

Re: API change for bus_dma

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/arch/2003-06/0332.html>

From: John-Mark Gurney (gurney_j_at_efn.org)

Date: 06/28/03

Date: Sat, 28 Jun 2003 01:04:48 -0700

To: Terry Lambert <tlambert2@mindspring.com>

Terry Lambert wrote this message on Sat, Jun 28, 2003 at 00:52 -0700:

> *Andrew Gallatin wrote:*

> > *Scott Long writes:*

> > > *The approach taken with busdma is that you don't assume coherency. The*

> >

> > *Unfortunately, in our application we must assume coherency in some*

> > *situations. We have kernel memory mmap'ed into user space for*

> > *zero-copy io of small messages. Doing a system call to force the dma*

> > *sync would add unacceptable latency. (we're talking sub 10us latencies*

> > *here, without syscalls).*

>

> *"contigmalloc"*

>

> *You have to do the same thing for BT848 and fram buffers that eat*

> *host memory instead of having their own to play in. In general,*

> *this has to be done in the device driver, very early on in the life*

> *of the system to stand any chance of succeeding: because we do not*

> *properly use ELF section attribute tags, it's not possible to defrag*

> *physical memory in FreeBSD to do these allocations later. The tags*

> *are necessary to identify code in the defragmentation code path so*

> *that it does not attempt to relocate itself while it's running itself.*

>

> *NB: FreeBSD doesn't support kernel paging and discard of init routines*

> *in device drivers, once the driver is operational, for the same reason:*

> *lack of section tags indicating "paging path" in the first instance,*

> *and lack of section tags indicating "discard after initialization" in*

> *the second.*

I'm sorry, no, this will not solve the problem he is talking about.

You need to reread the information that Andrew has provided before.

In a previous email you got confused on the STREAMING/COHERENT flag's

meaning. Using contigmalloc only gives you a linear address space,

but does not guarantee that the processor will snoop the memory write

cycles by the bridge or device to keep the cache of the cpu the same

with the memory. For what Andrew needs, he needs the processor to have

the same information as in memory. On multiprocessor systems, it can

get expensive if every processor has to snoop every memory write that

freebsd-arch: Re: API change for bus_dma

happens.

--

John-Mark Gurney

Voice: +1 415 225 5579

"All that I will do, has been done, All that I have, has not."

freebsd-arch@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-arch>

To unsubscribe, send any mail to "freebsd-arch-unsubscribe@freebsd.org"