
Subject: [PATCH 13/15] Revert "netns: Fix device renaming for sysfs"

Posted by [ebiederm](#) on Fri, 04 Jul 2008 01:21:28 GMT

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

This reverts commit aaf8cdc34ddba08122f02217d9d684e2f9f5d575.

Drivers like the ipw2100 call device_create_group when they are initialized and device_remove_group when they are shutdown. Moving them between namespaces deletes their sysfs groups early.

In particular the following call chain results.

netdev_unregister_kobject -> device_del -> kobject_del -> sysfs_remove_dir

With sysfs_remove_dir recursively deleting all of it's subdirectories, and nothing adding them back.

Ouch!

Therefore we need to call something that ultimate calls sysfs_mv_dir as that sysfs function can move sysfs directories between namespaces without deleting their subdirectories or their contents. Allowing us to avoid placing extra boiler plate into every driver that does something interesting with sysfs.

Currently the function that provides that capability is device_rename. That is the code works without nasty side effects as originally written.

So remove the misguided fix for moving devices between namespaces. The bug in the kobject layer that inspired it has now been recognized and fixed.

Signed-off-by: Eric W. Biederman <ebiederm@xmission.com>

```
net/core/dev.c      | 4 +---
net/core/net-sysfs.c | 7 +-----
net/core/net-sysfs.h | 2 +-
3 files changed, 3 insertions(+), 10 deletions(-)
```

```
diff --git a/net/core/dev.c b/net/core/dev.c
```

```
index fca23a3..585584d 100644
```

```
--- a/net/core/dev.c
```

```
+++ b/net/core/dev.c
```

```
@@ -3806,7 +3806,6 @@ int register_netdevice(struct net_device *dev)
{
}
```

```
- netdev_initialize_kobject(dev);
ret = netdev_register_kobject(dev);
if (ret)
```

```

    goto err_uninit;
@@ -4239,8 +4238,7 @@ int dev_change_net_namespace(struct net_device *dev, struct net
*net, const char
}

/* Fixup kobjects */
- netdev_unregister_kobject(dev);
- err = netdev_register_kobject(dev);
+ err = device_rename(&dev->dev, dev->name);
  WARN_ON(err);

/* Add the device back in the hashes */
diff --git a/net/core/net-sysfs.c b/net/core/net-sysfs.c
index 90e2177..4e7b847 100644
--- a/net/core/net-sysfs.c
+++ b/net/core/net-sysfs.c
@@ -449,6 +449,7 @@ int netdev_register_kobject(struct net_device *net)
    struct device *dev = &(net->dev);
    struct attribute_group **groups = net->sysfs_groups;

+ device_initialize(dev);
    dev->class = &net_class;
    dev->platform_data = net;
    dev->groups = groups;
@@ -469,12 +470,6 @@ int netdev_register_kobject(struct net_device *net)
    return device_add(dev);
}

-void netdev_initialize_kobject(struct net_device *net)
-{
- struct device *device = &(net->dev);
- device_initialize(device);
-}
-
int netdev_kobject_init(void)
{
    return class_register(&net_class);
diff --git a/net/core/net-sysfs.h b/net/core/net-sysfs.h
index 14e7524..f5f108d 100644
--- a/net/core/net-sysfs.h
+++ b/net/core/net-sysfs.h
@@ -4,5 +4,5 @@
@@ -4,5 +4,5 @@
int netdev_kobject_init(void);
int netdev_register_kobject(struct net_device *);
void netdev_unregister_kobject(struct net_device *);
-void netdev_initialize_kobject(struct net_device *);
+
#endif

```

--

1.5.3.rc6.17.g1911

Containers mailing list

Containers@lists.linux-foundation.org

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