EdgeStream

Network Latency and Its Effect on Video Streaming

Randy Chung
Chief Technology Officer

Introduction: Streaming Speed and Network Quality

Streaming video is commonly transmitted using UDP, TCP, or HTTP protocols. Using the TCP or HTTP streaming protocols will provide error free streaming, because the underlying TCP protocol retransmits missing or erroneous data until the data is received correctly.

EdgeStream technology uses HTTP streaming for error free delivery and because it is the most "firewall friendly" of the three protocols. The HTTP protocol is the same one used to deliver web pages. HTTP actually is a layer above the TCP protocol, so the data transmission speed for both HTTP and TCP is limited by how the TCP transmission speed.

TCP transmission uses a "slow start" algorithm when data transfer first starts. The transfer rate increases until data congestion is detected which limits the transfer rate. Then TCP transmission switches to a congestion avoidance algorithm for the bulk of the transfer.

The congestion avoidance algorithm will cut the transmission rate in half when congestion is detected, and will increase the transmission incrementally up to a maximum determined by the network latency (the time it takes for data to travel from source to destination) and TCP window size (which determines the number of data packets that can be transmitted as a group) if no congestion occurs.

The maximum TCP transmission rate is shown in the equation below, where S_{Window} is the maximum number of packets in the window, S_{Packet} is the packet size in bytes, and RTT is the network round trip time from server to client. On the Internet, the TCP transmission rate is affected by packet loss, which is not part of this simple equation.

$$RATE_{TCP} = \frac{S_{Window} * S_{packet}}{RTT}$$

Empirical Results

To determine the actual TCP window size for a typical transfer, a Windows 2000 server and Windows 2000 client were set up, both connected to a 100 base T local area network. The latency and packet loss between the server and client was set using the Cloud network emulation product from Shunra Software. A TCP video transfer was started and the maximum transmission rate was measured.

From the measurements with zero packet loss, S_{Window} was determined to be 22 packets for a packet size of about 1500 bytes, or about 32 Kbytes. For a network round trip time of 80 milliseconds, the maximum rate was measured to be about 3000 Kbit/second.

Currently within the United States, network round trip times of under 100 milliseconds with near zero packet loss can be achieved during many times of the day, so there is a short video clips can often be delivered over the Internet at rates at or above 1 Mbps — high enough to be "TV quality". Of course, network round trip times and packet loss can increase substantially at unpredictable times, affecting delivery speeds and service quality especially for content that might last 30 minutes or more.

We conducted additional throughput measurements with higher than normal latencies and packet losses. We found that the maximum TCP transmission rate drops dramatically when network impairments from a busy Internet cause increases in latency and packet loss. Higher than normal network latencies and packet loss can occur unpredictably and can cause video interruptions if they last for more than 30 seconds.

The figure below shows TCP streaming speeds measured at a network round trip time of 80 milliseconds and increasing packet loss. The measurements show that packet loss approaching 1% will drop the TCP streaming speed below the necessary rate for 1 Mbps TV quality video.

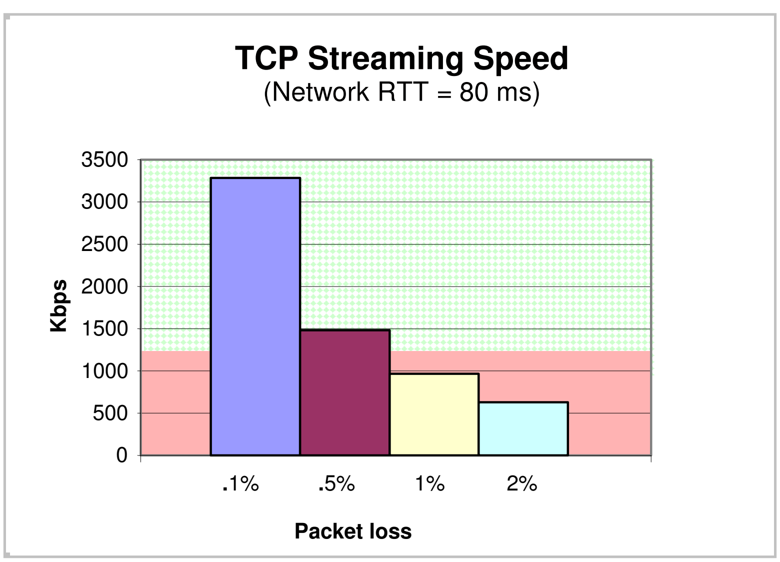

Network impairments may affect even major network providers at any time. Network operators work to keep the network impairments to a minimum, but cannot maintain perfection every second of every day.

