

## Re: Consistent performance issues at high bandwidths, UDP.

---

*Source:*

<http://www.tech-archive.net/Archive/Development/microsoft.public.win32.programmer.networks/2008-11/msg00000>

---

- *From:* "m" <[m@xxx](mailto:m@xxx)>
  - *Date:* Fri, 31 Oct 2008 21:03:58 -0400
- 

IIRC if multiple buffers are passed to WSASendTo, they are assembled into a single packet. The advantage of using WSASendTo with UDP is not the use of gather IO so much as the use of overlapped IO. Using overlapped IO may require a major change in your application, but it has the following two advantages:

- 1) reduced load on the OS because fewer threads are required. A single thread can do the work of many that are waiting on blocking IO calls. This saves memory and reduces context switches.
- 2) since packets are 'posted' to the IO subsystem, instead of sent one at a time, where batching operations makes sense the driver can do it. Of course if the drivers don't, there is nothing you can do about it – but at least you gave them a chance!

I agree that this problem is likely related to a hardware / driver issue and once that is resolved, the problem will likely go away. My comments on the use of overlapped IO merely constitute advice on good programming practice for Windows ;)

"Jason Cipriani" <[JasonCipriani@xxxxxxxxxxxxxxxxxxxxxxxxxxxx](mailto:JasonCipriani@xxxxxxxxxxxxxxxxxxxxxxxxxxxx)> wrote in message <news:83614550-8BFC-4CE2-BA06-A7D3BD517496@xxxxxxxxxxxxxxxxxxxx>

"m" wrote:

How are you sending the packets? Calling WSASendTo and using OVERLAPPED IO?

No, I am just calling sendto(), no overlapped IO. For the 927 and 59 socket setup, each socket is connect()ed first. For the 1 socket setup, I pass the address through sendto().

Re: Consistent performance issues at high bandwidths, UDP.

In WSASendTo is each WSABUF sent as a separate UDP packet?

It seems like a lot of overhead if I'm just sending 174 bytes with each call to WSASendTo but maybe I'm wrong. I'll try it.

I would expect the best performance from a single socket on each host since all creating multiple sockets should do is use more non-paged memory since the hardware can only send as fast as it can send regardless of how many endpoints are trying. If you see a significant change in performance by varying the number of sockets, then it is not just a hardware problem

I feel like it's an issue with the network driver. Hopefully the results of testing on other machines gets back to me soon.

Thanks,  
Jason

"Jason Cipriani" <JasonCipriani@xxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message [news:03C35CF4-650F-48FC-A1A8-4EE6251CC429@xxxxxxxxxxxxxxxxxxxx](mailto:news:03C35CF4-650F-48FC-A1A8-4EE6251CC429@xxxxxxxxxxxxxxxxxxxx)

"m" wrote:

Your seeing a performance impact from changing the number of sockets on the SENDING side when sending to the same host from all sockets?

Almost. Yes, I am seeing a distinct performance impact from changing the number of sockets on the sending side. However, I am not

Re: Consistent performance issues at high bandwidths, UDP.

sending to the  
same  
host from all sockets. There are 59 devices on the LAN, each  
with their  
own  
IP address.

In the configuration with 927 sockets on the sending side,  
there are  
14–16  
sockets per destination. No socket sends to more than one  
destination.

In the configuration with 59 sockets on the sending side,  
there is 1  
socket  
per destination. No socket sends to more than one  
destination.

In the configuration with 1 socket on the sending side, that 1  
socket  
is  
used to send to all 59 destinations.

Thanks again for your replies,  
Jason

"Jason Cipriani"  
<JasonCipriani@xxxxxxxxxxxxxxxxxxxxxxxxxxxx>  
wrote in  
message  
[news:A6BD1237-1258-4EEB-B581-D3C7FA8A3373@xxxxxxxxxxxxxxxxxxxx](mailto:news:A6BD1237-1258-4EEB-B581-D3C7FA8A3373@xxxxxxxxxxxxxxxxxxxx)

Thanks for your reply.

