

Re: [RFC PATCH] sched_domains: Make SD_NODE_INIT per-arch

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-09/9182.html>

From: Andi Kleen (ak_at_suse.de)

Date: 09/30/04

Date: Thu, 30 Sep 2004 22:45:02 +0200
To: Matthew Dobson <colpatch@us.ibm.com>

On Thu, Sep 30, 2004 at 11:36:52AM -0700, Matthew Dobson wrote:

> *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?*

It doesn't help me directly – what i need is the same thing
for SD_SIBLING_INIT for the CMP changes.

But it seems I need to do some other work to properly support the K8
CMP first, so I'm defering attacking this a bit.

> *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! ;)*

It would be nice if there was a SD_DEFAULT_NODE_INIT and a

Linux-Kernel: Re: [RFC PATCH] sched_domains: Make SD_NODE_INIT per-arch

SD_DEFAULT_SIBLING_INIT in some generic file that architecture code can use as a base for tweaking. For the CMP change I currently only want to remove SD_SHAREPOWER from SIBLING_INIT to get rid of SMT nice.

Later we'll probably want a SD_DEFAULT_CMP_INIT too that gives generic values for a dual core. Dual cores should be soon pretty common and tuning for them will be needed on several architectures (ppc64, ia64, x86, x86-64, sparc, parisc? ...). But figuring out good values for this will require a lot of benchmarking first.

-Andi

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to 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/>