

## Re: PCI Express MMCONFIG and BIOS Bug messages..

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-04/msg11575.html>

---

- *From:* Robert Hancock <[hancockr@xxxxxxx](mailto:hancockr@xxxxxxx)>
  - *Date:* Sun, 29 Apr 2007 12:39:09 -0600
- 

Jesse Barnes wrote:

On Sunday, April 29, 2007, Robert Hancock wrote:

Problem is that even if we read the MMCONFIG table location from the hardware registers, that doesn't mean we can trust the result. It could be that the BIOS hasn't lied about where it put the table, it just stuck it someplace completely unsuitable like on top of RAM or other registers. It seems that with some of those 965 chipsets the latter is what the BIOS is actually doing, and so when we think we're writing to the table we're really writing to random chipset registers and hosing things. (Jesse Barnes ran into this while trying to add chipset support for the 965).

Right, I've updated the BIOS since, but at least that version was totally buggy wrt MMconfig support. I haven't yet looked at the new one to see if it properly reserves MCFG space in ACPI\_CRS yet or properly programs it.

Likely what we need to do is:

- If chipset is known, take table address from registers, otherwise check the MCFG table
- Take the resulting area (Ideally not just the first minimum part as we check now, but the full area based on the expected length) and make sure that the entire area is covered by a reservation in ACPI motherboard resources.
- If that passes, then we still need to sanity check the result by making sure it hasn't been mapped over top of something else important. How to do this depends on exactly how they've set up the ACPI reservations on these broken boxes.. Does someone have a full dmesg from one on a recent kernel that shows all the pnpacpi resource reservation output?
- If these checks fail, we don't use the table, and the chipset is known, we should likely try to disable decoding of the region so that it won't get in the way of anything else.

Re: PCI Express MMCONFIG and BIOS Bug messages..

Yeah, that sounds like a good algorithm.

I'm not sure how to handle the fact that we don't have access to the \_CRS until late in boot though... Len?

We'd likely have to split the MMCONFIG initialization into two parts. The early part enables MMCONFIG only on systems where we require it (like the Macs that Andi mentioned). On all other systems we defer enabling it (use the regular PCI configura