
Subject: Re: [PATCH 06/11] sysfs: Implement sysfs tagged directory support.
Posted by [ebiederm](#) on Tue, 01 Jul 2008 09:20:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

Tejun Heo <htejun@gmail.com> writes:

> Hello, Eric.
>
> Eric W. Biederman wrote:
>>> It's still dynamic from sysfs's POV and I think that will make
>>> maintenance more difficult.
>>
>> Potentially. I have no problem make it clear that things are more static.
>
> Great. :-)
>

> Having enumed tag types limits that a sb can have map to only one tag
> but it doesn't really prevent multiple possibly visible entries which is
> the real unnecessary degrees of freedom. That said, I don't really
> think it's an issue.

Having a single tag type per directory and thus a single tag visible per directory does prevent multiple possible visible entries.

That is we can check when we add the sd if there will be a conflict in the directory.

>> The number of tag types will be low as it is the number of subsystems
>> that use the feature. Simple enough that I expect statically allocating
>> the tag types in an enumeration is a safe and sane way to operate.
>> i.e.
>>
>> enum sysfs_tag_types {
>> SYSFS_TAG_NETNS,
>> SYSFS_TAG_USERNS,
>> SYSFS_TAG_MAX
>> };
>
> I still would prefer something which is more generic. The abstraction
> is clearer too. A sb shows untagged and a set of tags. A sd can either
> be untagged or tagged (a single tag).

That is the abstraction now.

The only difference is how we represent the set of tags.
I use an array of the valid tags.
You use a bitmap.

And array allows the lookup of the tag I am looking for before I search for the sd. An bitmap requires me to compare each entry.

For me that is a deal breaker. Currently in certain pathological cases we have scaling issues with sysctl and sysfs that we can have enormous directories that start running slowly. To fix lookup performance requires that we know the full name before we do the directory search which is the name string and the tag.

So I having a type of tag as being of fundamental importance in the interface now so we don't need to refactor all of the users later. In addition to the fact that we need the type to know how to set the tags when mounting a superblock and when given a new kobject to create an sd for.

We could make the types dynamic rather than a static enumeration but that seems needless complexity for now.

> Using ida (or idr if a pointer for private data is necessary) is really
> easy. It'll probably take a few tens of lines of code. That said, I
> don't think I have enough rationale to nack what you described. So, as
> long as the tags are made static, I won't object.

Sounds good. The only justification I can think of for ida tags is that they are smaller, and so can keep the sysfs_dirents smaller. Which occasionally is a significant concern. Still that should be an optimization that we can apply later, as it is not a structural difference in the code.

Just to confirm. Do you the two operations:

mount_tag - called only when the sb is mounted

kobject_tag - called when we create new sd or rename an sd

Cause you to view an the tags as dynamic?

Eric

Containers mailing list

Containers@lists.linux-foundation.org

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