

# [patch 09/12] syslets: x86, mark async unsafe syscalls

---

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-02/msg10245.html>

---

- From: Ingo Molnar <mingo@xxxxxxx>
  - Date: Wed, 28 Feb 2007 22:42:07 +0100
- 

From: Ingo Molnar <mingo@xxxxxxx>

mark clone() and fork() as not available for async execution.  
Both need an intact user context beneath them to work.

Signed-off-by: Ingo Molnar <mingo@xxxxxxx>  
Signed-off-by: Arjan van de Ven <arjan@xxxxxxxxxxxxxxxxxx>

-----  
arch/i386/kernel/ioport.c | 6 ++++++  
arch/i386/kernel/ldt.c | 3 +++  
arch/i386/kernel/process.c | 6 ++++++  
arch/i386/kernel/vm86.c | 6 ++++++  
4 files changed, 21 insertions(+)

Index: linux/arch/i386/kernel/ioport.c

```
=====
--- linux.orig/arch/i386/kernel/ioport.c
+++ linux/arch/i386/kernel/ioport.c
@@ -62,6 +62,9 @@ asmlinkage long sys_ioperm(unsigned long
struct tss_struct * tss;
unsigned long *bitmap;

+ if (async_syscall(current))
+ return -ENOSYS;
+
if ((from + num <= from) || (from + num > IO_BITMAP_BITS))
return -EINVAL;
if (turn_on && !capable(CAP_SYS_RAWIO))
@@ -139,6 +142,9 @@ asmlinkage long sys_iopl(unsigned long u
unsigned int old = (regs->eflags >> 12) & 3;
struct thread_struct *t = &current->thread;

+ if (async_syscall(current))
+ return -ENOSYS;
+
if (level > 3)
return -EINVAL;
/* Trying to gain more privileges? */
```

Index: linux/arch/i386/kernel/ldt.c

```
----- linux.orig/arch/i386/kernel/ldt.c
+++ linux/arch/i386/kernel/ldt.c
@@ -233,6 +233,9 @@ asmlinkage int sys_modify_ldt(int func,
{
int ret = -ENOSYS;

+ if (async_syscall(current))
+ return -ENOSYS;
+
switch (func) {
case 0:
ret = read_ldt(ptr, bytecount);
```

Index: linux/arch/i386/kernel/process.c

```
----- linux.orig/arch/i386/kernel/process.c
+++ linux/arch/i386/kernel/process.c
@@ -750,6 +750,9 @@ struct task_struct fastcall * __switch_t

asmlinkage int sys_fork(struct pt_regs regs)
{
+ if (async_syscall(current))
+ return -ENOSYS;
+
return do_fork(SIGCHLD, regs.esp, &regs, 0, NULL, NULL);
}
```

```
@@ -759,6 +762,9 @@ asmlinkage int sys_clone(struct pt_regs
unsigned long newsp;
int __user *parent_tidptr, *child_tidptr;

+ if (async_syscall(current))
+ return -ENOSYS;
+
clone_flags = regs.ebx;
newsp = regs.ecx;
parent_tidptr = (int __user *)regs.edx;
Index: linux/arch/i386/kernel/vm86.c
```

```
----- linux.orig/arch/i386/kernel/vm86.c
+++ linux/arch/i386/kernel/vm86.c
@@ -209,6 +209,9 @@ asmlinkage int sys_vm86old(struct pt_reg
struct task_struct *tsk;
int tmp, ret = -EPERM;

+ if (async_syscall(current))
+ return -ENOSYS;
+
tsk = current;
if (tsk->thread.saved_esp0)
```

[patch 09/12] syslets: x86, mark async unsafe syscalls

```
goto out;
@@ -239,6 +242,9 @@ asmlinkage int sys_vm86(struct pt_regs r
int tmp, ret;
struct vm86plus_struct __user *v86;
```

```
+ if (async_syscall(current))
+ return -ENOSYS;
+
tsk = current;
switch (regs.ebx) {
case VM86_REQUEST_IRQ:
```

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