White Paper

TCP Network Latency and Throughput

Or 'Why your customer doesn't receive the Throughput they paid for'

Introduction

A frequent complaint raised by customers of a particular service is that they have been sold a particular rate or bandwidth but when the customer downloads one of the many applications that can check the actual throughput they then find that they are only receiving a fraction of the bandwidth that they expected. This leads to complaints, finger pointing and customer churn as the disgruntled customer seeks a better available bandwidth elsewhere. In reality, the available bandwidth may be no better elsewhere but by that time it's too late. The customer has gone and probably won't be coming back. One of the reasons is that throughput is adversely affected by latency and packet loss. This white paper explains why this is the case and shows the effect of latency and packet loss on service throughput.

TCP versus UDP Background

Traditionally, video was carried over User Datagram Protocol (UDP). UDP is a connectionless protocol. Packets are pushed to the destination with no regard for packet loss or packet order. The problem with UDP is that there is no guarantee that the packets will reach the end destination.

Now, most of the streaming video on the internet uses Transmission Control Protocol (TCP). Video streaming services or applications such as Netflix, Amazon and YouTube all use TCP for streaming video. These services account for the bulk of video streaming in the Western world. TCP was designed to ensure delivery of all packets and minimize packet loss. It is a connection oriented protocol which ensures quality of service by re-transmitting packets until all packets are received correctly.

TCP Network Latency and Throughput

Or 'Why your customer doesn't receive the Throughput they paid for'

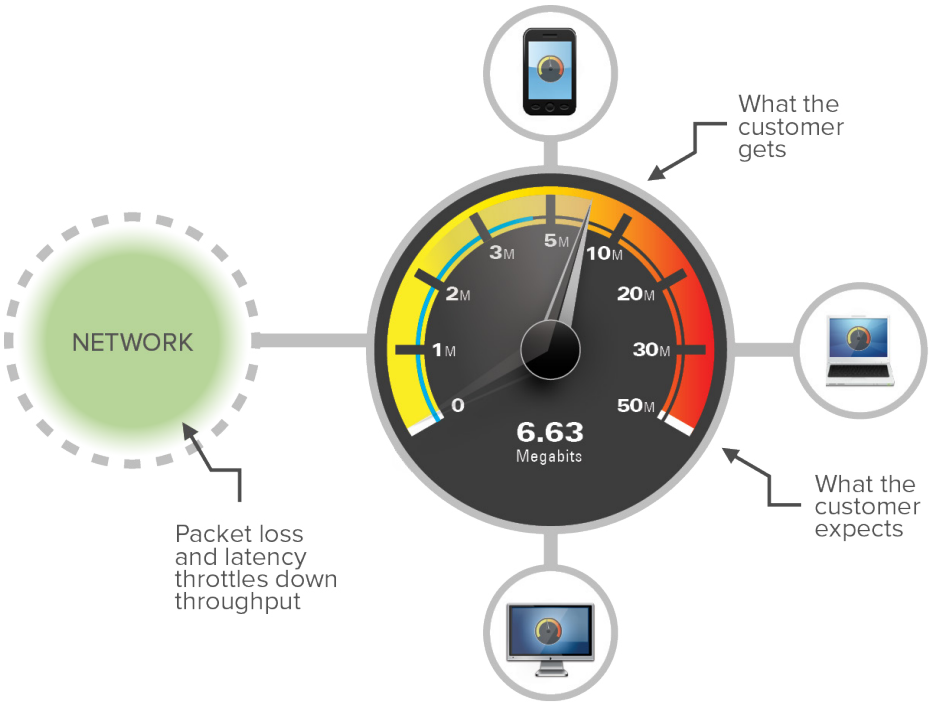

An unintended consequence is that TCP is extremely sensitive to network delay as it is a reliable protocol, which means that significant re-transmission can occur, and this along with the TCP handshake process is made worse by long transmission delays.

