

Re: A question regarding MTU: how it can effect TCP performance + other queries

Re: A question regarding MTU: how it can effect TCP performance + other queries

Source:

<http://www.tech-archive.net/Archive/Development/microsoft.public.development.device.drivers/2006-12/msg00098.html>

- *From:* Pankaj Garg <pankajg@xxxxxxxxxxxxxxxxxxxxxxxxxxxx>
 - *Date:* Wed, 29 Nov 2006 23:38:17 -0800
-

Removing multiple newsgroups as cross posting usually generates tons of extra traffic.

What packet size IO are you doing in case of TCP? Can you check if your physical NIC has TCP large send offload enabled?

The reason for the large difference in TCP performance may be due to the offload happening in case of physical NIC.

I can't think of anything for the UDP case however, that just seems strange to me. You probably need to dig deeper and tell us what exactly is happening here. How many packets are being sent, what was the size of these packets? Are you grouping multiple UDP packets in one TCP packet?

--

Pankaj Garg
2006-11-29 at 11:32pm
<http://www.intellectualheaven.com>

On Wed, 29 Nov 2006, Steve Dispensa wrote:

I'm sorry, I misread the UDP part of your question. You're right, something seems off there. Are you buffering somewhere somehow? Are some packets hitting the bit bucket after being delivered to the virtual NIC, perhaps?

As for TCP case, you said:

What I found that for TCP, throughput that I am getting on physical interface is around 900 Mbps and on virtual interface i am getting around 90 Mbps (only 10% of the physical one). All the tests were performed with Iperf with 15 parallel streams.

Re: A question regarding MTU: how it can effect TCP performance + other queries

and

So, for TCP packets over virtual interface, I am sending a TCP packet encapsulated within another TCP packet when passed to physical interface, while for UDP I am sending UDP packet encapsulated within TCP packet when passed to physical interface.

Nagle would matter for UDP packets encapsulated within TCP packets, depending (a great deal) on the characteristics of the traffic you are testing with. Nagle's has an adverse affect on transmission speed, because it introduces an artificial queuing delay, regardless of what it's carrying. I suggested turning on NODELAY for the "outside" TCP connection, not (necessarily) the "inside" connection.

At any rate, your results do seem strange to me; not a whole lot else comes to mind though. You might want to instrument your miniport and track where the packets spend their time between receipt at your send handler and final send() to the other system.

Good luck.

-sd

On 2006-11-29 07:37:48 -0600, "Arsalan Ahmad" <arsal__@xxxxxxxxxxxx> said:

My question was that Iperf is showing less UDP throughput over physical interface (around 550Mbps for 15 parallel streams) while for virtual NIC it is showing greater throughput (around 1Gbps for 15 parallel streams). So why is this? because what I thought that UDP thoughput should be high in case of physical interface which is not I am getting. And how Nagle's. Turn on NODELAY for TCP can effect UDP throughput performance?

Thanks,

Arsalan

"Steve Dispensa" <dispensa@xx>
wrote in message
news:2006112900064250073-dispensa@xx

On 2006-11-28 04:03:23 -0600, "Arsalan Ahmad"
<arsal__@xxxxxxxxxxxx> said:

Re: A question regarding MTU: how it can effect TCP performance + other queries

Hello all,

I have been doing some throughput tests with Iperf for my VNIC and physical NIC. I connect two P-4 3.4 GHz computers with 1Gbps NIC cards together via cross-over cable and then performed my experiments.

What I found that for TCP, throughput that I am getting on physical interface is around 900 Mbps and on virtual interface i am getting around 90 Mbps (only 10% of the physical one). All the tests were performed with Iperf with 15 parallel streams.

However, for UDP its totally different. With Iperf with 15 parallel streams, I am getting only around 500 Mbps throughput on physical interface while approximately 1000Mbps (1Gbps) on virtual interface. I am unable to figure out why am I getting this strange behaviour. Can you guess?

Yes, I can guess; are you tunneling the packets? Do you have Nagle's turned on? Did you do the math on the bandwidth-delay product and make your TCP window was set appropriately?

My virtual NIC pass all the packets to a user-level application which then send all the IP packets of virtual interface to user-level application running on the other computer. The communication between the two user-level application is via TCP.

Windows is sensitive to the exact structure of your sockets code. To get the best possible performance from the OS, you may want to look into using I/O Completion for your tunnel code on both ends, and make sure you keep enough buffers pending such that the protocol is never waiting on you.

So, for TCP packets over virtual interface, I am sending a TCP packet encapsulated within another TCP packet when passed to physical interface, while for UDP I am

Re: A question regarding MTU: how it can effect TCP performance + other queries

sending UDP packet encapsulated within
TCP packet when passed to physical
interface.

Given this, I'm voting for Nagle's. Turn on NODELAY and
re-test.

-sd