

Re: [PATCH 2/2] cciss: disable dma prefetch for P600

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-10/msg07314.html>

- *From:* Andrew Morton <akpm@xxxxxxxx>
 - *Date:* Fri, 20 Oct 2006 13:23:50 -0700
-

On Fri, 20 Oct 2006 14:56:18 -0500

"Mike Miller (OS Dev)" <mikem@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote:

..

What is the actual problem which is being fixed here?

Sorry, I should have been clearer. There is a bug in the DMA engine that that may result in prefetching data from beyond the end of memory or falling off into one the holes on IPF and AMD. It causes a machine check when that happens. It doesn't happen on Proliant because the last 4kB (or so) of memory is mapped out by the BIOS and Pentium guarantees contiguous memory.

I think that this:

```
#if defined(CONFIG_IA64) ||  
defined(CONFIG_X86_64)
```

is nowhere near strong enough and is probably inappropriate.

It *could* be that CONFIG_DISCONTIGMEM|CONFIG_SPARSEMEM will be closer, but even CONFIG_FLATMEM systems can have holes.

Re: [PATCH 2/2] cciss: disable dma prefetch for P600

I'm poking around on some IPF platforms. It looks like CONFIG_DISCONTIGMEM is set on them, but not the others you mention. Would that be sufficient?

I don't think so. All machines in all memory models can and do have holes in their memory map. I think the problem is that some machines object to having those holes read from and others do not. It could be that this problem is purely an ia64 thing.

And it's not just holes: we had a problem a year or so back where CPU prefetching was walking off the end of real mmeory and into the AGP region and was causing weird cache coherency problems on x86_64 (or something like that).

On what machines can/does this card exist? Things like powerpc?

This problem was found on Itanium. We don't try to support powerpc.

Well the CCISS driver presently has no architecture Kconfig dependencies, so anyone can build it on anything. I don't know whether it's physically possible to put a cciss controller into a power/sparc/whatever machine – are these controllers only ever integrated onto the main board?

Anyway, I'd suggest the best way of sorting this out is to come up with a complete description of the problem, decide which architectures are affected and to then ask the relevant architecture maintainers to recommend a solution.

I think the description would be

There is a bug in the DMA engine that that may result in prefetching data from beyond the end of memory or falling off into one the holes on IPF and AMD. It causes a machine check when that happens.

It doesn't happen on Proliant because the last 4kB (or so) of memory is mapped out by the BIOS and Pentium guarantees contiguous memory.

If the platform is vulnerable to this then driver's prefetching needs to be disabled at compile-time or, preferably, initialization-time. What is the best means by which we can determine whether the platform needs this treatment?

(the patch didn't compile, btw: there's no definition of I200_DMA1_CFG)

–

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in

Re: [PATCH 2/2] cciss: disable dma prefetch for P600

Re: [PATCH 2/2] cciss: disable dma prefetch for P600

the body of a message to majordomo@xxxxxxxxxxxxxxxxx

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

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