

---

Subject: Re: [PATCH] iptables compat code error way & module refcounting fix

Posted by [Mishin Dmitry](#) on Mon, 30 Oct 2006 15:57:35 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

On Monday 30 October 2006 18:45, Patrick McHardy wrote:

> Dmitry Mishin wrote:

> > This patch fixes bug in iptables modules refcounting on compat error way.

> >

> > As we are getting modules in check\_compat\_entry\_size\_and\_hooks(), in case

> > of later error, we should put them all in translate\_compat\_table(), not

> > in the compat\_copy\_entry\_from\_user() or compat\_copy\_match\_from\_user(), as

> > it is now.

> >

> > Signed-off-by: Dmitry Mishin <dim@openvz.org>

> > Acked-by: Vasily Averin <vvs@openvz.org>

> > Acked-by: Kirill Korotaev <dev@openvz.org>

> >

> > ip\_tables.c | 36 ++++++-----

> > 1 file changed, 11 insertions(+), 25 deletions(-)

> >

> > ----

> > @@ -1618,7 +1603,7 @@ translate\_compat\_table(const char \*name,

> > unsigned int \*hook\_entries,

> > unsigned int \*underflows)

> > {

> > - unsigned int i;

> > + unsigned int i, j;

> > struct xt\_table\_info \*newinfo, \*info;

> > void \*pos, \*entry0, \*entry1;

> > unsigned int size;

> > @@ -1636,21 +1621,21 @@ translate\_compat\_table(const char \*name,

> > }

> >

> > duprintf("translate\_compat\_table: size %u\n", info->size);

> > - i = 0;

> > + j = 0;

> > xt\_compat\_lock(AF\_INET);

> > /\* Walk through entries, checking offsets. \*/

> > ret = IPT\_ENTRY\_ITERATE(entry0, total\_size,

> > check\_compat\_entry\_size\_and\_hooks,

> > info, &size, entry0,

> > entry0 + total\_size,

> > - hook\_entries, underflows, &i, name);

> > + hook\_entries, underflows, &j, name);

> > if (ret != 0)

> > goto out\_unlock;

> >

> > ret = -EINVAL;

```
>> - if (i != number) {
>> + if (j != number) {
>>   duprintf("translate_compat_table: %u not %u entries\n",
>> -   i, number);
>> +   j, number);
>>   goto out_unlock;
>> }
>>
>> @@ -1709,6 +1694,7 @@ translate_compat_table(const char *name,
>> free_newinfo:
>> xt_free_table_info(newinfo);
>> out:
>> + IPT_ENTRY_ITERATE(entry0, total_size, cleanup_entry, &j);
>> return ret;
>> out_unlock:
>> xt_compat_unlock(AF_INET);
>
> This doesn't look right, if we fail at xt_alloc_table_info for
> example the module references are not released.
module references will be released because of 'goto out;' line just after
xt_compat_unlock() - it is missed because of standard +-3 lines patch
context, sorry.
```

--  
Thanks,  
Dmitry.

---