

## RE: Problem with twa in HEAD

**Source:** <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/current/2005-04/1450.html>

---

**From:** Vinod Kashyap (vkashyap\_at\_amcc.com)

**Date:** 04/29/05

To: "Bjoern A. Zeeb" <bz@FreeBSD.org>

Date: Thu, 28 Apr 2005 20:50:47 -0700

> -----Original Message-----

> From: Bjoern A. Zeeb [mailto:bz@FreeBSD.org]

> Sent: Tuesday, April 26, 2005 3:26 AM

> To: Vinod Kashyap

> Subject: RE: Problem with twa in HEAD

>

>

> On Mon, 25 Apr 2005, Vinod Kashyap wrote:

>

> Hi,

>

>> -----Original Message-----

>> From: Bjoern A. Zeeb [mailto:bz@FreeBSD.org]

>> Sent: Monday, April 25, 2005 6:45 AM

>> To: Vinod Kashyap

>> Subject: Re: Problem with twa in HEAD

>>>

>>>

>>> On Fri, 22 Apr 2005, Bjoern A. Zeeb wrote:

>>>

>>> Hi,

>>>

>>>> scottl redirected me to you.

>>>>

>>>> I am currently debugging "hangs" on reboot and shutdown on a

>>>> SMP machine with 12 discs at a

>>>>

>>>> 3ware device driver for 9000 series storage controllers,

>>>> version: 3.60.00.016

>>>> twa0: <3ware 9000 series Storage Controller> port

>>>> 0x9800-0x98ff mem 0xfe8ffc00-0xfe8ffcff,0xfb800000-0xfbffffff

>>>> irq 28 at device 6.0 on pci3

>>>> twa0: [FAST]

>>>> twa0: INFO: (0x15: 0x1300): Controller details:: 12 ports,

>>>> Firmware FE9X 2.06.00.009, BIOS BE9X 2.03.01.051

>>>>

```
>>>>
>>>> What I know so far is that Giant is held by sync.
>>>>
>>>> Things a "spinning" in cam/cam_xpt.c around:
>>>>
>>>> --- cam_xpt.c 31 Mar 2005 21:42:49 -0000 1.152
>>>> +++ cam_xpt.c 22 Apr 2005 18:42:43 -0000
>>>> @@ -3643,6 +3643,7 @@ xpt_polled_action(union ccb *start_ccb)
>>>> != CAM_REQ_INPROG)
>>>> break;
>>>> DELAY(1000);
>>>> printf("XXX status=%02x\n",
>>> start_ccb->ccb_h.status);
>>>> }
>>>> if (timeout == 0) {
>>>> /*
>>>>
>>>>
>>>> with status being 0x200.
>>>>
>>>> Seems the twa has a command stuck in it.
>>>>
>>>> I have seen the comment in dev/twa/tw_osl_cam.c ~ line 253 about
>>>> queuing and CAM_SIM_QUEUED but I don't know enough about cam.
>>>> I seems no all patches out of this functions seem to
> clear that from
>>>> status?
>>>>
>>>> Any help apreaciated ;) I can try patches; as long as I
> can break
>>>> to db> to reboot.
>>>>
>>>> further debugging shows that is seems to be spinning in twa_poll.
>>>> see debug output from TWA_DEBUG 3. The problem is that at
> this point
>>>> I am no longer able to break to debugger.
>>>>
>>>> twa0: tw_osli_execute_scsi: XPT SCSI_IO: Single virtual address!
>>>> twa0: tw_osli_execute_scsi: XPT SCSI_IO: Single virtual address!
>>>> unmount of /dev failed (BUSY)
>>>> twa0: tw_osli_execute_scsi: XPT SCSI_IO: Single virtual address!
>>>> twa0: tw_osli_execute_scsi: XPT SCSI_IO: Single virtual address!
>>>> Uptime: 2m57s
>>>> twa0: tw_osli_execute_scsi: XPT SCSI_IO: Single virtual address!
>>>> twa0: twa_poll: entering; sc = 0xc57bb200
>>>> twa0: twa_poll: exiting; sc = 0xc57bb200
>>>> twa0: twa_poll: entering; sc = 0xc57bb200
>>>> twa0: twa_poll: exiting; sc = 0xc57bb200
>>>> twa0: twa_poll: entering; sc = 0xc57bb200
>>>> twa0: twa_poll: exiting; sc = 0xc57bb200
>>>> twa0: twa_poll: entering; sc = 0xc57bb200
```

