

Re: [PATCH] 2.6 SGI Altix I/O code reorganization

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-10/1612.html>

From: Jesse Barnes (jbarnes_at_engr.sgi.com)

Date: 10/06/04

To: Grant Grundler <iod00d@hp.com>
Date: Wed, 6 Oct 2004 13:27:28 -0700

On Wednesday, October 6, 2004 12:54 pm, Grant Grundler wrote:

> *Colin,*
> *thanks for ACKing the feedback.*
> *I think there is still some confusion...*
>
> *On Wed, Oct 06, 2004 at 02:09:54PM -0500, Colin Ngam wrote:*
> ...
>
>> *Mathew explained replacing the raw_pci_ops pointer is the Right Thing*
>> *and I suspect it's easier to properly implement.*
>>
>> *I believe we did just that. We did not touch pci_root_ops.*
>
> *Correct. The patch ignores/overrides pci_root_ops with sn_pci_root_ops*
> *(which is what I originally suggested).*
>
> *Mathew's point was only raw_pci_ops needs to point at a different*
> *set of struct pci_raw_ops (see include/linux/pci.h).*

Though now what's there seems awfully redundant, wouldn't you say? Just allowing direct access to pci_root_ops is a much simpler approach and gets rid of a bunch of extra, unneeded code (i.e. closer to Pat's original version).

>> *Yes, would anybody allow us to make a platform specific callout*
>> *from within generic pcibios_fixup_bus()???*
>
> *If it can be avoided, preferably not. But that's up to Jesse/Tony I think.*

If it was made a machine vector that's a no-op on everything but sn2, I think it would be fine. Doing it for the general sn_pci_init routine would let us get rid of the check for ia64_platform_is("sn2") in one of the routines, I think (which is nice if only for the consistency).

> *Can you quote the bit of the patch which implements "if the bus does not*
> *exist" check?*
> *I can't find it.*

Linux-Kernel: Re: [PATCH] 2.6 SGI Altix I/O code reorganization

In the current code it's:

```
for (i = 0; i < PCI_BUSES_TO_SCAN; i++)  
    if (pci_bus_to_vertex(i))  
        pci_scan_bus(i, &sn_pci_ops, controller);
```

which causes the next loop to only fixup existing busses. But I don't see it in the new code.

> > *One favour. Would you agree to letting this patch be included by Tony
> > and we will come up with another patch to fix the 2 obvious items listed
> > above? It will be great to avoid spinning this big patch.*

The patch is ok with me, I think it's a big improvement over what's there in terms of readability.

I just checked out `sn_set_affinity_irq()` and it's a bit hard to see what's going on. Why does a new interrupt have to be allocated? Also, it looks like the `kfree()` is one line too high, if `sn_intr_alloc` fails, we'll leak `new_sn_irq_info`.

Jesse

-

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/>