"m" wrote:

What is  
different  
between  
config 1 &  
2?

The only difference is the

## Re: Consistent performance issues at high bandwidths, UDP.

number of sockets and how many packets I'm sending with each socket. I also tried a test with a single socket, sending all data with just that socket. In that test each frame took 40ms to send, and every 3.1 seconds a frame took 3.8 seconds to send (wow).

So, with the same frame size (~160kB), the same number of packets per frame (927 UDP packets, 174 bytes each), I can see distinctly different patterns depending on how many sockets I split those packets between:

927 sockets x 1 packet:  
10ms/frame, 0.7 seconds normal, 1.8 second delays

59 sockets x ~16 packets:  
20ms/frame, 0.1 seconds normal, 0.13 second delays

1 socket x 927 packets:  
40ms/frame, 3.1 seconds normal, 3.8 second delays.

I also found something interesting, if I increase the packet size to 512 bytes (over a single socket, still testing the same amount of data),

Re: Consistent performance issues at high bandwidths, UDP.

while  
the frame rate was slow  
(took 150ms to send 160kB  
for some reason),  
it  
was  
consistent and there were no  
regular delays. I can't  
explain that.  
Also  
that  
test is not that useful since  
increasing packet sizes is  
not an  
option...  
the  
protocol used to talk to the  
network devices relies on  
UDP packet  
boundaries.

The most  
likely cause  
is the NIC  
driver /  
hardware is  
not able to  
keep  
up  
with your  
sending rate  
and the IO  
is stalling  
until it can  
catch  
up.  
Lots  
of gigabit  
NICs,  
especially  
early ones,  
cannot  
achieve full  
gigabit  
speed.  
Another  
strong  
possibility is  
that the bus

Re: Consistent performance issues at high bandwidths, UDP.

/ bridge that  
feeds the  
NIC  
is  
being  
saturated.

That is how I feel, too. I  
can't think of any other good  
explanation.  
Testing with other hardware  
can help verify this.

In either  
case, test  
with  
alternate  
hardware  
and see  
what  
happens.

Tests on other hardware  
should be run today or  
tomorrow, when I get  
the  
results back hopefully it will  
answer some questions.

The are  
also some  
registry  
parameters  
that control  
the IP stack  
behaviour  
in  
some  
situations  
that might  
affect you  
and you can  
reduce host  
load by  
using  
overalpped  
IO & scatter

Re: Consistent performance issues at high bandwidths, UDP.

/ gather IO

I may have to start digging  
in to registry tweaks but I'd  
rather get  
as  
close as I can through less  
hacky means first. With luck  
we'll  
find  
out  
it's just unique to the  
machine we were running it  
on.

Machine causing problems  
was a MacBook Pro, Core 2  
Duo T7700 @ 2.6  
(i  
think)  
GHz. Not sure what NIC.

Next tests will be on a Mac  
Mini and an as-of-yet  
unidentified PC. I  
have  
to  
get the hardware specs from  
the people on site.

Development machine,  
which never showed any  
problems, was a Thinkpad  
T60  
Core Duo T2600 @  
2.16GHz, but the network  
setup was significantly  
different  
(Intel 3945ABG Wireless,  
not really relevant). If none  
of the other  
machines  
work I'll have to take the  
Thinkpad there and test with  
that, I  
guess.  
Remote  
work is such a pain.

Thanks,

Re: Consistent performance issues at high bandwidths, UDP.

Jason

"Jason  
Cipriani"  
<JasonCipriani@xxxxxxxxxxxxxxxxxxxxxxxxxxxx>  
wrote in  
message  
[news:29C16DAE-BFB7-44C1-A97B-D2E3E60C4295@xxxxxxxxxx](mailto:news:29C16DAE-BFB7-44C1-A97B-D2E3E60C4295@xxxxxxxxxx)

I  
wrote:

In  
configuration  
1,  
when  
sending  
frames  
as  
fast  
as  
possible,  
each  
frame  
generally  
takes  
about  
10ms  
to  
send.  
The  
real  
problem  
is  
every  
0.7  
seconds  
(~65  
frames,  
~10MB),  
a  
frame  
takes  
a  
whopping  
1.8  
seconds  
to  
send.  
I

Re: Consistent performance issues at high bandwidths, UDP.

can't  
explain  
this  
massive  
delay,  
but  
it  
is  
at  
extremely  
regular  
intervals.

