
Subject: Re: [PATCH 1/2] iptables 32bit compat layer
Posted by [Andi Kleen](#) on Mon, 20 Feb 2006 21:23:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

Mishin Dmitry <dim@openvz.org> writes:

> Hello,
>
> This patch set extends current iptables compatibility layer in order to get
> 32bit iptables to work on 64bit kernel. Current layer is insufficient
> due to alignment checks both in kernel and user space tools.
>
> This patch introduces base compatibility interface for other ip_tables modules

Nice. But some issues with the implementation

```
+#if defined(CONFIG_X86_64)
+#define is_current_32bits() (current_thread_info()->flags & _TIF_IA32)
```

This should be `is_compat_task()`. And we don't do such `ifdefs` in generic code. And what you actually need here is a `is_compat_task_with_funny_u64_alignment()` (better name sought)

So I would suggest you add macros for that to the `ia64` and `x86-64` `asm/compat.h`s and perhaps a `ARCH_HAS_FUNNY_U64_ALIGNMENT` `#define` in there.

```
+ ret = 0;
+ switch (convert) {
+ case COMPAT_TO_USER:
+   pt = (struct ipt_entry_target *)target;
```

etc. that looks ugly. why can't you just define different functions for that? We don't really need in kernel `ioctl`

```
+#ifdef CONFIG_COMPAT
+ down(&compat_ipt_mutex);
+#endif
```

Why does it need an own lock?

Overall the implementation looks very complicated. Are you sure it wasn't possible to do this simpler?

-Andi