freebsd-current: RE: Problem with twa in HEAD

```
> > > twa0: twa_poll: exiting; sc = 0xc57bb200
> > > twa0: twa_poll: entering; sc = 0xc57bb200
> > > twa0: twa_poll: exiting; sc = 0xc57bb200
> > > twa0: twa_poll: entering; sc = 0xc57bb200
> > > twa0: twa_poll: exiting; sc = 0xc57bb200
> > > twa0: twa_poll: entering; sc = 0xc57bb200
> > > twa0: twa_poll: exiting; sc = 0xc57bb200
> > > ...
> > >
> >
> > I am in the middle of an office move right now.
> > I will get back to you once I have some time to look into this.
>
>
> thanks for the information; I'll be able to test at least until end of
> this week and hopefully next week too.
>
```

I looked into this, and this is what is happening:

On reboot/halt, the following function calling sequence happens:  
... --> dashutdown --> xpt\_polled\_action --> twa\_poll.  
But, the interrupt handler in twa is still active at this time,  
since twa\_detach/twa\_shutdown hasn't been called yet. Before  
twa\_poll can fetch the response for the posted command, the ISR  
gets called when the firmware posts the response. The ISR clears  
the interrupt bit on the controller, registers a taskqueue handler like  
it always does, and exits. Meanwhile, xpt\_polled\_action continues  
to call twa\_poll, which cannot determine that the command has completed,  
since the interrupt bit on the controller is already cleared. So,  
we get into a (near) never-ending loop (the timeout for scsi\_synchronize\_cache,  
which is what is being tried here, is, for whatever reason, 60 minutes,  
and so, the system is as good as hung).

Now, does anyone know why xpt\_polled\_action is being called from  
dashutdown, even before the ISR has been unregistered (via twa\_detach)?

Bjoern, this patch should work-around your problem, although it's not  
the fix. Also, it still leaves a window for the race condition described  
above.

```
diff -u -r ../twa.cur/tw_osl.h ./tw_osl.h
--- ../twa.cur/tw_osl.h Fri Apr 8 12:43:45 2005
+++ ./tw_osl.h Thu Apr 28 20:28:40 2005
@@ -71,6 +71,7 @@
 /* Possible values of sc->state. */
 #define TW_OSLL_CTLR_STATE_OPEN (1<<0) /* control device is open */
 #define TW_OSLL_CTLR_STATE_SIMQ_FROZEN (1<<1) /* simq frozen */
+#define TW_OSLL_CTLR_STATE_POLLING (1<<2) /* polling for ctrl response */

#ifdef TW_OSL_DEBUG
```

freebsd-current: RE: Problem with twa in HEAD

```
diff -u -r ../twa.cur/tw_osl_cam.c ./tw_osl_cam.c
--- ../twa.cur/tw_osl_cam.c Fri Apr 8 12:43:57 2005
+++ ./tw_osl_cam.c Thu Apr 28 20:29:22 2005
@@ -482,6 +482,7 @@
     struct twa_softc *sc = (struct twa_softc *) (cam_sim_softc(sim));

     tw_osli_dbg_dprintf(3, sc, "entering; sc = %p", sc);
+ sc->state |= TW_OSلي_CTLR_STATE_POLLING;
     if (tw_cl_interrupt(&(sc->ctrl_handle)))
         tw_cl_deferred_interrupt(&(sc->ctrl_handle));
     tw_osli_dbg_dprintf(3, sc, "exiting; sc = %p", sc);
diff -u -r ../twa.cur/tw_osl_freebsd.c ./tw_osl_freebsd.c
--- ../twa.cur/tw_osl_freebsd.c Fri Apr 8 12:44:12 2005
+++ ./tw_osl_freebsd.c Thu Apr 28 20:31:25 2005
@@ -964,6 +964,8 @@
     struct twa_softc *sc = (struct twa_softc *) arg;

     tw_osli_dbg_dprintf(10, sc, "entered");
+ if (sc->state & TW_OSلي_CTLR_STATE_POLLING)
+ return;
     if (tw_cl_interrupt(&(sc->ctrl_handle)))
         taskqueue_enqueue_fast(taskqueue_fast,
            &(sc->deferred_intr_callback));
```

---

freebsd-current@freebsd.org mailing list

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

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