
Subject: Re: [PATCH] proc: Do not invalidate dentries with submounts

Posted by [den](#) on Thu, 06 Dec 2007 10:30:45 GMT

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

you have changed the behavior of revalidation by shadows. I think it will be better to restore it and keep new one for shadows (and below) only, which has been done by my yesterday patch.

Regards,

Den

Eric W. Biederman wrote:

```
> If the dcache path to a mount point is ever broken it becomes  
> impossible to umount it, and we leak a vfsmount. Therefore it is not  
> valid to invalidate dentries with mount points at or below them.  
>  
> This patch uses the have_submounts test as the other network  
> filesystem revalidate routines do.  
>  
> Signed-off-by: Eric W. Biederman <ebiederm@xmission.com>  
> ---  
> fs/proc/base.c | 9 ++++++++  
> fs/proc/generic.c | 5 +++++  
> 2 files changed, 14 insertions(+), 0 deletions(-)  
>  
> diff --git a/fs/proc/base.c b/fs/proc/base.c  
> index 0e71707..552d752 100644  
> --- a/fs/proc/base.c  
> +++ b/fs/proc/base.c  
> @@ -1216,6 +1216,9 @@ static int pid_revalidate(struct dentry *dentry, struct nameidata *nd)  
>     put_task_struct(task);  
>     return 1;  
> }  
> /* Force validity if something is mounted under us */  
> + if (inode && S_ISDIR(inode->i_mode) && have_submounts(dentry))  
> + return 1;  
>     d_drop(dentry);  
>     return 0;  
> }  
> @@ -1393,6 +1396,9 @@ static int tid_fd_revalidate(struct dentry *dentry, struct nameidata  
 *nd)  
> }  
>     put_task_struct(task);  
> }  
> /* Force validity if something is mounted under us */  
> + if (inode && S_ISDIR(inode->i_mode) && have_submounts(dentry))  
> + return 1;  
>     d_drop(dentry);
```

```
>     return 0;
> }
> @@ -2056,6 +2062,9 @@ static int proc_base_revalidate(struct dentry *dentry, struct
nameidata *nd)
>     put_task_struct(task);
>     return 1;
> }
> /* Force validity if something is mounted under us */
> + if (inode && S_ISDIR(inode->i_mode) && have_submounts(dentry))
> + return 1;
>     d_drop(dentry);
>     return 0;
> }
> diff --git a/fs/proc/generic.c b/fs/proc/generic.c
> index 4abd568..233dc dc 100644
> --- a/fs/proc/generic.c
> +++ b/fs/proc/generic.c
> @@ -370,6 +370,11 @@ static int proc_delete_dentry(struct dentry * dentry)
>
> static int proc_revalidate_dentry(struct dentry *dentry, struct nameidata *nd)
> {
>     struct inode *inode = dentry->d_inode;
> +
> /* Force validity if something is mounted under us */
> + if (inode && S_ISDIR(inode->i_mode) && have_submounts(dentry))
> + return 1;
>     d_drop(dentry);
>     return 0;
> }
```
