
Subject: Re: [PATCH 00/11] sysfs tagged directories V5

Posted by [serue](#) on Tue, 10 Jun 2008 22:50:24 GMT

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

Quoting Benjamin Thery (benjamin.thery@bull.net):

> Greg,

>

> Here is an updated version of the sysfs tagged directories that improves
> a bit the situation over the previous one.

>

> I've modified the patch 09 ("Enable tagging for net_class directories in
> sysfs") to be a bit less intrusive in sysfs core. I removed the #ifdef'd
> parts you didn't like in fs/sysfs/mount.c, and replaced it by a generic
> routine sysfs_ns_exit() that is called, if needed, by the namespace when
> it exits. This routine goes through every sysfs super blocks and calls
> the callback passed by the namespace to clean its tag.

>

> The patch is now splitted in two:

> * 09/11: the generic routine,

> * 10/11: the remaining network parts.

>

> The generic part can be merge with patch 05 ("sysfs: Implement sysfs
> tagged directory support.") but I left it separate for now to ease
> reviews.

Thanks. Not nearly as radical as I'd feared.

Greg, does this address your concern? It stops the need to put any non-braindead code (i.e. the _exit notifier stuff) in fs/sysfs/mount.c for any new namespaces needing sysfs tagging, leaving only simple comparison/assignments.

> Serge's patch for user namespace is modified to use this new service too.

> No more #ifdef CONFIG_NET or #ifdef CONFIG_USER_NS in fs/sysfs/mount.c
> now.

>

>

> But, currently, a new namespace that wants to add its tag to sysfs dirs
> still need to modify fs/sysfs/mount.c in a few routines to manage the
> new tag member added in struct sysfs_tag_info: sysfs_fill_super(),
> sysfs_test_super(), sysfs_kill_sb() (see the last two patches).
> These changes are only the initialization and a bunch of comparisons.

>

> If we really want to go further, to get rid of these, I've thought
> about:

>

> * Extending sysfs_tagged_dir_operations with super blocks operations:

> - fill_sb_tag, test_sb_tag, kill_sb_tag

>
> * Add routines in sysfs to allow registration/unregistration of these
> operations structs in a list:
> - sysfs_register_tagged_dir_ops()...
>
> * Each subsystem concerned implements and registers its operations at
> boot.
>
> * In sysfs_fill_super(), sysfs_test_super() and sysfs_kill_sb(), add
> loops to go through all registered operations structs and calls the
> corresponding operations if it's present.
>
> But... I thought it was a bit overkill for the few namespaces that will
> actually need sysfs tagged directories.

Agreed. We expect device namespaces... anything else?

thanks,
-serge

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