

---

Subject: Re: [PATCH 1/3] Introduce cpuid\_on\_cpu() and cpuid\_eax\_on\_cpu()  
Posted by [hpa](#) on Mon, 02 Apr 2007 16:20:16 GMT

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

---

Andi Kleen wrote:

> On Monday 02 April 2007 13:38, Alexey Dobriyan wrote:

>> They will be used by cpuid driver and powernow-k8 cpufreq driver.

>>

>> With these changes powernow-k8 driver could run correctly on OpenVZ kernels

>> with virtual cpus enabled (SCHED\_VCPU).

>

> This means openvz has multiple virtual CPU levels? One for cpuid/rdmsr and one

> for the rest of the kernel? Both powernow-k8 and cpuid attempt to schedule

> to the target CPU so they should already run there. But it is some other CPU,

> but when they ask your \_on\_cpu() functions they suddenly get a "real" CPU?

> Where is the difference between these levels of virtualness?

>

The CPUID and MSR drivers do not schedule to the target CPU; instead, on hardware, they rely on IPI'ing the target processor if it is not the one that's currently running.

There were a lot of discussion back when about which was the better solution. Alan Cox, in particular, really preferred the interrupt solution as being less likely to cause implicit deadlock.

I do want to add that it's been on my list for some time -- in fact, I keep implementing it half-way and then having other things to do -- to add MSR and CPUID ioctls() that allow the full register file to be set and read back, in order to support architecturally broken MSR and CPUID levels.

-hpa

---