
Subject: Re: [PATCHSET] 2.6.20-rc4-mm1-lxc2
Posted by [ebiederm](#) on Wed, 17 Jan 2007 01:46:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

Daniel Lezcano <dlezcano@fr.ibm.com> writes:

>
> Hi Dmitry,
>
> we are experiencing NULL address access when using the nsproxy in
> push_net_ns function without any unshare.
>
> It appears the exit_task_namespace function sets current->nsproxy to
> NULL and we are interrupted by an incoming packet. The netif_receive_skb
> does push_net_ns(dev->net_ns). The push_net_ns function retrieves the
> current->nsproxy to use it. But it was previously set to NULL by the
> exit_task_namespace function.
>
> The bug can be reproduced with the following command launched from
> another host.
>
> while \$(true); do ssh myaddress ls > /dev/null && echo -n .; done
>
> After a time (between 1 second - 3 minutes), the kernel panics.
>
> I think this will be very hard to fix and perhaps we should redesign
> some part. Instead of using nsproxy swapping, perhaps we should pass
> net_ns as parameter to functions, but that will breaks a lot of API.
>
> What is your feeling on that ?

After looking at several things primarily ramifications of file descriptor passing I have concluded that a magic global variable in the task struct is almost certainly the wrong thing to do. And the more I look at it the task is usually the wrong location to look to see what network namespace you are in.

To that effect I have been preparing a patchset for discussion targeting the end of this week to have it ready, in an easily reviewable format.

Eric

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