One of the key reasons for the switch form UDP to TCP is that content providers want TCP for reliability/quality of picture etc. Service providers believe that today's networks are much higher bandwidth and therefore able to cope with longer delays and re-transmissions. However, throughput can be much less than the provisioned capacity of the link. The issues affecting TCP throughput are very often out with the network operator's control but as these issues are often significant it's important that each network operator can be certain that their network is not causing throughput issues thereby reducing time consuming finger pointing and troubleshooting. Traditional testing is of the RFC-2544 variety i.e. load testing which does not capture problems associated with latency and buffer size. Studies have shown that a correctly configured TCP network can perform up to 10 times faster*.

** (Pittsburgh Supercomputing Center and https://en.wikipedia.org/wiki/TCP_tuning#cite_note-1)*

What Causes Network Latency?

- Propagation Delay: This is the time it takes to send a packet thousands of kilometres. For example, the propagation delay from New York to London, via optical fiber, is approximately 26ms.
- Routing/Switching Latency: An IP router adds approximately 200us per device. This would add another 1.4ms to the above New York to London link based on assumptions that there is one router per 800km.
- Queuing Latency: This is the amount of time a packet can spend in a queue awaiting transmission due to congestion after routing/switching. This can add another 20ms of latency.

Why is Network Latency such an issue for TCP?

As discussed above, Latency matters for reliable connection oriented protocols such as TCP. TCP is a session-based protocol that ensures packet delivery, and minimizes packet loss. As a result throughput can often be much lower than provisioned capacity. Unlike UDP, TCP requires the receiving end to acknowledge the successful receipt of packets. However, it is inefficient to wait for successful acknowledgement of every individual packet before sending the next packet. This is especially true over long distance networks with long latency. TCP uses the notion of a "window size" to determine how many packets are likely to be sent without loss. TCP waits on an ACK handshake to establish connection. If packets are missing TCP uses the ACK to request re-transmission. TCP uses a "windowing" technique by waiting on ACK to slowly increase the transmission rate to match the available bandwidth. The transmission rate is rapidly reduced when packets are lost or not acknowledged within an expected time frame. As a result, throughput is greatly affected by the latency and also by any packet loss in the link.

Network Latency can be such a problem that some satellite links implement TCP spoofing to terminate the link at the Tx site to pretend that the link has lower delay. This only works if the link is uncongested and error free.

High Speed Network—Effect of Buffering

TCP performs badly in the presence of latency and tiny amounts of packet loss. In addition, TCP windowing and latency can be an even bigger problem in high speed networks. At 100GbE, the TCP windowing buffer fills up much more quickly than at 10GbE or 40GbE as the interface input and output buffers tend not to be scaled up sufficiently to cope with higher speed traffic. This means (if you set the same wait time) that you get a throughput drop-off at even lower latencies than at 10GbE or 40GbE.

Some switches are now combating this issue by adding many GBytes worth of packet buffer memory. This is many orders of magnitude greater than the amount of buffering in traditional switch/routers. This helps minimize packet loss, preventing re-transmission delays and the build up of huge latencies.

As a result, different switches/routers will perform markedly differently within a network and there can be a huge resultant difference to the observed practical traffic throughput within the network.

TCP Network Latency and Throughput

Or 'Why your customer doesn't receive the Throughput they paid for'

TCP Windowing (effect of latency and packet loss on throughput)

The figure opposite shows the TCP windowing 'Slow Start' in action. It is inefficient to send a single frame and then wait for acknowledgement before sending another. So, TCP 'Slow Start' groups the frames together into a 'congestion window' (CWND) to increase throughput. The sender sends out one packet then waits for ACK before doubling the number of packets then repeats. When the TCP window registers that a packet is lost (for example due to congestion) then the packet window size is halved. The slow start kicks in again but this time it will not double the rate but increase more slowly by a segment at a time. This is done to detect optimal throughput. The whole process is very sensitive to delay through the network. A simple rule in a zero packet loss environment is that:

TCP Rate = MSS/RTD

where:

MSS = Maximum Segment Size

RTD = Round Trip Delay

TCP is also bursty in nature and spikes exceeding the CIR (Committed Information Rate) can equate to lost packets, leading to re-transmission with a smaller window size and consequently reduced throughput. There is a double hit here as the "ACK" process leads to more messages going back and forth which is the very process that takes longer when there is more latency in the network.

