

[PATCH 4/4] stack overflow safe kdump (2.6.15-rc3-i386) – fault

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-11/9665.html>

From: Fernando Luis Vazquez Cao (*fernando_at_intellilink.co.jp*)

Date: 11/30/05

To: "Eric W. Biederman" <ebiederm@xmission.com>

Date: Wed, 30 Nov 2005 16:37:47 +0900

When we have a bloated stack it is likely that it ends up making an invalid memory access that causes a page fault. Take this case into account in the page fault code.

```
diff -urNp linux-2.6.15-rc3/arch/i386/mm/fault.c linux-2.6.15-rc3-sov/arch/i386/mm/fault.c
--- linux-2.6.15-rc3/arch/i386/mm/fault.c      2005-11-30 14:51:49.000000000 +0900
+++ linux-2.6.15-rc3-sov/arch/i386/mm/fault.c  2005-11-30 14:56:04.000000000 +0900
@@ -245,6 +245,11 @@ fastcall void __kprobes do_page_fault(st
         local_irq_enable();
```

```

        tsk = current;
+       /* We may have invalid 'current' due to a stack overflow. */
+       if (!virt_addr_valid(tsk)) {
+           printk("do_page_fault: Discarding invalid 'current' struct task_struct * = 0x%p\n",
+                  tsk = NULL;
+       }
+   }
```

```
        si_code = SEGV_MAPERR;
```

```
@@ -271,7 +276,14 @@ fastcall void __kprobes do_page_fault(st
        goto bad_area_nosemaphore;
    }
}
```

```
-       mm = tsk->mm;
+       mm = NULL;
+       /* We may have invalid 'tsk' due to a i386 stack overflow */
+       if (tsk)
+           mm = tsk->mm;
+       if (mm && !virt_addr_valid(mm)) {
+           printk("do_page_fault: Discarding invalid current->mm struct mm_struct * = 0x%p\n",
+                  mm = NULL;
+       }
+   }
```

```
/*
 * If we're in an interrupt, have no user context or are running in an
-   */
-   */
```

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/>