The bursts of high latency and packet loss can affect groups of users at random times and at random locations.

The example below shows a 24 hour record of network latency and packet loss for a particular location in Texas, downstream from a major network backbone. It shows the average latency is about 60 milliseconds, but the peaks can reach 150 milliseconds. More importantly, while the average packet loss is close to zero, it can occasionally reach the 1% to 2% level for several minutes at a time. These periods of increased latency and packet loss can cause interruptions in high speed video streaming.

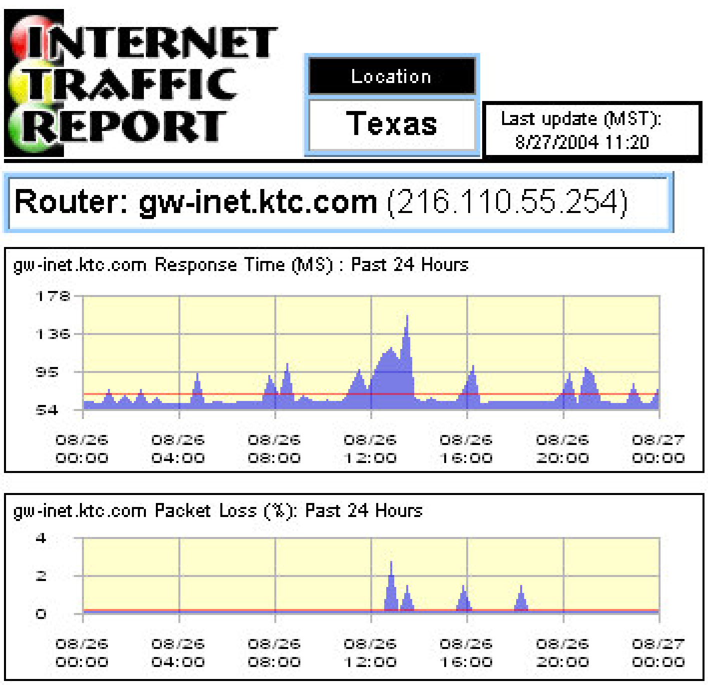

(This report information was obtained from www.internettrafficreport.com).

Network impairments between major backbones can occur at random times and affect only certain locations. The report below shows higher than normal latencies on certain links between UUNet and AT&T, Verio and AT&T, and Qwest and Level3. (Packet loss information is not available). It is unpredictable as to which networks may be affected, when these impairments occur, or how long these network impairments may last.

KEYNOTE
The Internet Performance Authority

The Internet Health Report (Last Hour)

Generated Fri Aug 27 20:47:33 2004 GMT

	UUNET	SAVVIS	QWEST	VERIO	SPRINT	LEVEL3	AT&T	INTERNAP	WILTEL	COGENT
UUNET	45	28	39	35	33	37	32	42	33	39
SAVVIS	28	17	36	29	22	26	23	25	25	43
QWEST	39	36	48	35	45	53	36	52	44	45
VERIO	35	29	35	76	31	33	34	38	35	39
SPRINT	33	22	45	31	38	35	31	37	33	39
LEVEL3	37	26	53	33	35	46	35	39	40	41
AT&T	32	23	36	34	31	35	37	36	29	32
INTERNAP	42	25	52	38	37	39	36	65	44	39
WILTEL	33	25	44	35	33	40	29	44	39	36
COGENT	39	43	45	39	39	41	32	39	36	41

UUNet to AT&T (Last Hour)

