Packet reordering and blocking problem at gigabit with 2.4 kernel

Source: http://linux.derkeiler.com/Mailing−Lists/Kernel/2003−09/6045.html

---

From: Hyunje Park (*hjpark_at_zooin.net*)
Date: 09/23/03

```
To: <linux−kernel@vger.kernel.org>
Date:   Tue, 23 Sep 2003 22:27:14 +0900
```

I am developing application level multicast router with Xeon processor.
The system lost lots of packets and reordered the packets, and finally
was blocked with heavy load.
Some of packets are lost and the lost packet appeared after a few
seconds from its normal sequence.

The specification of the systems are:
Intel SE7501VW (7501 chipset with 533MHz interface) or Tian S2723(7500
chipset with 400MHz,)
2.4 GHz or 3.0 GHz Xeon (two CPUs, one CPU without hyperthreading was
tested)
Intel 82544SX Gigabit card

Software vesions experimented:
Linux kernel 2.4.18, 2.4.20, 2.4.22
E1000 drivers: 4.x, 5.1.x, 5.2.16 (with/without NAPI option)

Every channel sends 6Mbps with 2048 bytes packet. The senders sends upto
100~120 channel to the multicast
relay and the relay forwards the channels to the receiver using unicast
or multicast.

The L3 switch, Alteon 780 cannot support IGMP snooping currently. It
means the multicast traffics may be bounced
to the multicast relay as broadcast. Every channel are generating about
800 pps: 100 channel will generate 80Kpps.

I tried to experiment a lot of diffenent combinations with software,
hardware, or configuration.
I heard that SMP machines has inherant reordering issues, so I turned
off SMP, but no improvement.
I also tried to turn on NAPI, but no improvment.
I tried to Interrupt mitigation by setting (RX, TX)Intdelay between 100
and 500. The performance becomes better, but the
reordering problems are not discarded.

With the faster CPU, there are no differences either.

I don't know what else I have to try to fix the bugs.
Is it inherant problems with current version of Linux kernel?
Or is there a bug of E1000 driver?
Am I giving too much stress to the system?
I will test with Sysconnect gigabit ethernet card tomorrow, but I am not sure if it solve the problem.

Pease CC me with any suggestions or comments since I am not subscribed to the list.
Thanks in advance.

Hyunje Park

–