
Subject: [PATCH 00/11] sysfs tagged directories V6
Posted by [Benjamin Thery](#) on Wed, 18 Jun 2008 17:07:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

Greg, Andrew,

Here is yet another updated version of Eric Biederman's patchset to implement tagged directories in sysfs ported on top of 2.6.26-rc5-mm3. This may be easier for you to review as 2.6.26-rc2-mm1 was getting a bit old now.

There is no major changes since the last version which introduced some changes to address concerns from Greg about the patch "Enable tagging for net_class directories" being to intrusive in sysfs core.

Refer to the Changelog below and <http://thread.gmane.org/gmane.linux.kernel/690799> to read the details about the proposed changes (I didn't copy them here to save some space).

Andrew,

Can you consider merging this patchset in -mm until Greg has time to re-review it and take it (or reject it)?

Thanks,
Benjamin

(Below you'll find the traditional introduction for sysfs tagged dirs and the updated changelog)

--

With the introduction of network namespaces, there can be duplicate network interface names on the same machine. Indeed, two network interfaces can have the same name if they reside in different network namespaces.

- * Network interfaces names show up in sysfs.
- * Today there is nothing in sysfs that is currently per namespace.
- * Therefore we need to support multiple mounts of sysfs each showing a different network namespace.

We introduce tagged directories in sysfs for this purpose.

Of course the usefulness of this feature is not limited to network stuff: Serge Hallyn wrote a patch to fix a similar issue with user namespaces based on this patchset. His patch is included at the end of the patchset.

Tested with and without SYSFS_DEPRECATED. No regression found so far.

Changelog

* V6:

- Ported to 2.6.26-rc5-mm3
- Patch 11 (users) Removed an unused kset member from struct user_namespace left from a previous version of the patch.

* V5:

- Make namespace tags a bit less intrusive in sysfs core:
 - New patch 09: Added a generic sysfs_ns_exit routine called by exiting namespaces. A callback is passed to this routine to execute the subsystem specific code.
 - Modified patches 09 and 10 (now 10 and 11) ("netns tagging" and "users tagging") to use this new routine instead of adding #ifdef'd code in fs/sysfs/mount.c.
- Added missing -ENOMEM in fs/sysfs/dir.c:prep_rename() (Roel Kluin)

* V4:

- Ported to 2.6.26-rc2-mm1
- Updated patch for user namespace by Serge Hallyn (patch 10).

* V3:

- Removed patch 10 ("avoid kobject name conflict with different namespaces"), a better one was provided by Eric.
- Removed patch 11 ("sysfs: user namespaces: add ns to user_struct"), Serge needs to rework some parts of it.
- Change Acked-by: to Signed-off-by:, someone told me it is more appropriate (as I'm in the delivery path).

Here is the announcement Eric wrote back in December to introduce his patchset:

"

Now that we have network namespace support merged it is time to revisit the sysfs support so we can remove the dependency on !SYSFS.

[...]

The bulk of the patches are the changes to allow multiple sysfs superblocks.

Then comes the tagged directory sysfs support which uses information captured at mount time to decide which object with which tag will appear in a directory.

Then the support for renaming and deleting objects where the source may be ambiguous because of tagging.

Then finally the network namespace support so it is clear how all of this tied together.

"

Regards,
Benjamin

--

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