

Re: USB regression (and other failures) in 2.6.2[45]*

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-02/msg08403.html>

- *From:* Andrew Buehler <abuehler.kernel@xxxxxxxxxx>
 - *Date:* Sat, 16 Feb 2008 20:12:40 -0500
-

On 2/16/2008 6:11 PM, Alan Stern wrote:

On Sat, 16 Feb 2008, Andrew Buehler wrote:

For another, getting two copies of a message is no big deal

I disagree.

Everyone has his own taste. Obviously there's no world-wide consensus, possibly because different people have different workflow habits and so are affected by duplicate messages to varying extents.

I am well aware that this particular point is opinion. I have had justifications for and arguments in favor of it in the past, but none of them come readily to mind at the moment, except for the one gone over briefly below.

When I receive a message sent directly to me in a discussion which is on a list, I expect that it is because someone either considered it important enough to warrant making certain it came to my attention specifically, or wanted to continue the discussion but felt that it should not continue to take place on the mailing list.

Sometimes that is the case but often it isn't. Your expectations are at variance with other people's behavior; you shouldn't expect everyone else to change just to match your personal ideals.

Messages sent to my address directly are explicitly not filtered into

Re: USB regression (and other failures) in 2.6.2[45]*

the folders I have set up for various mailing lists, so that if someone does send me a "heads up" reply for a specific topic on a list to which I am subscribed it does not get caught by the list filter and fail to come to my attention. If a message fails to be filtered into any mailing-list folder, then I should be able to conclude that it is specifically intended for me, and not part of normal mailing-list traffic. The practice of sending replies to both addresses renders this an invalid conclusion. I do not think that it is unreasonable to expect that conclusion to be valid.

On the other hand, I would be perfectly happy to edit your name out of the reply list -- but since you said you aren't receiving all the messages in this thread via the list that might not be a good thing to do at the moment...

It's not that I'm not receiving all of this thread's messages via the list -- it's that I'm not receiving *any* of them via the list, and I suspect that the reason is that my address is in both the To:/Cc: and the list itself. Something is filtering it such that I do not receive "duplicate" replies in this way, but it is doing so by filtering out the list copy rather than the direct copy. I have seen mailing lists which do this before, but I see no other indication that the LKML is one of them, and I would not be in the least surprised if this turned out to be yet one more problem with gmail.

As far as I am aware, I am seeing all messages posted to the list which do not have me in To: or Cc:. I suspect that if a reply in this thread were posted to the list but not sent to me, I would see it on the list. It might be worth an experiment, but since it would increase traffic for other list members to no purpose it is probably not worth it overall.

People on LKML who are more familiar with interrupt routing problems might be able to offer more help. For now, you can try things like turning on CONFIG_USB_DEBUG, posting the output from dmesg, posting the contents of /proc/interrupts (say before and after a new USB device is plugged in).

In my current testing kernel, which I believe is the one with which I captured the sole successful non-2.6.23.1 dmesg so far, CONFIG_USB_DEBUG is on. The associated dmesg (obtained yesterday from booting with the Flash drive connected) is attached. (The flood of 'no version magic, tainting kernel' messages between line 600 and line 1160 are a side effect of Novell's custom environment

Re: USB regression (and other failures) in 2.6.2[45]*

which I have not yet made the effort to fix; the boot scripts attempt to detect the network card by modprobing every network driver available until they find one which works. Here, because the correct one fails, they wind up trying each one twice.)

The line saying:

```
ehci_hcd 0000:00:1d.7: Unlink after no-IRQ? Controller is probably using the wrong IRQ.
```

is an indication that interrupt routing is indeed not working right. Or possibly your EHCI controller isn't working. You could try blacklisting or unloading ehci-hcd to see if that helps. Of course then none of your USB devices would be able to run at high speed.

ehci-hcd is not modular in my current kernel, and if there is a way to turn it off without its being modular I am not aware of it. I will have to jump through a few hoops to be able to obtain a copy of the boot CD with an updated kernel while not at work, but I will try to do so sometime tomorrow.

In practical terms, I am frankly not especially bothered by the lack of support for high-speed USB in Linux on this machine; the primary reason I am interested in USB there at the moment, aside from a general philosophy of "unsupported devices are bad and anything I can do to help them become supported is good", is because getting it working would allow me to easily get the necessary information out to be able to properly report the other problems, with AHCI and networking.

I have transcribed the contents of /proc/interrupts both before and after plugging in the Flash drive I have been using for testing, and they are also attached. I have been as careful as I could to be sure that the contents of the attached 'r61-interrupts-[12].txt' files is the same as what was shown on the laptop, but cannot absolutely guarantee that I have not missed something. For the record, the '1' is from before connecting the drive, and the '2' is from after.

Notice that the interrupt count for IRQ 11 doesn't change when you plug in the device. Obviously something is wrong there.

In fact, it's a little surprising that almost all the USB controllers are routed to the same IRQ. However this is beyond my area of expertise. You could try posting a message on the linux-acpi mailing list; the people there should know a lot more about these issues.

Re: USB regression (and other failures) in 2.6.2[45]*

Until this thread, I was not even aware that ACPI was related to USB; I had largely conflated it with a similar acronym which I think is related to power management and which I can suddenly not even find in my kernel config. I will, however, look into linux-acpi.

Assuming that the 2.6.23 kernel works on your computer, you can go the extreme route of installing git and doing a bisection to find the first patch causing your difficulty.

That would require me to learn enough of how git works, as distinct from more traditional VCSes, to be able to use it with some confidence. This is not impossible – indeed I want to do it at some point – but for the time being I have no idea where to start, and indeed I am not especially clear on exactly what (from a user's perspective) the differences between git and e.g. CVS or Subversion are. I know that the entire concept relies around a lack of centralization, but I have not been able to get my head around what that means in a practical sense.

There are some excellent tutorials on the web, with detailed explanations of how to do a bisection to track down a kernel bug.

I have found at least a place to start, and am reading up on the subject. I will most likely not be able to make a practical start on this until at least Tuesday, as not having direct access to the machine I will in the long term be building on makes some things impractical, but if no solution is forthcoming in the meantime I will expect to do this.

That will not be helpful for the other two problems, however, since neither of them was ever working as far as I am aware. That also leaves me hesitant to conclude that they are rooted in the same IRQ issue as the USB problem appears to be.

Which lists or other addresses would be appropriate for reporting problems with AHCI/libata and with networking, specifically with the e1000/e1000e drivers? I see a mailing list for e1000 in MAINTAINERS, but only the maintainer's address for SATA/libata/whatever else may be involved there, and I am reflexively reluctant to bother a maintainer directly with as little information as I presently have.

--

Andrew Buehler

--

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

Re: USB regression (and other failures) in 2.6.2[45]*

Re: USB regression (and other failures) in 2.6.2[45]*

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

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