Subject: Re: [PATCH 1/3] Introduce cpuid_on_cpu() and cpuid_eax_on_cpu()
Posted by Andi Kleen on Tue, 03 Apr 2007 13:42:50 GMT
View Forum Message <> Reply to Message

> > 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?
>
> *_on_cpu functions do some work on given physical CPU.
> set_cpus_allowed() in openvz operates on VCPU level, so process doing
> set_cpus_allowed() still could be scheduled anywhere.

Ok so you have multple levels.

> > Also it has weird semantics. For example if you have multiple
> > virtual CPUs mapping to a single CPU then would the powernow-k8 driver
> > try to set the frequency multiple times on the same physical CPU?
>
> If core cpufreq locking is OK, why would it?

It won't know about multiple CPUs mapping to a single CPU.

> apply_microcode() looks small enough to convert it to IPIs, but so far
> nobody asked for microcode updates in openvz.

Well if they try it they will probably have problems.

> > Before adding any hacks like this I think your vcpu concept
> > needs to be discussed properly on l-k. For me it doesn't look like it is
> > something good right now though.
>
> Andi, I think it all relies on correctness of core cpufreq locking.

I have my doubts it will cope with you changing all reasonable expected semantics
under it.

-Andi