

## [PATCH] Fix kprobes handling of simultaneous probe hit/unregister

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-08/8053.html>

---

*From:* Jim Keniston ([jkenisto\\_at\\_us.ibm.com](mailto:jkenisto_at_us.ibm.com))

*Date:* 08/31/05

To: LKML <[linux-kernel@vger.kernel.org](mailto:linux-kernel@vger.kernel.org)>, Andrew Morton <[akpm@osdl.org](mailto:akpm@osdl.org)>

Date: 31 Aug 2005 14:53:37 -0700

This patch fixes a bug in kprobes's handling of a corner case on i386 and x86\_64. On an SMP system, if one CPU unregisters a kprobe just after another CPU hits that probepoint, `kprobe_handler()` on the latter CPU sees that the kprobe has been unregistered, and attempts to let the CPU continue as if the probepoint hadn't been hit. The bug is that on i386 and x86\_64, we were neglecting to set the IP back to the beginning of the probed instruction. This could cause an oops or crash.

This bug doesn't exist on ppc64 and ia64, where a breakpoint instruction leaves the IP pointing to the beginning of the instruction. I don't know about sparc64. (Dave, could you please advise?)

This fix has been tested on i386 and x86\_64 SMP systems. To reproduce the problem, set one CPU to work registering and unregistering a kprobe repeatedly, and another CPU pounding the probepoint in a tight loop.

Please apply.

Acked-by: Prasanna S Panchamukhi <[prasanna@in.ibm.com](mailto:prasanna@in.ibm.com)>

Signed-off-by: Jim Keniston <[jkenisto@us.ibm.com](mailto:jkenisto@us.ibm.com)>

```
--- linux-2.6.13/arch/i386/kernel/kprobes.c 2005-08-30 12:27:35.000000000 -0700
+++ linux-fixed/arch/i386/kernel/kprobes.c 2005-08-30 15:33:03.000000000 -0700
@@ -220,7 +220,10 @@
     * either a probepoint or a debugger breakpoint
     * at this address. In either case, no further
     * handling of this interrupt is appropriate.
+ * Back up over the (now missing) int3 and run
+ * the original instruction.
     */
+ regs->eip -= sizeof(kprobe_opcode_t);
     ret = 1;
```

Linux-Kernel: [PATCH] Fix kprobes handling of simultaneous probe hit/unregister

```
    }
    /* Not one of ours: let kernel handle it */
--- linux-2.6.13/arch/x86_64/kernel/kprobes.c 2005-08-30 12:27:35.000000000 -0700
+++ linux-fixed/arch/x86_64/kernel/kprobes.c 2005-08-30 15:32:31.000000000 -0700
@@ -360,7 +360,10 @@
        * either a probepoint or a debugger breakpoint
        * at this address. In either case, no further
        * handling of this interrupt is appropriate.
+ * Back up over the (now missing) int3 and run
+ * the original instruction.
        */
+ regs->rip = (unsigned long)addr;
        ret = 1;
    }
    /* Not one of ours: let kernel handle it */
```

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