

## Re: Van Jacobson's net channels and real-time

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-04/msg05008.html>

---

- *From:* "linux-os \ (Dick Johnson)" <[linux-os@xxxxxxxxxxxxx](mailto:linux-os@xxxxxxxxxxxxx)>
  - *Date:* Mon, 24 Apr 2006 14:12:01 -0400
- 

On Mon, 24 Apr 2006, Rick Jones wrote:

That's right. This will be made a non issue  
with early demuxing  
on the NIC and MSI (or was it MSI-X?)  
which will select  
the right CPU based on hardware channels.

MSI-X. with MSI you still have only one cpu handling all  
MSI interrupts and  
that doesn't look any different than ordinary interrupts.  
MSI-X will allow  
much better interrupt handling across several cpu's.

Auke  
-

Message signaled interrupts are just a kudge to save a trace on a  
PC board (read make junk cheaper still). They are not faster and  
may even be slower. They will not be the salvation of any interrupt  
latency problems. The solutions for increasing networking speed,  
where the bit-rate on the wire gets close to the bit-rate on the  
bus, is to put more and more of the networking code inside the  
network board. The CPU get interrupted after most things (like  
network handshakes) are complete.

if the issue is bus vs network bitrates would offloading really buy that  
much? i suppose that for minimum sized packets not DMA'ing the headers  
across the bus would be a decent win, but down at small packet sizes  
where headers would be 1/3 to 1/2 the stuff DMA'd around, I would think  
one is talking more about CPU path lengths than bus bitrates.

and up and "full size" segments, since everyone is so fond of bulk

Re: Van Jacobson's net channels and real-time

transfer tests, the transfer saved by not shoving headers across the bus is what 54/1448 or ~3.75%

spreading interrupts via MSI-X seems nice and all, but i keep wondering if the header field-based distribution that is (will be) done by the NICs is putting the cart before the horse – should the NIC essentially be telling the system the CPU on which to run the application, or should the CPU on which the application runs be telling "networking" where it should be happening?

rick jones

Ideally, TCP/IP is so mature that one should be able to tell some hardware state-machine "Connect with 123.555.44.333, port 23" and it signals via interrupt when that happens. Then one should be able to say "send these data to that address" or "fill this buffer with data from that address". All the networking could be done on the board, perhaps with a dedicated CPU (as is now done) or all in silicon.

So, the driver end of the networking software just handles buffers. There are interrupts that show status such as completions or time-outs, trivial stuff.

Cheers,

Dick Johnson

Penguin : Linux version 2.6.16.4 on an i686 machine (5592.89 BogoMips).

Warning : 98.36% of all statistics are fiction, book release in April.

–

\*\*\*\*\*

The information transmitted in this message is confidential and may be privileged. Any review, retransmission, dissemination, or other use of this information by persons or entities other than the intended recipient is prohibited. If you are not the intended recipient, please notify Analogic Corporation immediately – by replying to this message or by sending an email to DeliveryErrors@xxxxxxxxxxxxx – and destroy all copies of this information, including any attachments, without reading or disclosing them.

Thank you.

–

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

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

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