

Re: [PATCH 2/2] Initial generic hypertransport interrupt support.

Re: [PATCH 2/2] Initial generic hypertransport interrupt support.

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-07/msg04364.html>

- *From:* Dave Olson <olson@xxxxxxxxxxxxx>
 - *Date:* Wed, 12 Jul 2006 20:56:46 -0700 (PDT)
-

On Wed, 12 Jul 2006, Eric W. Biederman wrote:

| Dave Olson <olson@xxxxxxxxxxxxx> writes:

|

| > On Tue, 11 Jul 2006, Eric W. Biederman wrote:

| > | There is a hypertransport capability that implements a rough equivalent
| > | of a per device ioapic. It is quite similar to MSI but with a different
| > | register level interface.

| >

| > It's really just the same as MSI, and is set up and handled pretty
| > much the same way.

|

| No it is not just the same. There is not global enable bit, only

...

| But from the perspective of using them in a driver the concept really
| is the same.

I meant, from the perspective of a driver. Sorry for not being clear.

| The code that breaks it is only in -mm. It's scheduled for 2.6.19.

| All of the MSI magic in ioapic land on i386 and x86_64 is deleted.

| The code just needs to age a bit and let the few unexpected
| corner case crop up, and get sorted out.

Well, if that set of patches is accepted into 2.6.19, it will likely break other people with HyperTransport chips and cards as well. I do know of other HTX-slot cards in development, and some of them, at least, do use interrupts.

So I think it needs to age enough to not break existing drivers.

Unfortunately, I still haven't had time to try your HT path with our ipath driver, because I'm in fire-fighting mode on some customer issues. You might be able to test it on some of your LS-1 systems in your lab, if you need early feedback. I'll try it, as soon as I can.

| > This part I never really quite understood. Why do you want a separate
| > interface than the existing request_irq().

Re: [PATCH 2/2] Initial generic hypertransport interrupt support.

Re: [PATCH 2/2] Initial generic hypertransport interrupt support.

| request_irq is still needed. The question is how do you get the irq.

| > and pci_enable_msi()?

| The HT and msi semantics are moderately different, but I have
| implemented the equivalent of pci_enable/disable_msi. So the
| code is not a pci standard but just a ht standard I didn't use the
| pci prefix.

My point was that many other pci_* functions have to be called by any driver for HyperTransport-attached chips (that aren't chipset), so why break these out separately, rather than somehow fitting them into the existing routines, perhaps by looking at the bus-type the device is attached to?

Unless we add a full set of ht_* routines (not something I'd like to see!) I don't see a reason to do it for just these routines, other than convenience for early testing of patch proposals, as opposed to final code going into the mainline.

I'm not suggesting re-using existing (old) msi.c code. I'm suggesting making the new MSI code work for all 3 of PCI, PCIe, and HT, with the same exact exported routine.

| Of course I don't expect the interface exported to drivers to change.

But if you add new ht_* routines that they have to call, that's definitely a change, a new set of exported routines. I'm questioning whether that's really the right direction. If you had already intended to have them not exported in the final patches, that wasn't clear to me, and I apologize for misunderstanding.

Dave Olson
dave.olson@xxxxxxxxxxx

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxxxx

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

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

Re: [PATCH 2/2] Initial generic hypertransport interrupt support.