
Subject: Re: [PATCH 3/20] Introduce MS_KERNMOUNT flag
Posted by [Oleg Nesterov](#) on Sat, 11 Aug 2007 11:20:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 08/11, Christoph Hellwig wrote:

>
> On Fri, Aug 10, 2007 at 03:47:55PM +0400, xemul@openvz.org wrote:
> > This flag tells the .get_sb callback that this is a kern_mount() call
> > so that it can trust *data pointer to be valid in-kernel one. If this
> > flag is passed from the user process, it is cleared since the *data
> > pointer is not a valid kernel object.
> >
> > Running a few steps forward - this will be needed for proc to create the
> > superblock and store a valid pid namespace on it during the namespace
> > creation. The reason, why the namespace cannot live without proc mount
> > is described in the appropriate patch.
>
> I don't like this at all. We should never pass kernel and userspace
> addresses through the same pointer. Maybe add an additional argument
> to the get_sb prototype instead. But this whole idea of mounting /proc
> from kernelspace sounds like a really bad idea to me. /proc should
> never be mounted from the kernel but always normally from userspace.

Can't comment because I don't understand vfs at all, and perhaps I just misunderstood you.

But could you clarify? We already create internal proc mount from kernel space, proc_root_init() does this. With this series we are doing the same when a new namespace is created.

Thanks,

Oleg.
