

Re: kswapd in tight loop 2.6.9-rc3-bk-recent

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-10/2077.html>

From: Nick Piggin (nickpiggin_at_yahoo.com.au)

Date: 10/08/04

Date: Fri, 08 Oct 2004 10:34:47 +1000

To: Andrew Morton <akpm@osdl.org>

Andrew Morton wrote:

>

>*This probably won't fix it.*

>

>*It looks like the code will lock up if all zones are out of unreclaimable
>memory, but you won't be hitting that.*

>

>

Out of `_reclaimable_` memory?

It shouldn't if `all_unreclaimable` is being set correctly.

>*I also wonder if it'll lock up if just the first zone has `->all_unreclaimable`.*

>

>

Well, not if the `all_unreclaimable` flag is set, but if it should be and isn't, then probably it will lock up.

>*I think a good starting point here will be to revert the most recent
>change.*

>

That may fix it for the simple fact that kswapd will just go through its priority loop once then stop.

I think that resetting `all_unreclaimable` in `free_pages_bulk` is the wrong idea though, because that will keep it clear if a bit of kernel memory is being pinned and freed in the background, won't it?

I had a look and decided that `all_unreclaimable` should probably be cleared only if `vmscan.c` frees some memory – but I couldn't really come up with any hard numbers to back me up :P

Linux-Kernel: Re: kswapd in tight loop 2.6.9-rc3-bk-recent

—

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/>