

[PATCH][2.6.8.1-mm4] phys_proc_id change for x86-64

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-09/2873.html>

From: James Cleverdon (jamesclv_at_us.ibm.com)

Date: 09/09/04

To: linux-kernel@vger.kernel.org, Andi Kleen <ak@suse.de>

Date: Wed, 8 Sep 2004 17:23:16 -0700

The value that `cpuinfo` returns for command 1 in `ebx` is the physical APIC ID latched when the system comes out of reset.

Ordinarily, this is identical to the value in the local APIC's ID register, because nearly all BIOSes accept the HW assigned value.

Our systems, made up of individual building blocks, can't do that. Each node boots as a separate system and is joined together by the BIOS. Thus, the BIOS rewrites the local APIC ID register with a new value.

Potomac and Nocona chips have a mechanism by which the BIOS writer can change bits 7:5 to match the assigned cluster ID.

Bits 2:0 come from the thread ID. However, bits 4:3 are still those latched at reset. Oops!

Summary: Large clustered systems can't use `cpuid` to derive the sibling information.

Fix: Use the local APIC ID. That's the value we use to online the CPUs, so it had better be OK. For non-clustered systems, `cpuid == local APIC`, so nothing but large boxes should be affected.

```
diff -pruN 2.6.8.1-mm4/arch/x86_64/kernel/setup.c w8.1m4/arch/x86_64/kernel/setup.c
--- 2.6.8.1-mm4/arch/x86_64/kernel/setup.c 2004-08-25 14:48:21.000000000 -0700
+++ w8.1m4/arch/x86_64/kernel/setup.c 2004-09-02 15:51:55.000000000 -0700
@@ -724,7 +724,7 @@ static void __init detect_ht(struct cpui
     }
     if (index_lsb != index_msb )
         index_msb++;
- initial_apic_id = ebx >> 24 & 0xff;
+ initial_apic_id = hard_smp_processor_id();
     phys_proc_id[cpu] = initial_apic_id >> index_msb;
```

Linux-Kernel: [PATCH][2.6.8.1-mm4] phys_proc_id change for x86-64

```
printk(KERN_INFO "CPU: Physical Processor ID: %d\n",
```

--

James Cleverdon

IBM LTC (xSeries Linux Solutions)

{jamesclv(Unix, preferred), cleverdj(Notes)} at us dot ibm dot comm

-

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-diff attachment: [cpu_sibling_map_2004-08-31_2.6.8.1-mm4](#)