TCP Slow Start

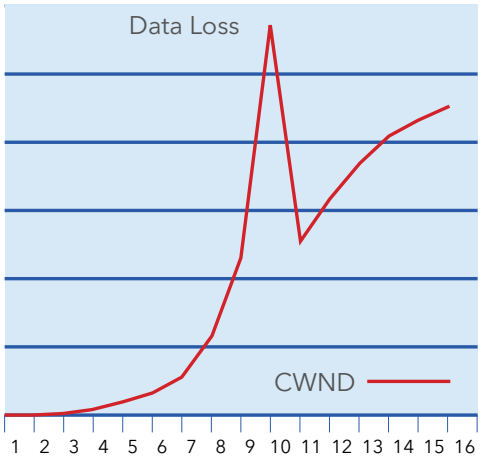

- Slow Start uses a congestion window (CWND)
- CWND is initialized to the segment size announced by the other end of a new connection

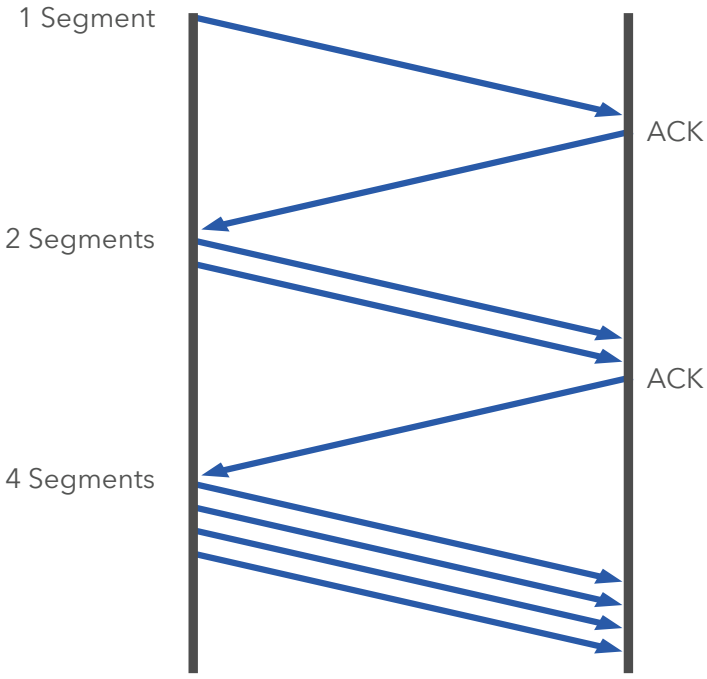

How is TCP effected by Latency?

The graph below shows the effect of different amounts of latency on throughput. The left hand side shows the effect of "windowing" with 30ms round trip delay (RTD). Packets are continuously halted as the window size is full and the sender needs to see ACK. 30ms could be the typical RTD found across the USA. On the right hand side, with the exact same bandwidth, but with the latency increased to 256ms RTD. This could represent a typical RTD for a satellite link. You can see that the throughput drops dramatically to almost zero even though the bandwidth has not decreased.

