

## [PATCH] (Longhaul 1/5) PCI: Protect bus master DMA from Longhaul by rw semaphores

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-06/msg09698.html>

---

- *From:* Rafał Bilski <[rafalbilski@xxxxxxxxxx](mailto:rafalbilski@xxxxxxxxxx)>
  - *Date:* Wed, 28 Jun 2006 20:25:43 +0200
- 

You mean the longhaul driver can change the frequency of the PCI bus? Oh, that's a recipe for disaster...

No. Sorry. My English is bad. I mean changing CPU frequency.

No, it's a hack :)

Again :-)

No, this is not acceptable. What exactly do you want to do here? Make sure the PCI drivers are not doing DMA when the longhaul driver wants to change the pci bus speed?

I'm trying not to break DMA. Current version of longhaul (marked broken in 2.6.16.2) simply clears bus master bit on every device.

Does it really save battery?

Yes. And CPU temperature is lower.

And what about PCI devices that always do DMA? (think USB controllers, they can easily saturate the PCI bus all the time).

This is worst for SATA. USB (this is strange) seems to work correctly. I know that this is 10% coverage, but it is better than nothing. It is always possible to add support for longhaul to driver.

[PATCH] (Longhaul 1/5) PCI: Protect bus master DMA from Longhaul by rw semaphores

Why not just suspend all PCI devices make the bus change, and then resume them? That would require no PCI core, or driver changes.

This was my first idea. But trust me in current kernel this is simply worst idea.

greg k-h

Though currently in the driver, voltage scaling is missing, so we never save any power, and just run at the maximum voltage the whole time.

I added this to longhaul, but it only works on non EPGA CPU's. EPGA CPU's (at least Nehemiah) seem to have voltage scaling disabled.

It needs there to be no bus mastering occuring at the time of a CPU speed transition. Though I'm unable to find the part that mentions this in the specs I have right now.

Dave

"Once this is set, the processor will switch to the value in [26:23] on the next AUTOHALT transition. The duration of the AUTOHALT should be  $\geq 1$ ms to ensure the CPU's internal PLL is resynchronized. For AUTOHALT, this means interrupts must be disabled except for the time tick, which should be reset to  $\geq 1$ ms. Care must be taken to avoid other system events that could interfere with this operation. A few examples are snooping, NMI, INIT, SMI and FLUSH."

For CPU's with Longhaul MSR this time is equal to 200us.

RafaB

---

PS. Fajny portal... >>> <http://link.interia.pl/f196a>

-

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

[PATCH] (Longhaul 1/5) PCI: Protect bus master DMA from Longhaul by rw semaphores