumbersome to setup but it's also fast and easy to toy with for prototypes at least.

Thanks.

--

tejun

---

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

---