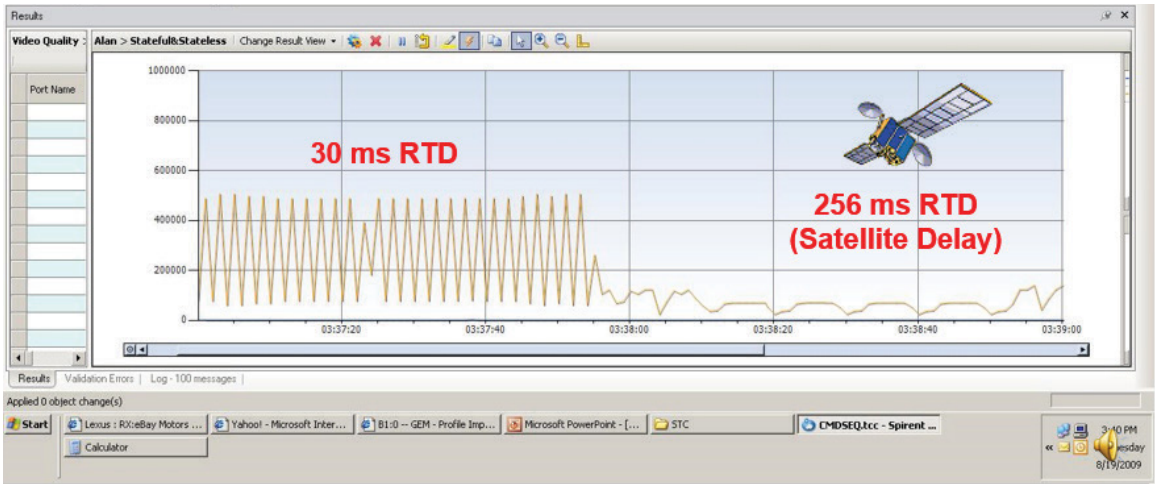

How is TCP effected by Packet Loss and Latency?

The graph below provides an indication of the effect of 0% packet loss and no latency on the left with almost full bandwidth utilization. With 3% packet loss, congestion control & slow start process TCP struggles to fill the link even with no latency. Adding 30ms latency combined with 3% packet loss further impacts throughput leading to low link utilization. Adding 256ms latency combined with 3% packet loss completely kills throughput. Therefore, the effect of packet loss and latency on throughput can be extremely severe.

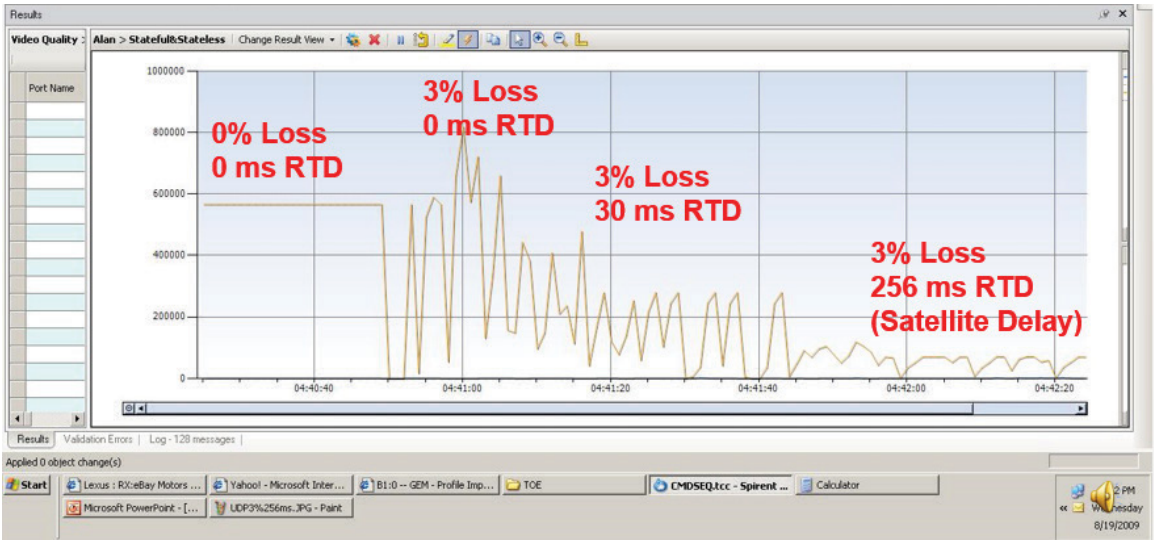

TCP Network Latency and Throughput

Or 'Why your customer doesn't receive the Throughput they paid for'

How is UDP effected by Latency?

