
Subject: Re: [RFC][PATCH 1/2] remove proc_mnt's use or killing inodes

Posted by [ebiederm](#) on Tue, 20 Feb 2007 20:08:48 GMT

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

Dave Hansen <hansendc@us.ibm.com> writes:

> We use proc_mnt as a shortcut to find a superblock on which
> to go killing /proc inodes. This will break if we ever have
> more than one /proc superblock. So, use the superblock list
> to go find each of the /proc sb's and kill inodes on each
> superblock.

>

> This does introduce an extra lock grab from what was there
> before, but the list should be only 1 long 99% of the time,
> and we don't exactly remove proc entries in hot paths. Note
> that this *isn't* the path that we use to get rid of the
> actual /proc pid entries. Those are a different beast.

Looks good.

The only other thing we could do is turn proc/generic.c into
it's own filesystem that we mount with magic mounts like
nfs does when crossing mount points.

I think that would give us similar problems with killing inodes
when a proc entry is removed.

Eric

Containers mailing list

Containers@lists.osdl.org

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