

Re: [PATCH 3/6] sched: high-res preemption tick

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-10/msg12335.html>

- *From:* Andi Kleen <andi@xxxxxxxxxxxxxx>
 - *Date:* Wed, 31 Oct 2007 22:53:04 +0100
-

Peter Zijlstra <a.p.zijlstra@xxxxxxxxxx> writes:

Use HR-timers (when available) to deliver an accurate preemption tick.

The regular scheduler tick that runs at 1/HZ can be too coarse when nice level are used. The fairness system will still keep the cpu utilisation 'fair' by then delaying the task that got an excessive amount of CPU time but try to minimize this by delivering preemption points spot-on.

This might be costly when hrtimers happen to use an more expensive to reprogram time source. Even an APIC timer access is fairly slow. And you'll potentially add the to lots of context switces.

Not sure that is a good idea for performance in general.

-Andi

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>