I  
should  
note  
that  
I  
sent  
data  
as  
fast  
as  
possible  
in  
the  
test  
application,  
but  
in  
the  
real  
application  
even  
when  
I  
limit  
as  
low  
as  
30FPS,  
I  
still  
see  
the  
same  
problem,  
except  
the  
intervals

Re: Consistent performance issues at high bandwidths, UDP.

are  
longer.  
For  
example,  
at  
30FPS  
(approx  
average  
bandwidth  
30–40mbps)  
it  
takes  
10–20ms  
to  
send  
a  
frame  
and  
I  
still  
see  
the  
1.8  
second  
delays,  
but  
it  
happens  
every  
3  
or  
4  
seconds  
rather  
than  
every  
0.7  
seconds.

Thanks,  
Jason

Original  
message:

"Jason  
Cipriani"  
wrote:

Re: Consistent performance issues at high bandwidths, UDP.

I  
have  
an  
application  
that  
streams  
data  
over  
the  
network  
at  
high  
bandwidths  
that  
is  
having  
a  
lot  
of  
performance  
issues.

I  
narrowed  
it  
down  
to  
a  
very  
minimal  
case,  
and  
am  
leaving  
out  
all  
the  
rest  
of  
the  
details  
here:

The  
application  
is  
a  
multimedia  
application  
that  
must  
stream

Re: Consistent performance issues at high bandwidths, UDP.

frames  
of  
data  
using  
UDP  
at  
consistent  
frame  
rates.  
Each  
frame  
is  
about  
160kB  
split  
into  
927  
UDP  
packets  
(of  
174  
bytes  
each),  
with  
each  
packet  
sent  
to  
one  
of  
59  
destination  
devices  
(max  
16  
packets  
per  
destination).  
It  
is  
a  
dedicated  
gigabit  
LAN  
and  
I've  
verified  
that  
a  
gigabit  
connection

Re: Consistent performance issues at high bandwidths, UDP.

was  
established  
between  
all  
network  
devices.  
The  
performance  
issue  
I'm  
experiencing  
is  
highly  
inconsistent  
frame  
rates  
caused  
by  
a  
periodic  
delay  
in  
the  
call  
to  
sendto()  
at  
extremely  
regular  
intervals.  
For  
the  
purposes  
of  
testing  
I've  
turned  
off  
the  
frame  
rate  
limiter  
and  
am  
sending  
data  
as  
fast  
as  
possible  
(with

Re: Consistent performance issues at high bandwidths, UDP.

the  
limiter  
set  
to  
60FPS  
average  
bandwidth  
is  
roughly  
75mbps).

I  
have  
two  
configurations  
that  
I've  
tested:

1)  
927  
sockets.  
Each  
frame  
sends  
174  
bytes  
over  
each  
of  
these  
sockets.

2)  
59  
sockets.  
Each  
frame  
sends  
2436  
to  
2784  
bytes  
over  
each  
of  
these  
sockets.

At  
gigabit

Re: Consistent performance issues at high bandwidths, UDP.

speed  
each  
frame  
\*should\*  
take  
about  
1.2ms  
to  
send  
(160kB  
/  
1000mb  
~=  
0.0012,  
mind  
your  
bits  
and  
bytes).  
In  
practice  
I  
can't  
even  
get  
close  
to  
1.2ms  
per  
frame  
(see  
below).

In  
configuration  
1,  
when  
sending  
frames  
as  
fast  
as  
possible,  
each  
frame  
generally  
takes  
about  
10ms  
to  
send.

