
Subject: Re: 2.6.20-lxc8: kernel panic with af_unix as module
Posted by [Pierre Peiffer](#) on Fri, 04 May 2007 08:26:24 GMT
[View Forum Message](#) <> [Reply to Message](#)

>
>> In my case, when registering per_net_unix_root_table (in function
>> unix_sysctl_register in file net/unix/sysctl_net_unix.c), it tries to access the
>> child (which is per_net_unix_net_table) at the address 0x00000080 which is the
>> address of this symbol in the section .data.pernet of the module object, but not
>> a virtual address.
>
> Hmm. It should be a global offset in the .data.pernet section.

Yes, indeed, it is.

If I do some printk (in unix_sysctl_register) just before the panic, I get:

```
per_net(unix_root_table, net) = 0xf7e2dce0
per_net(unix_root_table, net)->child = 0x00000080  <= cause the panic
```

=> should be equal to this:

```
__per_net_base(unix_net_table)= 0xc03c7f40
```

And:

```
=====
$ objdump -D -j .data.pernet net/unix/unix.ko
```

```
net/unix/unix.ko:    file format elf32-i386
```

Disassembly of section .data.pernet:

```
00000000 <per_net_sysctl_unix_max_dgram_qlen>:
 0: 0a 00          or    (%eax),%al
```

```
...
```

```
00000020 <per_net_unix_root_table>:
 20: 03 00          add    (%eax),%eax
[...]
```

```
00000080 <per_net_unix_net_table>:
 80: 04 00          add    $0x0,%al
 82: 00 00          add    %al,(%eax)
```

```
...
```

The init of per_net(unix_root_table, net)->child is done with the offset of

`__per_net_base(unix_net_table)` in the module, and never translated to its virtual address.

It looks like the module loader translates/relocates correctly the address of `__per_net_base(unix_net_table)`, but the not the address (value of) `per_net(unix_root_table, net)->child` ???

But if I do:

```
$ objdump -r -j .data.pernet net/unix/unix.ko
```

net/unix/unix.ko: file format elf32-i386

RELOCATION RECORDS FOR [.data.pernet]:

OFFSET	TYPE	VALUE
00000024	R_386_32	.rodata.str1.1
00000034	R_386_32	.data.pernet <== per_net__unix_root_table->child (*)
00000070	R_386_32	.data.pernet
00000084	R_386_32	.rodata.str1.1
00000094	R_386_32	.data.pernet <== per_net__unix_net_table->>child (?)
000000e4	R_386_32	.rodata.str1.1
000000e8	R_386_32	per_net__sysctl_unix_max_dgram_qlen
000000f8	R_386_32	proc_dointvec

(*) If I well read/understand, this relocation entry should correspond to `per_net__unix_root_table->child` and should be translated at load time, but it isn't ?

>
> Thanks. I will go back and look but I don't plan on back porting
> anything for 2.6.20. I'm lazy and do not have enough hours in the
> day. :)

No problem.
Thanks.

--
Pierre

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