
Subject: [PATCH] lost content of /proc/sys/fs/binfmt_misc

Posted by [den](#) on Wed, 05 Dec 2007 14:35:18 GMT

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

/proc/sys/fs/binfmt_misc dentry disappeared during d_revalidate.

d_revalidate only dentries from shadowed one and below.

http://bugzilla.kernel.org/show_bug.cgi?id=9504

CC: Eric W. Biederman <ebiederm@xmission.com>

CC: Marcus Better <marcus@better.se>

Signed-off-by: Denis V. Lunev <den@openvz.org>

```
diff --git a/fs/proc/generic.c b/fs/proc/generic.c
index 5fccfe2..1497ac4 100644
--- a/fs/proc/generic.c
+++ b/fs/proc/generic.c
@@ -380,12 +380,17 @@ static int proc_revalidate_dentry(struct dentry *dentry, struct
nameidata *nd)
    return 0;
}

-static struct dentry_operations proc_dentry_operations =
+static struct dentry_operations proc_dentry_shadow_operations =
{
    .d_delete = proc_delete_dentry,
    .d_revalidate = proc_revalidate_dentry,
};

+static struct dentry_operations proc_dentry_operations =
+{
+    .d_delete = proc_delete_dentry,
+};
+
/*
 * Don't create negative dentries here, return -ENOENT by hand
 * instead.
@@ -394,6 +399,7 @@ struct dentry *proc_lookup(struct inode * dir, struct dentry *dentry, struct
nam
{
    struct inode *inode = NULL;
    struct proc_dir_entry * de;
+    int use_shadow = 0;
    int error = -ENOENT;

    lock_kernel();
@@ -406,8 +412,10 @@ struct dentry *proc_lookup(struct inode * dir, struct dentry *dentry, struct
nam
        if (!memcmp(dentry->d_name.name, de->name, de->namelen)) {
```

```
unsigned int ino;

- if (de->shadow_proc)
+ if (de->shadow_proc) {
    de = de->shadow_proc(current, de);
+   use_shadow = 1;
+ }
    ino = de->low_ino;
    de_get(de);
    spin_unlock(&proc_subdir_lock);
@@ -423,6 +431,8 @@ struct dentry *proc_lookup(struct inode * dir, struct dentry *dentry, struct
nam

if (inode) {
    dentry->d_op = &proc_dentry_operations;
+   dentry->d_op = use_shadow ?
+   &proc_dentry_shadow_operations : dentry->d_parent->d_op;
    d_add(dentry, inode);
    return NULL;
}
```