Generated Fri Aug 27 20:47:09 2004 GMT

UUNET \ AT&T	NY	DC	SANFRAN	DALLAS
SANFRAN	74	126	11	55
DALLAS	47	38	50	6
NY	9	13	68	47
ATLANTA	27	19	61	25

Verio to AT&T (Last Hour)

Generated Fri Aug 27 20:47:22 2004 GMT

VERIO \ AT&T	NY	DC	SANFRAN	DALLAS
LA	77	131	13	37
PHILLY	9	12	69	44

Qwest to Level3 (Last Hour)

Generated Fri Aug 27 20:47:17 2004 GMT

QWEST \ LEVEL3	PHILLY	DETROIT	SANDIEGO	STLOUIS
DC	11	37	79	36
TAMPA	50	71	66	65
LA	85	94	24	87
DENVER	56	77	47	75

(This report information was obtained from www.internetpulse.net)

The network latency and packet loss measurements were repeated with higher network latency and packet loss, and compared the streaming performance between standard TCP transfers and EdgeStream transfers. We waited for the video transfer to stabilize, then took a measurement over a 60 second average. In both cases, the video was a 1 Mbps average clip. The EdgeStream implementation allows the video transfer rate to peak up to three times the normal clip rate, if the connection bandwidth supports it.

While the standard TCP streaming bandwidth dropped sharply with increasing packet loss, the EdgeStream bandwidth was practically unchanged, even when the packet loss reached 2%.

When the network latency is increased to 160 milliseconds round trip time, streaming speed using EdgeStream technology remains very high, while speed under standard streaming cannot sustain 1 Mbps even with 0.5% packet loss.

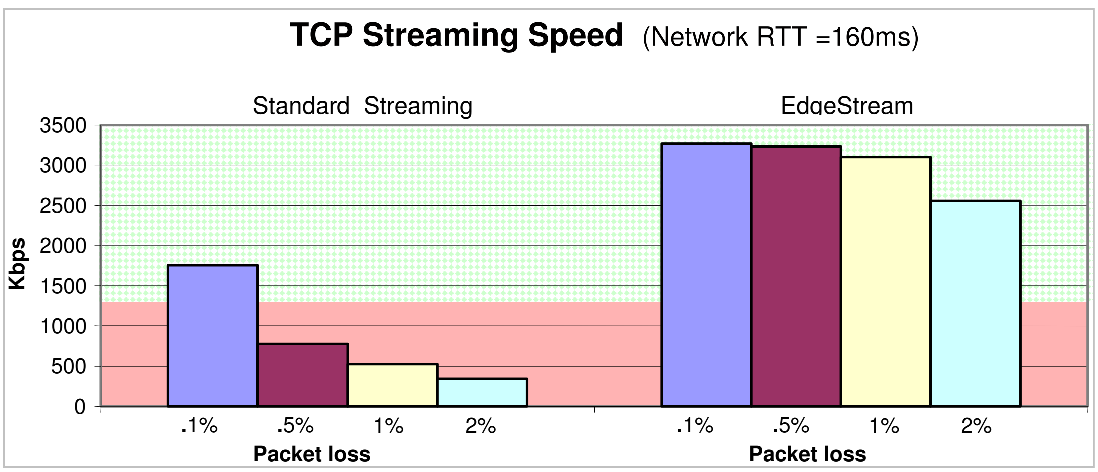

Packet loss with 80 msec latency	*0.10%*	*0.50%*	*1%*	*2%*
Standard throughput, Kbps	1758	775	526	343
EdgeStream throughput, Kbps	3266	3231	3102	2557

Conclusion

EdgeStream consistently delivers high bit rate streams even when the network path has substantial latency and packet loss. This becomes a critical ability when TV and DVD quality content must be delivered to a demanding audience for extended viewing. Furthermore, as the use of these high bit rate services increases, network traffic will also increase. The result will be greater latency and packet loss, further driving demand for the EdgeStream solution.

The capability to deliver streaming video through network impairments like high latency and packet loss, combined with the ability to automatically and transparently route around network congestion, makes EdgeStream technology indispensable in applications where the objective is to consistently provide the highest quality video to the largest possible Internet audience.