

Re: Syscall/Sysret state on i386 arch

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/hackers/2005-08/0482.html>

From: John Baldwin (jhb_at_FreeBSD.org)

Date: 08/29/05

To: freebsd-hackers@freebsd.org

Date: Mon, 29 Aug 2005 11:36:14 -0400

On Sunday 28 August 2005 10:32 am, alexander wrote:

- > *The AMD64 arch is using the syscall/sysret opcodes instead of int80h to*
- > *perform a syscall (/usr/src/lib/libc/amd64/SYS.h). I just checked the*
- > *output my of dmesg and it says:*
- >
- > *CPU: AMD Duron(tm) Processor (1311.69-MHz 686-class CPU)*
- > *Origin = "AuthenticAMD" Id = 0x671 Stepping = 1*
- >
- > *Features=0x383f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV*
- > *,\ PAT,PSE36,MMX,FXSR,SSE>*
- > *AMD Features=0xc0400800<SYSCALL,MMX+,3DNow+,3DNow>*
- >
- > *I got a hold of the AMD document number 21086.pdf. It describes both*
- > *opcodes pretty well, but doesn't tell which CPUs support the new opcodes.*
- > *But since the first revision of that document is dated Sept 1997 quite a*
- > *lot of i386 CPU's should support the opcodes. The NASM manual only states*
- > *[P6,AMD] as the required CPU to perform those opcodes.*
- >
- > *I found some patches for Linux that replace the int80h syscall calling*
- >
- > *convention with syscall/sysret on i386 and the results look pretty*
- > *convincing:*
- > > *(INT \$0x80 based getpid(), got pid 497) latency:282 cycles*
- > > *(SYSENTER based getpid(), got pid 497) latency:138 cycles*
- > >
- > > *on a 266 MHz PII this is 0.51 usecs for a getpid(). (was 1.06 usecs)*
- >
- > *Quoted from: <http://www.ussg.iu.edu/hypermail/linux/kernel/9806.1/0878.html>*
- >
- > *Does anybody know more about this? Is it even possible to replace the*
- > *current syscall implementation that easily or would that require elaborate*
- > *changes to all the syscalls (libc), etc. And which CPU's support these new*
- > *opcodes? Doesn't anybody know if the Linux patches actually got committed to*
- > *the official kernel?*

Support for syscall/sysret is determined by a cpuid flag. I do believe someone has worked on either syscall/sysret or sysenter/sysexit support in a

freebsd-hackers: Re: Syscall/Sysret state on i386 arch

p4 branch. You can try asking jeff@ about it. I think it was
sysenter/sysexit and it didn't really improve things much.

--

John Baldwin <jhb@FreeBSD.org> <>< <http://www.FreeBSD.org/~jhb/>
"Power Users Use the Power to Serve" = <http://www.FreeBSD.org>

freebsd-hackers@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-hackers>

To unsubscribe, send any mail to "freebsd-hackers-unsubscribe@freebsd.org"