

## Re: BUG-lockdep and freeze (was: Arrr! Linux 2.6.18)

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-09/msg09141.html>

---

- *From:* Peter Zijlstra <[a.p.zijlstra@xxxxxxxxx](mailto:a.p.zijlstra@xxxxxxxxx)>
  - *Date:* Sat, 30 Sep 2006 21:49:14 +0200
- 

On Sat, 2006-09-30 at 21:20 +0200, Eric Rannaud wrote:

Hello,

On a 16-way Opteron (8 dual-core 880) with 8GB of RAM, vanilla 2.6.18 crashes early on boot with a BUG.

After many hours of git-bisecting, here is what I gathered.

This box had a history of oopses which were not fully investigated in the past, with older FC4 kernels. We're now doing a more complete analysis of the problems, and we tried running 2.6.18 on it. The whole memory was tested with memtest86+.

The first kernel known not to crash on boot was 2.6.15.4, and the first known to crash was 2.6.18.

Two directions were taken during the git-bisection:

- (1) BUG: message appears at some point (beginning of 2.6.18 cycle), but the kernel does not crash and seems to run fine (well enough to compile a kernel with -j 32). This one could be triggered by a different bug than in (2), but since the message is similar I thought it might be a good idea to look at its origin as well.
- (2) the kernel crashes very early on boot.

(traces and hardware info below, config on the web:

<http://engm.ath.cx/kernel/config-60be6b9a41cb0da0df7a9f11486da56baebf04cd>

<http://engm.ath.cx/kernel/config-d94a041519f3ab1ac023bf917619cd8c4a7d3c01>

<http://engm.ath.cx/kernel/config-2.6.18>

)

(1) is triggered by lockdep, and the BUG: is introduced by commit 60be6b9a41cb0da0df7a9f11486da56baebf04cd

[PATCH] lockdep: annotate on-stack completions, Signed-off-by: Ingo Molnar. Before that commit, and since its introduction in Linus' tree, lockdep was giving a trace and a warning ("INFO: trying to register non-static key. the code is fine but needs lockdep annotation. turning off the

Re: BUG-lockdep and freeze (was: Arrr! Linux 2.6.18)

locking correctness validator").

The BUG can be seen in every kernel I have tested between  
60be6b9a41cb0da0df7a9f11486da56baebf04cd and the first bad commit in

<http://kernel.org/pub/linux/kernel/people/akpm/patches/2.6/2.6.18/2.6.18-mm1/broken-out/slab-fix-lockdep-warning>  
<http://kernel.org/pub/linux/kernel/people/akpm/patches/2.6/2.6.18/2.6.18-mm1/broken-out/slab-fix-lockdep-warning>  
<http://kernel.org/pub/linux/kernel/people/akpm/patches/2.6/2.6.18/2.6.18-mm1/broken-out/slab-fix-lockdep-warning>

Those should rid you off the trace seen under:

```
----- console for (1) without numa=noacpi
Sep 30 15:54:06 liw64 kernel:
=====
Sep 30 15:54:06 liw64 kernel: [ INFO: possible recursive locking detected ]
Sep 30 15:54:06 liw64 kernel:
-----
Sep 30 15:54:06 liw64 kernel: swapper/1 is trying to acquire lock:
Sep 30 15:54:06 liw64 kernel: (&nc->lock){....}, at:
[<ffffffff8028a61a>] kmem_cache_free+0x15a/0x230
Sep 30 15:54:06 liw64 kernel:
Sep 30 15:54:06 liw64 kernel: but task is already holding lock:
Sep 30 15:54:06 liw64 kernel: (&nc->lock){....}, at:
[<ffffffff8028ad3a>] kfree+0x15a/0x240
Sep 30 15:54:06 liw64 kernel:
Sep 30 15:54:06 liw64 kernel: other info that might help us debug this:
Sep 30 15:54:06 liw64 kernel: 2 locks held by swapper/1:
Sep 30 15:54:06 liw64 kernel: #0: (&nc->lock){....}, at: [<ffffffff8028ad3a>]
kfree+0x15a/0x240
Sep 30 15:54:06 liw64 kernel: #1: (&parent->list_lock){....}, at:
[<ffffffff8028a995>] __drain_alien_cache+0x45/0xa0
Sep 30 15:54:06 liw64 kernel:
Sep 30 15:54:06 liw64 kernel: stack backtrace:
Sep 30 15:54:06 liw64 kernel:
Sep 30 15:54:06 liw64 kernel: Call Trace:
Sep 30 15:54:06 liw64 kernel: [<ffffffff8020b1ce>] show_trace+0xae/0x280
Sep 30 15:54:06 liw64 kernel: [<ffffffff8020b5e5>] dump_stack+0x15/0x20
Sep 30 15:54:06 liw64 kernel: [<ffffffff8024e462>] __lock_acquire+0x8f2/0xcf0
Sep 30 15:54:06 liw64 kernel: [<ffffffff8024ebeb>] lock_acquire+0x8b/0xc0
Sep 30 15:54:06 liw64 kernel: [<ffffffff8049c415>] _spin_lock+0x25/0x40
Sep 30 15:54:07 liw64 kernel: [<ffffffff8028a61a>] kmem_cache_free+0x15a/0x230
Sep 30 15:54:07 liw64 kernel: [<ffffffff8028a7ab>] slab_destroy+0xbb/0xf0
Sep 30 15:54:07 liw64 kernel: [<ffffffff8028a8f1>] free_block+0x111/0x170
Sep 30 15:54:07 liw64 kernel: [<ffffffff8028a9be>]
__drain_alien_cache+0x6e/0xa0
Sep 30 15:54:07 liw64 kernel: [<ffffffff8028ad4f>] kfree+0x16f/0x240
Sep 30 15:54:07 liw64 kernel: [<ffffffff8094d179>] free+0x9/0x10
Sep 30 15:54:07 liw64 kernel: [<ffffffff8094d63e>] huft_free+0x1e/0x30
Sep 30 15:54:07 liw64 kernel: [<ffffffff8094e808>] inflate_dynamic+0x4d8/0x610
Sep 30 15:54:07 liw64 kernel: [<ffffffff8094ee3d>]
```

Re: BUG-lockdep and freeze (was: Arrr! Linux 2.6.18)

```
unpack_to_rootfs+0x4ed/0x9c0Sep 30 15:54:07 liw64 kernel:
[<ffffff8094f3a9>] populate_rootfs+0x69/0x100
Sep 30 15:54:07 liw64 kernel: [<ffffff80207139>] init+0xd9/0x350
Sep 30 15:54:07 liw64 kernel: [<ffffff8020aa9e>] child_rip+0x8/0x12
Sep 30 15:54:07 liw64 kernel: it is
-----
```

—  
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/>