

## Re: [Lse-tech] [patch] sched-domain cleanups, sched-2.6.5-rc2-mm2-A3

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-03/7395.html>

---

**From:** Ingo Molnar ([mingo\\_at\\_elte.hu](mailto:mingo_at_elte.hu))

**Date:** 03/30/04

Date: Tue, 30 Mar 2004 10:45:01 +0200  
To: Nick Piggin <[nickpiggin@yahoo.com.au](mailto:nickpiggin@yahoo.com.au)>

\* Nick Piggin <[nickpiggin@yahoo.com.au](mailto:nickpiggin@yahoo.com.au)> wrote:

> *You're probably mostly right, but I really don't know if I'd start  
> with the assumption that threads don't share anything. I think they're  
> very likely to share memory and cache.*

it all depends on the workload i guess, but generally if the application scales well then the threads only share data in a read-mostly manner – hence we can balance at creation time.

if the application does not scale well then balancing too early cannot make the app perform much worse.

things like JVMs tend to want good balancing – they really are userspace simulations of separate contexts with little sharing and good overall scalability of the architecture.

> *Also, these additional system wide balance points don't come for free  
> if you attach them to common operations (as opposed to the slow  
> periodic balancing).*

yes, definitely.

the implementation in sched2.patch does not take this into account yet. There are a number of things we can do about the 500 CPUs case. Eg. only do the balance search towards the next N nodes/cpus (tunable via a domain parameter).

Ingo

–

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to [majordomo@vger.kernel.org](mailto: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/>