The next two figures show that unlike TCP, UDP is unaffected by latency and packet loss. UDP will continue to transmit at line rate regardless of latency in the connection. As mentioned above, TCP is a reliable connection oriented protocol and UDP is not. UDP is connectionless with no reliability built-in and unlike TCP there is no transmission window that can close. Therefore, UDP is completely unaffected by latency (there is no ACK process) as there is no re-transmission.

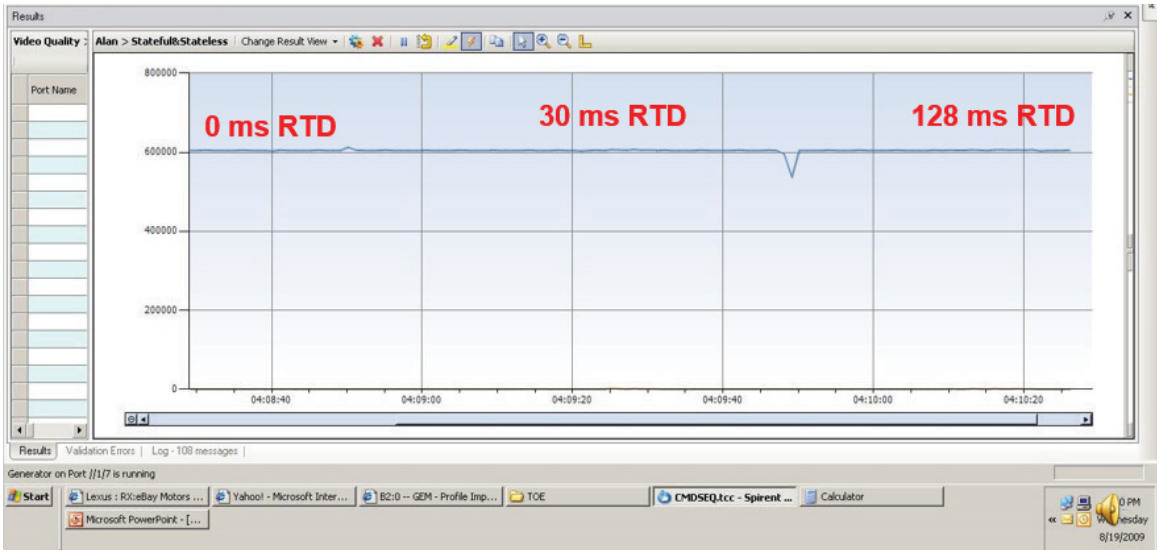

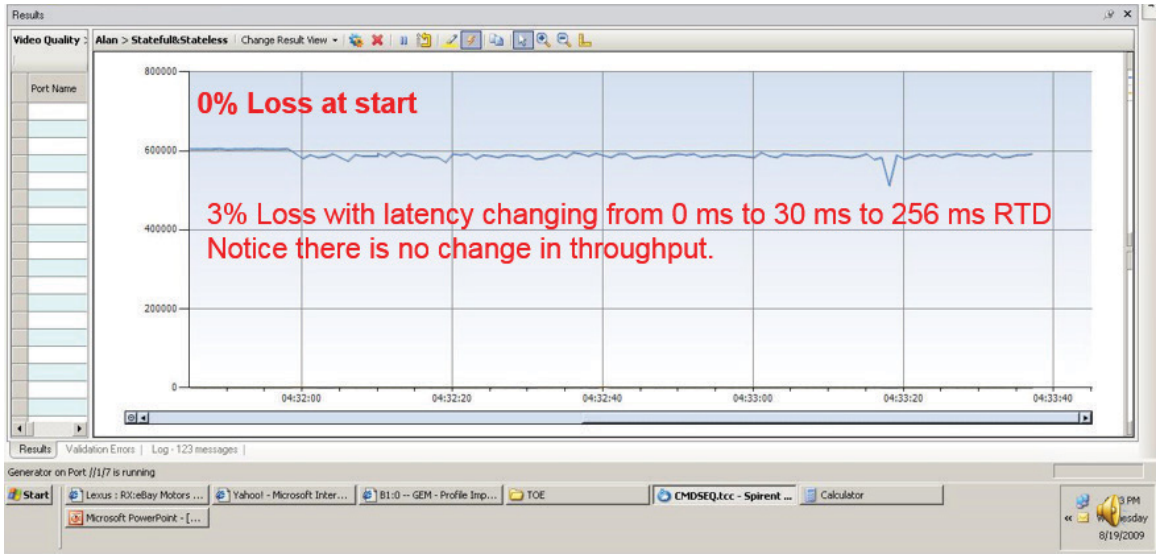

Simulating Latency

In order to know the performance of your TCP network or application/service, it's important to be able to simulate the network in your laboratory. Service providers will need to be able to add delay to test the impact on throughput and will need the ability to error packets and measure the impact on throughput due to packet loss. It is possible to use drums of fiber to simulate network delay but there are practical considerations with this approach.

- **Drums of fiber are too large:** Drums of fiber are very large, heavy and unwieldy. 200km of fiber is needed for 1ms of delay (5ns/metre). So, to simulate the latency from say Los Angeles to New York which is a distance of 4,500km then 22.5ms of delay is required. That's 22.5 drums of 200km fiber.
- **Amplifiers are needed:** Fiber drums need amplification. Single mode 1310nm fiber has a fiber loss of 0.4 dB per km which equates to 80dB loss for 200km. The specified supported range for a single 100GbE LR4 interface is only 100km. Therefore, optical amplifiers are needed for >100km. Simulating longer links such as the 4,500km link between LA and New York needs a lot of amplification. That's a lot more cost and complexity.
- **Lack of control:** If a test fails with 200km of fiber (1ms latency) it's not easy to finely adjust fiber length to find point of failure. Splicing fiber is one method but that's a time consuming practice and is a one time event unless you want to purchase more fiber. Adding additional drums to increase latency is also problematic. It doesn't lend itself to automation either. You may want or need to control latency via software that either adds or decreases the latency in small incremental steps.

TCP Network Latency and Throughput

