

Re: [RFC,PATCH] loopback: calls netif_receive_skb() instead of netif_rx()

Re: [RFC,PATCH] loopback: calls netif_receive_skb() instead of netif_rx()

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-03/msg11927.html>

- *From:* David Miller <davem@xxxxxxxxxxxxxx>
 - *Date:* Mon, 31 Mar 2008 03:08:48 -0700 (PDT)
-

From: Ingo Molnar <mingo@xxxxxxx>
Date: Mon, 31 Mar 2008 11:48:23 +0200

* Eric Dumazet <dada1@xxxxxxxxxxxxxx> wrote:

I noticed some paths in kernel are very stack aggressive, and on i386 with CONFIG_4KSTACKS we were really in a dangerous land, even without my patch.

What we call 4K stacks is in fact 4K - sizeof(struct task_struct), so a little bit more than 2K. [...]

that's just wrong - 4K stacks on x86 are 4K - sizeof(thread_info) - the task struct is allocated elsewhere. The patch below runs just fine on 4K-stack x86.

I don't think it's safe.

Every packet you receive can result in a sent packet, which in turn can result in a full packet receive path being taken, and yet again another sent packet.

And so on and so forth.

Some cases like this would be stack bugs, but wouldn't you like that bug to be a very busy cpu instead of a crash from overrunning the current stack?

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxx
More majordomo info at <http://vger.kernel.org/majordomo-info.html>
Please read the FAQ at <http://www.tux.org/lkml/>

Re: [RFC,PATCH] loopback: calls netif_receive_skb() instead of netif_rx()