

The Obsession with Bandwidth (Part 2)

“The Obsession with Bandwidth (Part 1)” white paper discussed the broader issues related to the popular misnomer that bandwidth is a *speed*, when in fact it is only a *rate*. Part 2 extends this discussion to introduce the importance of network quality and the user experience. Reading Part 1 is recommended for better understanding of key network measurement concepts discussed further in Part 2.

The demand for bandwidth has persistently increased year over year, driven by a massive increase in network dependent applications. What started as Morse code is now changing the very fabric of human culture. Online applications affect how we work, how we play, how we shop, how we think and how we communicate, socially and formally. Information is truly anywhere and everywhere.

High capacity network pipes have evolved to accommodate higher and higher demands. The high costs associated with high capacity pipes are justified against opportunity because high capacity pipes are expected to deliver business and social application services to hundreds, thousands and even millions of users, locally and globally. Simply put, networks are driven by opportunity. After all, necessity is the mother of invention.

The successful delivery of an application is naturally a critical business concern because failure can and will impact a broad set of customers, directly or indirectly. Underwriting a good user experience is therefore a critical element, and few would argue that the qualification and assessment of a pipe is a fundamental process to underwrite the success of any opportunity. This is not as simple a task as one might think, but it would be fair to say, the bigger the pipe the more critical the assessment.

Real-life case study

As an example, consider a real-life case study. The pipe in question was a dedicated 1Gbps Ethernet connection with an end-to-end latency of 21ms. The pipe connected a high school to an application service provider. Based on popular belief that bits per second is a speed, the 1Gbps was considered to be suitably “fast” by both the customer and the provider. Unfortunately, reality did not meet expectation and the connection performed poorly. Users would often experience unworkable application response times and worse, disconnects. The poor user experience naturally resulted in user complaints and the complaints went unresolved...

Running a connection speed test application reported a result close to 1Gbps. Based on this test result the provider naturally reported that the connection as ‘good’ and redirected blame to the last mile, namely the school’s internal network.

So why was there a problem and how do we resolve it?

The popular online speed test applications delivered little value to the resolution process, in fact the tests only served to delay resolution because ownership of the issue was passed back and forth between provider and customer. This is a common situation for cloud applications, the main reason being the misnomer that bits per second defines a speed when in fact it does not. Attaining an expected bit rate does not in itself define how an application will perform for a single user, only the concurrency of users.

To answer the obvious question of ‘why there was a problem?’ required a detailed assessment of: 1) What the connection is expected to deliver, and 2) How well the connection delivers on that expectation. The expectation applies to all users and all applications. This is essentially a logical process of mapping supply to demand.

The Case Study Requirement

The 1Gbps pipe was established to deliver a wide range of education content, including statutory examinations, to 5,500 high school students. The school facilities would host about 2,200 students online at any one moment in time.

User Demand

To understand what the network is expected to deliver requires a clear understanding of the average demand from a single student, and moreover how the application data for a single student utilizes the network. Application data flow correlates directly to the application need and the experience delivered to the student.

As a rule of thumb, the smallest amount of data a single PC will place on a network is 0.52 million bits (about 45 packets, 64KB). A PC can send more data, however for typical web applications it is unlikely to send less. The data dispatched on the wire is referred to as the *unacknowledged window* because the sender of the data will not send more data until the data receiver acknowledges the data as received. *This acknowledgement cycle is the prime reason why the trip time of a connection is so material to the user experience.* Longer trips increase the time taken to complete the task.

Extrapolating this data value for 2,500 online students equates to a demand of $2,500 \times 0.52$ million bits or **1,300 million bits**.

Connection Supply

The *capacity* of a pipe is defined by the product of width and latency end-to-end. A pipe expressed as 1 Gigabit per second is the same as 1 million bits per millisecond. The pipe’s minimum trip time end-to-end for the case study was measured at 21ms, therefore the product of width and length defines the pipe capacity at 21×1 million bits or 21 million bits. The minimum peak demand of 1,300 million bits is therefore 62 times larger than the 21-million-bit capacity.

What does the differential between demand and capacity matter?

As with any supply and demand model, when resources demanded exceed resources supplied then delays are to be expected. Ask anyone who drives to work in the rush hour, or anyone who tries to order a drink in a crowded bar on a hot summer afternoon. No matter the circumstance, the user experience will come under threat when demand exceeds supply because delays are inevitable. If the delays climb too high then a wide range of quality issues will likely occur, this includes disconnects.

