

Linux-Kernel: Re: [patch 1/1] Ptrace – i386: fix "syscall audit" interaction with singlestep

```
> --- linux-2.6.git/arch/i386/kernel/ptrace.c~sysaudit-singlestep-non-umlhost 2005-07-26
20:22:40.000000000 +0200
> +++ linux-2.6.git-paolo/arch/i386/kernel/ptrace.c 2005-07-26 20:23:44.000000000 +0200
> @@ -683,8 +683,19 @@ void do_syscall_trace(struct pt_regs *re
> /* do the secure computing check first */
> secure_computing(regs->orig_eax);
>
> - if (unlikely(current->audit_context) && entryexit)
> - audit_syscall_exit(current, AUDITSC_RESULT(regs->eax), regs->eax);
> + if (unlikely(current->audit_context)) {
> + if (entryexit)
> + audit_syscall_exit(current, AUDITSC_RESULT(regs->eax), regs->eax);
> +
> + /* Debug traps, when using PTRACE_SINGLESTEP, must be sent only
> + * on the syscall exit path. Normally, when TIF_SYSCALL_AUDIT is
> + * not used, entry.S will call us only on syscall exit, not
> + * entry ; so when TIF_SYSCALL_AUDIT is used we must avoid
> + * calling send_sigtrap() on syscall entry.
> + */
> + else if (is_singlestep)
> + goto out;
> + }
>
```

This appears to be a UML patch, applied to x86, which has no `is_singlestep`.

-

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/>