
Subject: Re: Purpose of `__net_exit` & friends
Posted by [davem](#) on Fri, 29 Feb 2008 20:23:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

From: Sam Ravnborg <sam@ravnborg.org>
Date: Fri, 29 Feb 2008 21:12:47 +0100

> On Fri, Feb 29, 2008 at 11:23:42AM -0800, David Miller wrote:
> > From: "Denis V. Lunev" <den@openvz.org>
> > Date: Fri, 29 Feb 2008 16:40:58 +0300
> >
> > > +static void __net_exit icmpv6_sk_exit(struct net *net)
> > > +{
> > >
> > > Since you call this from `__new_init` code, won't this
> > > create a section conflict when ipv6 is built statically
> > > into the kernel?
> >
> > I have noticed this `__net_init`/`__net_exit` stuff before
> > but never got around to ask about it.
> >
> > Whats the actual purpose and are there any strict
> > rules as to what may be called from where?
> > If there are we should apply the same checks as we do
> > for `__devinit`/`__devexit` and friends.

Basically, if network namespace support isn't enabled,
these functions can act like normal `__init` and `__exit`
functions.

Otherwise we have to keep them around so that they can
be called when namespaces are created and destroyed.

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
