Subject: Re: unlock iptables in netns
Posted by Alexey Dobriyan on Mon, 16 Jun 2008 11:04:34 GMT
View Forum Message <> Reply to Message

On Mon, Jun 16, 2008 at 12:26:03PM +0200, Patrick McHardy wrote:
> Patrick McHardy wrote:
>> Alexey Dobriyan wrote:
>>> Hi,
>>>
>>> Den basically banned iptables in netns via this patch
>>>
>>> --- a/net/netfilter/core.c
>>> +++ b/net/netfilter/core.c
>>> ...
>>> , however, at least some of netfilter pieces are ready for usage in netns
>>> and it would be nice to unlock them before release.
>>>
>>> If I'm deciphering chengelog correctly it's all about code which does
>>> nf_register_hook{,s} but not netns-ready itself:
>>>
>>> br_netfilter.c
>>> iptable_mangle (via ip_route_me_harder)
>>> conntracking (both IPv4 and IPv6)
>>> NAT
>>> ...
>>> Patch above can be applied and we can mark above list as "depends
>>> !NET_NS"
>>> and move on.
>>>
>>> Comments? Den, was there something else you're afraid of?
>> That might result in some bad surprises for people how have already
>> turned on NET_NS. I'd prefer a way that doesn't potentially disable
>> half the netfilter options in existing configs.
>
>
> By the way, is there already work done for conntrack/NAT namespace
> support? I have this patch that uses marks for something very similar
> that should be easy to adjust.

Yes, right now I'm fighting something which looks like double free
of conntrack during clone(CLONE_NEWNET)/exit test despite none created
in netns. And unknown to me dimensions of input and output packet codepaths.
:^)

Preliminaty details:
 struct nf_conn::ct_net which pins netns
 netns of expectation is netns of master conntrack by definition
 per-netns conntrack hash

per-netns expect hash
per-netns unconfirmed list