How a network copes with excess demand is singularly one of the most important characteristics that underwrites a stable network connection. *The inability to measure and assess delay variances is a core reason why popular speed test applications fail to identify and resolve issues that deliver a poor user experience.*

Understanding Delay

While accepting that 0.52 million bits represents the smallest unacknowledged payload for a single student, this does not represent the total demand needed to complete a student's application transaction.

For the case study the average application transaction size per student was assessed at 1,200 Kilobytes, which is 9.6 million bits. Therefore, a transaction demand of 9.6 million bits at 0.52Mb per trip will consume 19 trips to complete one transaction. This means the time taken to complete a transaction is around 400ms per student (19 trips x 21ms per trip, assuming latency remains constant!). This magic number means the pipe should be able to satisfy the demand of about 100 students per second, factoring in network protocol overheads. It is actually a bit less than a hundred but 100 is a nice round number.

The total application trip time demand on the pipe for all 2,500 students at 1 second per 100 students is therefore 25 seconds. Unfortunately, the TCP protocol timeout is circa 15 seconds which means around 40% of students will risk encountering a delay that exceeds the timeout -- resulting in a disconnect and a bad experience.

Is latency variance a material threat to the application?

Unfortunately, 100-student per second assumptions are optimistically flawed. First, the assessment has assumed that running at the capacity of the pipe does not increase latency end-to-end. In reality congestion delay will climb beyond 21ms, by how much is a key question. Second, if demand peaks persistently exceed supply, packets will eventually be discarded causing more extensive delays as a result of packet recovery. This defines a quality problem.

Consider what happens when congestion increases the latency of the pipe by just 4ms. One student transaction of 1,200KB took 19 trips at 21ms consuming 400ms of time. At 25ms per trip this climbs to 475ms, reducing the number of concurrent users per second from the optimistic 100 to a more realistic 80. The total application trip time demand for all 2,500 students at 1 second per 80 students is now

2,500/80 or 31 seconds and the number of disconnects climbs from 40% to over 50%. This is quite an impact for just 4ms of time!

The effectiveness and efficiency of a pipe is defined by how a connection performs in relation to the trip time as student demand climbs. Connections with poor efficiency will increase latency to higher values than more efficient connections. In essence, the efficiency of a pipe defines how consistently the pipe is able to deliver data on-time (minimum latency) relative to the utilization of the pipe's capacity.

Congestion versus Quality

What are quality issues all about and why do they occur?

There are many factors that cause packet delay on a network and any event that causes a delay to packets is a potential threat to the user experience. Putting aside specific delays for application processing and the end-to-end distance that defines the minimum latency, data packet delays essentially come from two core network attributes: network congestion and packet quality.

Congestion

Congestion delay variations materialize from all the devices that comprise a network connection end-to-end. This includes the client and server devices along with all the other management devices in between, such as routers, switches and firewalls. Packets can only stop in a device (buffer) because there are no stopping places on the wire. *As buffers fill, the latency delay climbs. The extent by which latency climbs separates a good connection from bad one.* A perfect connection is one where latency does not climb regardless of demand. Not all routers are created equal.

Quality delay variations materialize when a packet is deemed unfit for its purpose. For example, a packet in transit can become corrupt, out-of-order or lost. When a data packet becomes compromised, recovery is essential to underwrite data integrity and allow the flow to continue. Quality delays can be extreme in terms of time compared to congestion delays. Larger delays create a domino effect because data is allowed to drain from the pipe. One of the more severe penalties is incurred for a lost packet because all packets following the lost packet cannot be processed out of order. If packet(s) recovery time exceeds the trip time, the pipe will drain. Unfortunately, a packet does not have to actually be lost - to be lost, it just has to be delayed long enough to be considered lost. In this situation packets are unnecessarily retransmitted, further adding to delay and also bandwidth demand, a digital double whammy. The inability to detect good packets from bad packets is also the reason why monitoring solutions add little value and popular speed tests report misleading results

The discarding of packets has become a common network policy event invoked by service providers to manage the utilization of a network. Unfortunately, regulatory policies function similar to traffic lights on a road and threaten the user experience by adding delay and reducing demand. Because recovery delays are excessive, the majority of network problems that impact the application experience result more from poorly implemented policies than congestion events. In short, *the provisioning of network traffic is one of the most significant threats to the user experience.* Consider, in the case of the school

study, if a small 4ms congestion delay increased disconnects by 25%, what outcome would a 400ms quality policy delay have created?

