
Subject: [PATCH] p4-clockmod: switch to rdmsr_on_cpu/wrmsr_on_cpu
Posted by [Alexey Dobriyan](#) on Mon, 19 Mar 2007 16:10:09 GMT

[View Forum Message](#) <> [Reply to Message](#)

Dances with cpumasks go away.

Signed-off-by: Alexey Dobriyan <adobriyan@sw.ru>

arch/i386/kernel/cpu/cpufreq/p4-clockmod.c | 31 +++++-----
1 file changed, 6 insertions(+), 25 deletions(-)

```
--- a/arch/i386/kernel/cpu/cpufreq/p4-clockmod.c
+++ b/arch/i386/kernel/cpu/cpufreq/p4-clockmod.c
@@ -27,7 +27,6 @@ #include <linux/smp.h>
#include <linux/cpufreq.h>
#include <linux/slab.h>
#include <linux/cpumask.h>
-#include <linux/sched.h> /* current / set_cpus_allowed() */

#include <asm/processor.h>
#include <asm/msr.h>
@@ -62,7 +61,7 @@ static int cpufreq_p4_setdc(unsigned int
if (!cpu_online(cpu) || (newstate > DC_DISABLE) || (newstate == DC_RESV))
return -EINVAL;

- rdmsr(MSR_IA32_THERM_STATUS, l, h);
+ rdmsr_on_cpu(cpu, MSR_IA32_THERM_STATUS, &l, &h);

if (l & 0x01)
dprintk("CPU #%d currently thermal throttled\n", cpu);
@@ -70,10 +69,10 @@ static int cpufreq_p4_setdc(unsigned int
if (has_N44_O17_errata[cpu] && (newstate == DC_25PT || newstate == DC_DFLT))
newstate = DC_38PT;

- rdmsr(MSR_IA32_THERM_CONTROL, l, h);
+ rdmsr_on_cpu(cpu, MSR_IA32_THERM_CONTROL, &l, &h);
if (newstate == DC_DISABLE) {
dprintk("CPU #%d disabling modulation\n", cpu);
- wrmsr(MSR_IA32_THERM_CONTROL, l & ~(1<<4), h);
+ wrmsr_on_cpu(cpu, MSR_IA32_THERM_CONTROL, l & ~(1<<4), h);
} else {
dprintk("CPU #%d setting duty cycle to %d%%\n",
cpu, ((125 * newstate) / 10));
@@ -84,7 +83,7 @@ static int cpufreq_p4_setdc(unsigned int
*/
l = (l & ~14);
l = l | (1<<4) | ((newstate & 0x7)<<1);
```

```

- wrmsr(MSR_IA32_THERM_CONTROL, l, h);
+ wrmsr_on_cpu(cpu, MSR_IA32_THERM_CONTROL, l, h);
}

return 0;
@@ -111,7 +110,6 @@ static int cpufreq_p4_target(struct cpuf
{
    unsigned int newstate = DC_RESV;
    struct cpufreq_freqs freqs;
- cpumask_t cpus_allowed;
    int i;

    if (cpufreq_frequency_table_target(policy, &p4clockmod_table[0], target_freq, relation,
&newstate))
@@ -132,17 +130,8 @@ static int cpufreq_p4_target(struct cpuf
/* run on each logical CPU, see section 13.15.3 of IA32 Intel Architecture Software
 * Developer's Manual, Volume 3
 */
- cpus_allowed = current->cpus_allowed;
-
- for_each_cpu_mask(i, policy->cpus) {
- cpumask_t this_cpu = cpumask_of_cpu(i);
-
- set_cpus_allowed(current, this_cpu);
- BUG_ON(smp_processor_id() != i);
-
+ for_each_cpu_mask(i, policy->cpus)
    cpufreq_p4_setdc(i, p4clockmod_table[newstate].index);
- }
- set_cpus_allowed(current, cpus_allowed);

/* notifiers */
for_each_cpu_mask(i, policy->cpus) {
@@ -256,17 +245,9 @@ static int cpufreq_p4_cpu_exit(struct cp

static unsigned int cpufreq_p4_get(unsigned int cpu)
{
- cpumask_t cpus_allowed;
    u32 l, h;

- cpus_allowed = current->cpus_allowed;
-
- set_cpus_allowed(current, cpumask_of_cpu(cpu));
- BUG_ON(smp_processor_id() != cpu);
-
- rdmsr(MSR_IA32_THERM_CONTROL, l, h);
-
- set_cpus_allowed(current, cpus_allowed);

```

```
+ rdmsr_on_cpu(cpu, MSR_IA32_THERM_CONTROL, &l, &h);

if (l & 0x10) {
    l = l >> 1;
```