Re: Consistent performance issues at high bandwidths, UDP.

The  
real  
problem  
is  
every  
0.7  
seconds  
(~65  
frames,  
~10MB),  
a  
frame  
takes  
a  
whopping  
1.8  
seconds  
to  
send.  
I  
can't  
explain  
this  
massive  
delay,  
but  
it  
is  
at  
extremely  
regular  
intervals.

In  
configuration  
2,  
the  
same  
problem  
exists  
except  
each  
frame  
generally  
takes  
about  
20ms  
to  
send,  
and  
every

Re: Consistent performance issues at high bandwidths, UDP.

0.1  
seconds  
(~6  
frames,  
~1MB)

,  
a  
frame  
takes  
0.13  
seconds  
to  
send.  
Again,  
it  
is  
at  
extremely  
regular  
intervals.

I  
had  
thought  
that  
something  
else  
on  
the  
machine  
was  
interfering  
with  
network  
communications  
(checking  
usual  
culprits,  
making  
sure  
wireless  
networking  
disabled,  
etc.)  
but  
the  
thing  
is  
configuration  
1  
and

Re: Consistent performance issues at high bandwidths, UDP.

2  
experience  
distinctly  
different  
patterns  
of  
delay.  
One  
is  
1.8  
seconds  
every  
0.7  
seconds,  
the  
other  
is  
0.13  
seconds  
every  
0.1  
seconds.

Another  
piece  
of  
information  
that  
may  
be  
important  
is  
the  
delay  
does  
not  
seem  
to  
happen  
when  
the  
network  
cable  
is  
unplugged.  
This,  
of  
course,  
strongly  
suggests  
some

Re: Consistent performance issues at high bandwidths, UDP.

hardware  
issues  
at  
some  
point  
in  
the  
network.  
However,  
I  
have  
not  
been  
able  
to  
verify  
this  
yet  
as  
the  
software  
is  
for  
a  
client  
running  
things  
remotely,  
and  
communication  
is  
sometimes  
difficult  
(unfortunately,  
I  
can't  
be  
at  
the  
site  
to  
witness  
the  
actual  
problem,  
which  
does  
not  
occur  
on  
the

Re: Consistent performance issues at high bandwidths, UDP.

development  
machine  
:-(  
).

My  
test  
application  
does  
not  
do  
anything  
exotic.  
It  
simply  
creates  
the  
sockets  
like  
so:

```
socket(AF_INET,  
SOCK_DGRAM,  
IPPROTO_UDP);
```

And  
sends  
data  
packets  
to  
their  
appropriate  
destinations:

```
sendto(p.sock,  
data,  
datalen,  
0,  
paddr,  
addrlen);
```

All  
calls  
to  
sendto()  
are  
succeeding,  
all  
data  
is  
sent

Re: Consistent performance issues at high bandwidths, UDP.

and  
I  
have  
verified  
that  
it  
is  
correctly  
received  
by  
the  
remote  
devices.

1.  
What  
is  
causing  
the  
delays  
at  
regular  
intervals?

2.  
Why  
does  
the  
delay  
time  
and  
interval  
length  
depend  
on  
how  
many  
sockets  
I  
have  
open  
and/or  
how  
much  
data  
I'm  
sending  
per  
socket  
(note  
that  
in

Re: Consistent performance issues at high bandwidths, UDP.

both  
configurations  
each  
frame  
is  
the  
same  
amount  
of  
data  
total)?  
3.  
How  
can  
I  
troubleshoot  
/  
solve  
this?  
Are  
there  
some  
kind  
of  
socket  
options  
I  
can  
set  
to  
improve  
performance  
and  
consistency?

It's  
frustrating  
because  
the  
software  
goes  
live  
soon,  
and  
these  
problems  
didn't  
occur  
until  
the  
software

Re: Consistent performance issues at high bandwidths, UDP.

was  
run  
on  
site.  
Any  
advice,  
hints,  
info  
would  
be  
greatly  
appreciated.