

receive path with fragmented skbs

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-01/7799.html>

From: Kallol Biswas (kallol.biswas_at_s2io.com)

Date: 01/31/04

To: linux-kernel@vger.kernel.org

Date: 30 Jan 2004 15:12:24 -0800

Hello,

We have been developing drivers and networking software on a 10 gigabit ethernet adapter from S2io Inc (www.s2io.com). There is a requirement that the ethernet header, IP+TCP headers have to be cache aligned and the payload and the IP+TCP headers have to be in different fragments. So we have created receive path skbs with data size big enough to hold the ethernet header and two fragments, one fragment for the IP+TCP header and the other for payload. The card can directly dma into the three receive scatter buffers when a frame arrives.

We could not get ping working with this design of receive skbs, but if a skb is linearized with `skb_linearize()` before calling `netif_rx()`, ping works.

`/proc/net/snmp` was printed, no frame had any error. Probably no one has ever tested the receive path of the stack with fragmented skbs, am I right? One of the ways this problem can be debugged is to find out where exactly the packets get dropped. Any comment?

Kallol

—

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

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

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