

Re: CPU boot problem on 2.6.0-test3-bk8

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2003-08/5767.html>

From: Dave Hansen (haveblue_at_us.ibm.com)

Date: 08/21/03

To: Andrew Theurer <habanero@us.ibm.com>

Date: 20 Aug 2003 20:42:09 -0700

On Wed, 2003-08-20 at 18:13, Andrew Theurer wrote:

> On Wednesday 20 August 2003 20:02, Dave Hansen wrote:
> > On Wed, 2003-08-20 at 14:58, Andrew Theurer wrote:
> > > Maybe this is already known, but just in case:
> > > I cannot fully boot on an x440 system with 2.6.0-test3-bk8. The kernel
> > > tries to boot more than the 16 logical processors, and after failing (no
> > > response) on cpus 16, 17, and 18, it still thinks it has 19 cpus total.
> > > It finally gets stuck at "checking TSC synchronization across 19 CPUs:"
> > >
> > > Attached is the boot log. Any ideas? I'll try -test3-bk7 next
> > >
> > Can you see if it works without HT on? Did it work on plain -test3?
> > My 16-way x440 with no HT boots fine on test3.
> >
> I'll try without HT to see what happens. FWIW, it boots fine with HT if I set
> maxcpus=16. I am wondering if (apicid == BAD_APIC) test is not working in
> smp_boot_cpus.

Hmmm. This is looking like fallout from the massive wli-bomb. Here's the loop that controls the cpu booting, before and after cpumask_t:

```
- for (bit = 0; kicked < NR_CPUS && bit < BITS_PER_LONG; bit++) + for  
(bit = 0; kicked < NR_CPUS && bit < MAX_APICS; bit++)  
    apicid = cpu_present_to_apicid(bit);
```

"kicked" only gets incremented for CPUs that were successfully booted, so it doesn't help terminate the loop much. MAX_APICS is 256 on summit, which is *MUCH* bigger than BITS_PER_LONG.

cpu_2_logical_apicid[NR_CPUS] which is referenced from cpu_present_to_apicid() is getting referenced up to MAX_APICs, which is bigger than NR_CPUS. Overflow. Bang. garbage != BAD_APICID :)

Attached patch fixes it. We sure do have a lot of duplicate code in the subarches. <sigh>

--

Dave Hansen

Re: CPU boot problem on 2.6.0-test3-bk8

Linux-Kernel: Re: CPU boot problem on 2.6.0-test3-bk8

haveblue@us.ibm.com

—

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/>

- text/x-patch attachment:
[cpu to logical apicid-fix-2.6.0-test3-bk8-0.patch](#) charset ANSI_X3.4-1968