Or 'Why your customer doesn't receive the Throughput they paid for'

About Spirent Communications

Spirent Communications (LSE: SPT) is a global leader with deep expertise and decades of experience in testing, assurance, analytics, and security, serving developers, service providers, and enterprise networks.

We help bring clarity to increasingly complex technological and business challenges.

Spirent's customers have made a promise to their customers to deliver superior performance. Spirent assures that those promises are fulfilled.

For more information, visit: www.spirent.com

Conclusion

Streaming video is now predominantly carried by TCP (Amazon, Netflix and YouTube). Network latency which is caused by Propagation Delay, Switching/Routing and Buffering/Queuing is a big issue for TCP networks. Latency kills TCP throughput as TCP waits for ACK before sending the next packets. This process is greatly affected by delay. Re-transmission is also needed if packets are dropped/lost and this process is also severely impacted by latency. Latency is an even bigger problem in high speed networks as the TCP windowing buffer fills up much more quickly.

So, it's important to be able to check the effect of latency and dropped packets by emulating the network and or service in the lab. You need to be able to add delay to test the impact of latency on your network or service's throughput. Plus, you need to be able to error packets and measure the impact on throughput due to packet loss.

- TCP re-transmits packets to ensure packet delivery and therefore Quality of Service.
- As a result, TCP is very sensitive to Network Latency and Packet Loss.
- Packet Loss and Latency can completely kill TCP throughput.
- The biggest complaint of end users is being promised xGb/s of service and receiving significantly less.
- Check the effect of latency and dropped packets by emulating the network and or service in the lab.

spirent™
Promise. Assured.

Contact Us

For more information, call your Spirent sales representative or visit us on the web at www.spirent.com/ContactSpirent.

www.spirent.com

Americas 1-800-SPIRENT
+1-800-774-7368 | sales@spirent.com

Europe and the Middle East
+44 (0) 1293 767979 | emeainfo@spirent.com

Asia and the Pacific
+86-10-8518-2539 | salesasia@spirent.com

Rev B | 10/19