

---

Subject: Re: NFS causing oops when freeing namespace  
Posted by [Trond Myklebust](#) on Wed, 17 Jan 2007 22:46:42 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Wed, 2007-01-17 at 15:30 -0700, Eric W. Biederman wrote:

> Cedric Le Goater <clg@fr.ibm.com> writes:

> >

> > your first analysis was correct : exit\_task\_namespaces() should be moved

> > above exit\_notify(tsk). It will require some extra fixes for nsproxy

> > though.

>

> I think the only issue is the child\_reaper and currently we only have one of

> those. When we really do the pid namespace we are going to have to revisit

> this. My gut feel says that we won't be able to exit our pid namespace until

> the process is waited on. So we may need to break up exit\_task\_namespace into

> individual components.

It makes little sense, afaics, to have an interruptible sleep in something like lockd\_down() if you have no pid space or signal handling.

That isn't the only place where the process has to wait in an NFS unmount, BTW. Things like rpc client cleanup (waiting for all the RPC tasks to complete) also tend to lead to interruptible sleeps.

Trond

---

Containers mailing list

[Containers@lists.osdl.org](mailto:Containers@lists.osdl.org)

<https://lists.osdl.org/mailman/listinfo/containers>

---