

Re: [BUG 2.6.20-rc2-mm1] init segfaults when CONFIG_PROFILE_LIKELY=y

Re: [BUG 2.6.20-rc2-mm1] init segfaults when CONFIG_PROFILE_LIKELY=y

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-12/msg08186.html>

- *From:* Daniel Walker <dwalker@xxxxxxxxxx>
 - *Date:* Sun, 31 Dec 2006 13:11:26 -0800
-

On Sun, 2006-12-31 at 12:43 -0800, Randy Dunlap wrote:

On Sun, 31 Dec 2006 11:45:09 -0800 Daniel Walker wrote:

On Sun, 2006-12-31 at 23:04 +0800, Fengguang Wu wrote:

Hi,

The following messages keeps popping up when CONFIG_PROFILE_LIKELY=y:

```
init[1]: segfault at ffffffff8118c110 rip ffffffff8118c110 rsp
00007fff9a9d14d8 error 15
init[1]: segfault at ffffffff8118c110 rip ffffffff8118c110 rsp
00007fff9a9d14d8 error 15
init[1]: segfault at ffffffff8118c110 rip ffffffff8118c110 rsp
00007fff9a9d14d8 error 15
init[1]: segfault at ffffffff8118c110 rip ffffffff8118c110 rsp
00007fff9a9d14d8 error 15
init[1]: segfault at ffffffff8118c110 rip ffffffff8118c110 rsp
00007fff9a9d14d8 error 15
init[1]: segfault at ffffffff8118c110 rip ffffffff8118c110 rsp
00007fff9a9d14d8 error 15
init[1]: segfault at ffffffff8118c110 rip ffffffff8118c110 rsp
00007fff9a9d14d8 error 15
init[1]: segfault at ffffffff8118c110 rip ffffffff8118c110 rsp
00007fff9a9d14d8 error 15
```

Does this seem like an appropriate solution? This just reconstitutes Ingo's patch by removing the unlikely calls that got added recently.

How does this fix the problem? (if it does)
What is the real cause of the problem?

Re: [BUG 2.6.20-rc2-mm1] init segfaults when CONFIG_PROFILE_LIKELY=y

Re: [BUG 2.6.20-rc2-mm1] init segfaults when CONFIG_PROFILE_LIKELY=y

Well I tested it so I sure hope it fixes it (unless I've gone mad). I guess we can wait for Fengguang to test it tho.

Maybe a comment into vsyscall.c that says to stay away from all macro's and possible debug code that could be added might be helpful ?

Why?

I don't know very much about vsyscalls, but from what I've read they actually reside in userspace. So with and "unlikely" added into that code, and profiling on, you will end up calling do_check_likely() which is in kernel space that's how the segfault happens.

I imagine this goes for all debugging in kernel space, you can't add it into a vsyscall. That's my reasoning behind adding a comment.

Daniel

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>