Subject: Re: [PATCH 1/2][INET] (resend) Fix potential kfree on vmalloc-ed area of request_sock_queue
Posted by davem on Thu, 15 Nov 2007 10:58:03 GMT
View Forum Message <> Reply to Message

From: Eric Dumazet <dada1@cosmosbay.com>
Date: Thu, 15 Nov 2007 10:21:01 +0100

> On Thu, 15 Nov 2007 11:41:37 +0300
> Pavel Emelyanov <xemul@openvz.org> wrote:
>
> > The request_sock_queue's listen_opt is either vmalloc-ed or
> > kmalloc-ed depending on the number of table entries. Thus it
> > is expected to be handled properly on free, which is done in
> > the reqsk_queue_destroy().
> >
> > However the error path in inet_csk_listen_start() calls
> > the lite version of reqsk_queue_destroy, called
> > __reqsk_queue_destroy, which calls the kfree unconditionally.
> >
> > Fix this and move the __reqsk_queue_destroy into a .c file as
> > it looks too big to be inline.
> >
> > As David also noticed, this is an error recovery path only,
> > so no locking is required and the lopt is known to be not NULL.
> >
> > Signed-off-by: Pavel Emelyanov <xemul@openvz.org>
> >
>
> Acked-by: Eric Dumazet <dada1@cosmosbay.com>
>
> Thank you for finding this bug Pavel

Indeed.

I applied this, but what I did was I combined both changes
into one because to me they logically belong together.

Thanks again Pavel!