

## Re: [RFC PATCH] sched\_domains: Make SD\_NODE\_INIT per-arch

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-09/9130.html>

---

**From:** Matthew Dobson ([colpatch\\_at\\_us.ibm.com](mailto:colpatch_at_us.ibm.com))

**Date:** 09/30/04

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

Date: Thu, 30 Sep 2004 11:36:52 -0700

On Thu, 2004-09-30 at 01:15, Nick Piggin wrote:

> *Matthew Dobson wrote:*

> > IA64 already has their own version of SD\_NODE\_INIT, tuned for their  
> > extremely large machines. I think that all arches would benefit from  
> > having their own, arch-specific SD\_NODE\_INIT initializer, rather than  
> > the one-size-fits-all variant we've got now.

> >

>

> I suppose the patch is pretty good (IIRC Martin liked the idea).

> I guess it will at least increase the incidence of copy+paste,

> if not getting people to think harder ;)

Thanks! Martin does like the idea, and I think Andi Kleen likes the idea of being able to tune sched\_domains for x86\_64, too. Any comments, Andi?

The patch is pretty simple. I don't think it will increase any copy+pasting because I don't believe anyone has modified SD\_NODE\_INIT at all since it's been implemented, and certainly not for many kernel releases. I think part of the reason for that is that it is currently impossible to tweak the values for your architecture of choice because modifying the values now will change EVERYONE's sched\_domains timings. Which is bad. :( If anyone wants to tweak SD\_NODE\_INIT, they shouldn't be copying+pasting those values to all architectures. Besides, IA64 already gets their own SD\_NODE\_INIT to play with, why shouldn't everyone else! ;)

> *Can I be lame and ask that you keep this around until closer*

> *to 2.6.10? I have a few possible scheduler performance*

> *improvements that I'd like to get tested in -mm after 2.6.9*

> *and this would make things a bit harder :P*

>

> *I don't think anyone is looking at getting any tweaks in before*

> *then...*

Linux-Kernel: Re: [RFC PATCH] sched\_domains: Make SD\_NODE\_INIT per-arch

I would like to try to get this in before then, unless this will really make things difficult for you. 2.6.9 is looking to be a pickup point for distros, so getting this patch in now (pre 2.6.9) means that distros can add tiny patches to their builds to simply tweak individual architecture values without having to diverge from mainline by implementing this patch on their own. It also means that any tuning work that the distros do can easily be pushed back to mainline by simply sending a patch per architecture with new values. It makes those patches safe and minimizes conflicts.

Will this patch really make your scheduler improvements that much harder to test/implement? I don't think it should, unless your improvements are tweaking the values in SD\_NODE\_INIT, since that is all this touches... Even after this patch, SD\_NODE\_INIT is still picked up in include/linux/sched.h, so the changes required to cope with this patch should be minimal...

–Matt

–

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/>