How should a pipe be assessed to underwrite a good user experience?

A thorough assessment of a network connection for data performance must address packet quality, efficiency, throughput, and delay variance. Randomly throwing data at a pipe to measure the 'bps' achieved will not deliver useful results and only serve to mislead. There are three parts to a robust test model for data performance.

1. Equilibrium

The equilibrium throughput of a pipe defines a vital baseline measurement. The equilibrium result is defined by the product of 'trips per second' x 'unacknowledged data size'. To make this point clear, if the trip time of a connection under test was one whole second and the unacknowledged data window was 100 bits, then the equilibrium throughput test result must achieve 100bps (including overhead). If the test result is not 100bps, then clearly the connection has failed to deliver all the data at the trip latency -- the true definition of slow. The percentage variance from equilibrium defines the scale of connection's problem, and the difference between the test result and the equilibrium result represents the negative variation in latency.

2. Efficiency

Efficiency defines the ability for a pipe to maintain its trip latency end-to-end when demand is at or above capacity. A pipe's ability to maintain the equilibrium result, defined by the latency and the concurrency of the test sessions, up to the limit of the bandwidth, defines the efficiency graph.

In the 100-bit example discussed above, running 2 x 100bit sessions should reach 200bps or the bandwidth limit -- whichever is the smaller. In reality, as traffic climbs efficiency drops -- just like rush hour traffic. The percentage drop defines how well a connection will cope with peak demands. A very good result is 90% plus and a bad result starts below 70%. Fig 1. Shows a 30Mbps connection that drops efficiency to 70% at capacity, a loss of about 9Mbps. The optimum result was 80% at 16Mbps. A well-structured connection should be capable of reaching the bandwidth limit at an equilibrium value of +90%.

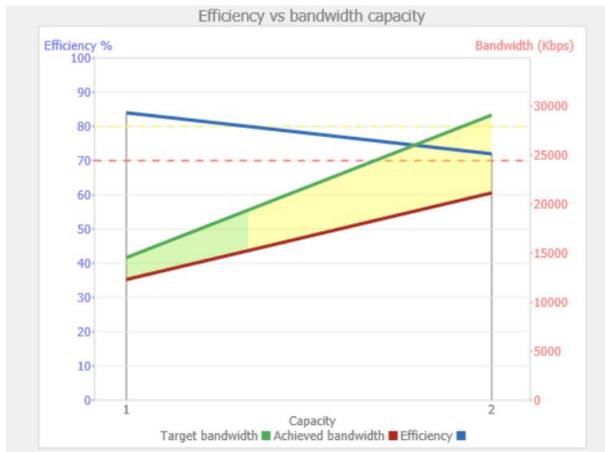


Fig 1. Efficiency vs. Bandwidth

3. Testing Limits

Testing limits of a connection is important and the final piece of the puzzle. To be reliable, a network test must not stop at bandwidth attainment, the truth is quite the opposite. *To underwrite a consistent user experience requires that a connection is able to sustain its bandwidth supply even when it is under a demand that exceeds the bandwidth supply.* Measuring 1Gbps connections from a single 1Gbps workstation will be unlikely to uncover quality issues associated with demand simply because demand cannot exceed supply. In the real world, as is the case with 2,500 students, peak demand is defined by the potential combined value of all students which will always exceed supply on a consistent basis. As bandwidths climb to 10Gbps and beyond, ensuring demand is coordinated to exceed supply by a real-world margin is essential to the purpose and value delivered by the test process.

In the case study, each student transaction (9.6million bits) was able to consume half the supply capacity. This means the maximum peak client demand for all 2,500 students could in theory exceed 1,250 times the pipe’s capacity. How a provider’s regulatory policies manage the excess demand poses the most significant threat to the user experience. If the provider policies and buffers cannot sustain the concurrent demand then data will have to be discarded and quality penalties are incurred. This will cause retransmission events which further increases demand and aggravates the problem. Contrary to popular belief, in the world of data discarding packets to reduce demand actually increases demand.

A robust assessment test will quickly and accurately identify connections that threaten the user experience. Conversely, a test that does not consider connection equilibrium, efficiency and testing limits is not reporting results that are meaningful to the user experience.