
Subject: Re: [PATCH 0/59] Cleanup sysctl
Posted by [ebiederm](#) on Wed, 17 Jan 2007 19:02:27 GMT
[View Forum Message](#) <> [Reply to Message](#)

Kirill Korotaev <dev@sw.ru> writes:

> Eric, really good job!
>
> Patches: 1-13, 15-24, 26-32, 34-44, 46-49, 52-55, 57 (all except below)
> Acked-By: Kirill Korotaev <dev@openvz.org>
>
> 14/59 - minor (extra space)
> 25/59 - minor note
> 33/59 - not sorted sysctl IDs
> 45/59 - typo
> 50/59 - copyright/file note
> 51/59 - copyright/file name/kconfig option notes
>
> 56,58,59/59 - will review tomorrow
>
> another issue I have to think over is removal of de->owner.
> Alexey Dobriyan has sent recently patching fixing /proc <-> modules refcounting.
> I guess w/o these patches your changes are not safe if proc_handler or strategy
> are functions from the module.

sysctl uses the logic in use_table/unuse_table to keep it safe from module remove while it is in use. And it does the logic in the generic code in either do_rw_proc or do_sysctl. This definitely works on the sys_sysctl path and it appears to work in the do_rw_proc case, things are a little trickier there so someone may have missed a race somewhere. In my rewrite of proc it works exactly like the binary case so we are good there.

It is certainly the intention of the sysctl implementation that users should not have to set de->owner. So if there is a problem with removing de->owner it is a bug in the sysctl implementation not in the code where it was removed.

Normal proc users definitely have to set de->owner to be safe, but sysctl has always been it's own thing, with different rules.

Eric

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
