

Re: Problem: NIC transmit timeouts

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-03/msg02149.html>

- *From:* "John Z. Bohach" <jzb@xxxxxxxxxxxxxx>
 - *Date:* Tue, 7 Mar 2006 08:48:00 -0800
-

On Tuesday 07 March 2006 04:08, PaNiC wrote:

Thanks for your reply.
I have an idea what you're talking about, but no more i'm afraid.
I'm no programmer and i don't know how to try what you suggested.
What i did try was applying some patches manually that i found in the mailing list archives.

This is one of them:

I'll try to elaborate a little, but since my experience in dealing with the NETDEV watchdog timeout issue is confined to x86 architectures, some translation may be necessary...

I've seen two causes for this: 1) driver bug, 2) firmware bug. I have yet to come across one where the kernel itself was at fault.

1) if its a driver bug, updating to the latest driver should have an effect, hopefully a positive one. Generally, the driver maintainer is at least aware of the problem, if this is the cause.

2) If its a firmware bug, you'll have to upgrade to the latest version of the BIOS, uboot, rommon, or whatever your platform uses to boot the OS.

Case 2) is the type of problem where many people can be using the exact same kernel and exact same driver, and only a few experience any problems. This is because its not the kernel or the driver, but the firmware that is at fault.

In particular, on the x86 platform where I've root-caused this issue, a different version of the BIOS was present (I did an upgrade, and the upgrade failed to implement the proper register programming, and that's when the problem showed up). The manifestation is that heavy PCI traffic (not just ethernet traffic) would overload the host-PCI bridge, and bad things started happening. In some cases, just ethernet traffic alone was enough to overpower the bridge, but that depended on the link speed. 100 Mbs took a lot longer to NETDEV watchdog timeout than 1000 Mbs. Another fun aspect of case 2) is that since its PCI traffic related, its very hard to recreate exactly, since there's usually other stuff on the PCI bus.

Re: Problem: NIC transmit timeouts

So if nobody seems to have a solution for your problem, it is possible that its firmware-related, so try getting a different firmware. In particular, if you have any experience at this, or know someone in your organization who does, try focusing on the MTT (multi-transaction timeout) register, and the equivalent (if any) of the ICH configuration register in the host-PCI bridge. These are both proprietary registers, and are not present in the standard 64 byte PCI header, so its pretty hard to "guess" where they are what they should be set to. If you have a similar platform that does not have any NETDEV watchdog timeout issues, it may be possible to dump the PCI config. space of the host-PCI bridge on that platform, and compare it with the failing platform, and start "guessing" from there, but the firmware is the responsible entity that should have these registers programmed correctly in the first place.

Sorry for being so verbose, I'm hoping that others googling in the future may save themselves some time in isolating this pretty obscure cause...

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx
More majordomo info at <http://vger.kernel.org/majordomo-info.html>
Please read the FAQ at <http://www.tux.org/lkml/>