
Subject: [PATCH 09/11] sysfs: add sysfs_ns_exit routine
Posted by [Benjamin Thery](#) on Fri, 06 Jun 2008 15:48:24 GMT
[View Forum Message](#) <> [Reply to Message](#)

Add sysfs routine sysfs_ns_exit() to allow a namespace to go away while sysfs is still mounted.

The exiting namespace calls this routine and pass it a callback to be called for every sysfs superblocks present. The callback contains the necessary code to clean the superblock tag data associated with this namespace.

Signed-off-by: Benjamin Thery <benjamin.thery@bull.net>

fs/sysfs/mount.c | 21 +++++
include/linux/sysfs.h | 8 +++++
2 files changed, 29 insertions(+)

Index: linux-mm/fs/sysfs/mount.c

=====

--- linux-mm.orig/fs/sysfs/mount.c

+++ linux-mm/fs/sysfs/mount.c

@@ -181,6 +181,27 @@ restart:

spin_unlock(&sb_lock);

}

/* Clean sysfs tags related to a given namespace when it exits */

void sysfs_ns_exit(void (*func)(struct sysfs_tag_info *, void *), void *data)

{

+ /* Allow the namespace to go away while sysfs is still mounted. */

+ struct super_block *sb;

+ mutex_lock(&sysfs_rename_mutex);

+ sysfs_grab_supers();

+ mutex_lock(&sysfs_mutex);

+ list_for_each_entry(sb, &sysfs_fs_type.fs_supers, s_instances) {

+

+ struct sysfs_super_info *info = sysfs_info(sb);

+ /* Call the cleaning routine provided by the namespace.

+ * data is the current namespace id passed by the namespace.

+ */

+ func(&info->tag, data);

+ }

+ mutex_unlock(&sysfs_mutex);

+ sysfs_release_supers();

+ mutex_unlock(&sysfs_rename_mutex);

+}

+

int __init sysfs_init(void)

```

{
    int err = -ENOMEM;
Index: linux-mm/include/linux/sysfs.h
=====
--- linux-mm.orig/include/linux/sysfs.h
+++ linux-mm/include/linux/sysfs.h
@@ -134,6 +134,9 @@ void sysfs_printk_last_file(void);
int sysfs_enable_tagging(struct kobject *kobj,
    const struct sysfs_tagged_dir_operations *tag_ops);

+void sysfs_ns_exit(void (*func)(struct sysfs_tag_info *, void *),
+    void *data);
+
extern int __must_check sysfs_init(void);

#else /* CONFIG_SYSFS */
@@ -247,6 +250,11 @@ static inline int sysfs_enable_tagging(s
    return 0;
}

+static inline void sysfs_ns_exit(void (*func)(struct sysfs_tag_info *, void *),
+    void *data)
+{
+}
+
static inline int __must_check sysfs_init(void)
{
    return 0;
}

--

```

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