
Subject: [PATCH -mm] Fix race between proc_readdir and remove_proc_entry
Posted by [Alexey Dobriyan](#) on Sun, 11 Mar 2007 11:52:44 GMT
[View Forum Message](#) <> [Reply to Message](#)

> -procfs-fix-race-between-proc_readdir-and-remove_proc_entry. patch
> +fix-race-between-proc_get_inode-and-remove_proc_entry.patch
>
> Updated. Looks sane.

Why have you dropped the first patch? Resending slightly fixed version of it.

[PATCH -mm] Fix race between proc_readdir and remove_proc_entry

From: "Darrick J. Wong" <djwong@us.ibm.com>

Fix the following race:

```
proc_readdir  remove_proc_entry
=====
```

```
spin_lock(&proc_subdir_lock);
[choose PDE to start filldir from]
spin_unlock(&proc_subdir_lock);
    spin_lock(&proc_subdir_lock);
    [find PDE]
    [free PDE, refcount is 0]
    spin_unlock(&proc_subdir_lock);
    /* boom */
```

```
if (filldir(dirent, de->name, ...
```

```
[de_put on error path --adobriyan]
```

Signed-off-by: Darrick J. Wong <djwong@us.ibm.com>

Signed-off-by: Alexey Dobriyan <adobriyan@sw.ru>

```
fs/proc/generic.c | 11 ++++++++--
1 file changed, 9 insertions(+), 2 deletions(-)
```

```
--- a/fs/proc/generic.c
+++ b/fs/proc/generic.c
@@ -478,14 +478,21 @@ int proc_readdir(struct file * filp,
 {

    do {
+ struct proc_dir_entry *next;
```

```

+
+ /* filldir passes info to user space */
+ de_get(de);
+ spin_unlock(&proc_subdir_lock);
+ if (filldir(dirent, de->name, de->namelen, filp->f_pos,
-     de->low_ino, de->mode >> 12) < 0)
+     de->low_ino, de->mode >> 12) < 0) {
+     de_put(de);
+     goto out;
+ }
+ spin_lock(&proc_subdir_lock);
+ filp->f_pos++;
- de = de->next;
+ next = de->next;
+ de_put(de);
+ de = next;
+ } while (de);
+ spin_unlock(&proc_subdir_lock);
+ }

```
