NetForecast®

Internet Connection Requirements for Effective Video Conferencing to Support Work from Home and eLearning

Alan Jones, Peter Sevcik, and Rebecca Wetzel

February 15, 2021

EXECUTIVE SUMMARY

Pandemic precautions and stay-at-home mandates have dramatically heightened dependence on video conferencing applications like Zoom, Microsoft Teams, WebEx, Google Meet, and Slack for remote work and eLearning. This dependence has raised questions about how internet connectivity affects a user's video conferencing experience. To answer these questions, NetForecast undertook a study to independently determine the internet connectivity and performance requirements to support popular video conferencing applications. We measured network bandwidth consumed during actual business and education video conference sessions and analyzed the data to determine the bandwidth required for an acceptable user experience. We also assessed how latency and packet loss affected the user experience.

Based on video conferencing data collected over more than 152 hours by a range of users in widely distributed locations, NetForecast determined that under real-world conditions the internet connectivity required for video conferencing is lower than other activities such as online gaming and streaming video. In general, acceptable video conferencing quality requires 5Mbps of down, and 3 Mbps of up bandwidth for a single user.

Video and voice are different. They affect the human experience differently and generate different network demands. Achieving good video and audio quality over an internet connection requires attention to multiple network performance parameters. In addition to sufficient network bandwidth, other network performance parameters—especially network latency and packet loss—affect video conferencing experience quality. Latency can introduce long delays which each session participant sees and hears. Latency can degrade video quality, and make interaction difficult when video and audio become out of sync. When packet loss reaches a certain threshold, it can cause video to blur and freeze.

NetForecast found notable differences in sensitivity to latency and loss among the applications tested, with some applications continuing to function well with Increased latency and loss, while others degraded. Our study found that low latency is actually more critical than bandwidth or loss for successful business and eLearning sessions.

NetForecast®

NetForecast Report NFR5137

USAGE DATA COLLECTION & ANALYSIS

Volunteers collected video conference usage data using a NetForecast-developed application. Before each video session, users started the application and selected the video conferencing service they would be using. During each video conference, users entered comments within the NetForecast application to tag the data with observations about their session experience, such as poor audio and freezing video. They also noted conditions that changed during the session, such as when a presenter began screen sharing.

NetForecast's application recorded network utilization from Microsoft Windows Performance Monitor. Users were instructed to minimize use of other applications during video conferencing. The application recorded the maximum count of megabits instantaneously sent and received, and the mean megabits sent and received during consecutive 10-second intervals. The application also recorded the number of packets that experienced error and loss conditions.

Only data from sessions with three or more participants with their video cameras active were used for analysis. Volunteers' provisioned internet connection down speeds ranged from 1Gbps to 3Mbps (DSL). NetForecast collected data from 12 users covering 152 hours from November 19, 2020 through January 7, 2021.

Data were aggregated and analyzed to determine minimum, maximum, and mean values, as well as standard deviation. NetForecast data analysts plotted the data to identify trends and unusual behaviors, and to validate conclusions drawn from statistical analysis. Graphs of the data were overlayed with user comments to correlate comments with network conditions.

VIDEO CONFERENCING NETWORK RESOURCE USAGE RESULTS

NetForecast found that the bandwidth consumed by the video conferencing applications we tested was surprisingly low. As video conferencing software has advanced, vendors have applied sophisticated compression algorithms and other mechanisms[1] to reduce downstream and upstream bandwidth consumption. The NetForecast data summarized in Figure 1 illustrate the results of these efforts.

Figure 1 shows the average bandwidth consumed by each application over the study period, and it shows the 95th percentile, which is the value below which 95% of the usage measurements fell for each application. The 95th percentile value reflects the maximum bandwidth each application can generally be expected to use. NetForecast recommends internet service packages which provide bandwidth exceeding the 95th percentile value shown in Figure 1 to consistently support videoconferencing for a single user. Because adaptive bitrate (ABR) technologies used by some of the applications reduce the per-user bandwidth required when available bandwidth is constrained, we caution against linearly extrapolating demand for multiple users.

Figure 1 also shows that the minimum vendor-recommended bandwidth is generally less than half the NetForecast-recommended minimum (excluding Google Meet). Unimpeded by bandwidth, the generally highest bandwidth (95th percentile) consumed by the applications is shown in Table 1. Measurements on low-bandwidth connections confirm that all of the applications tested can run at less than the vendors' recommended values, but with degraded performance.

[1] Many video conferencing applications use adaptive bitrate (ABR) techniques. See: Video optimization - Wikipedia

NetForecast believes the 95th percentile value represents the minimum bandwidth that allows the application to operate with the default video and audio quality. Google Meet recommends both down and up bandwidth above 3 Mbps, however, in our study Google Meet consistently used significantly less bandwidth.

To support any of the applications measured, NetForecast recommends internet access bandwidth of 5Mbps down and 3Mbps up for a single user.

Application	Downstream			Upstream		
	Measured Average (Mbps)	Measured 95th Percentile (Mbps)	Vendor Recommended (Mbps)	Measured Average (Mbps)	Measured 95th Percentile (Mbps)	Vendor Recommended (Mbps)
Zoom	1.0	2.1	1.0	0.4	1.0	0.8
Slack	0.9	2.5	1.2	0.5	1.4	0.6
Google Meet	1.2	2.0	3.2	0.4	0.8	3.2
WebEx	1.3	2.8	0.4	0.8	1.4	0.4
Teams	1.5	3.5	2.0	1.0	2.3	1.0

Figure 1 – Single Session Video Conferencing Bandwidth Consumption[2]

[2] Sources for vendor-recommended values in Table 1 are listed in the Reference section at the end of this report.

BANDWIDTH TEST RESULTS BY APPLICATION

The following charts all show two representative sessions of each application of the many measured. The number of participants in each meeting may exceed the number of users running the Observer program. Charts include all users including those with DSL service.

Zoom

Zoom, the most widely used video conferencing application in the US, performed well during our study. Zoom's up and down bandwidth requirements were at the low end of the pack. For meetings of three or more people, users' comments indicate that Zoom was preferred in terms of reliability over Slack.

> The blue and orange colors on this and subsequent graphs indicate individual sessions

Figure 3 compares Zoom sessions seven days apart with four users on video for each call. A volunteer with a 3Mbps DSL connection was on both calls. This user did not report any issues during the sessions.

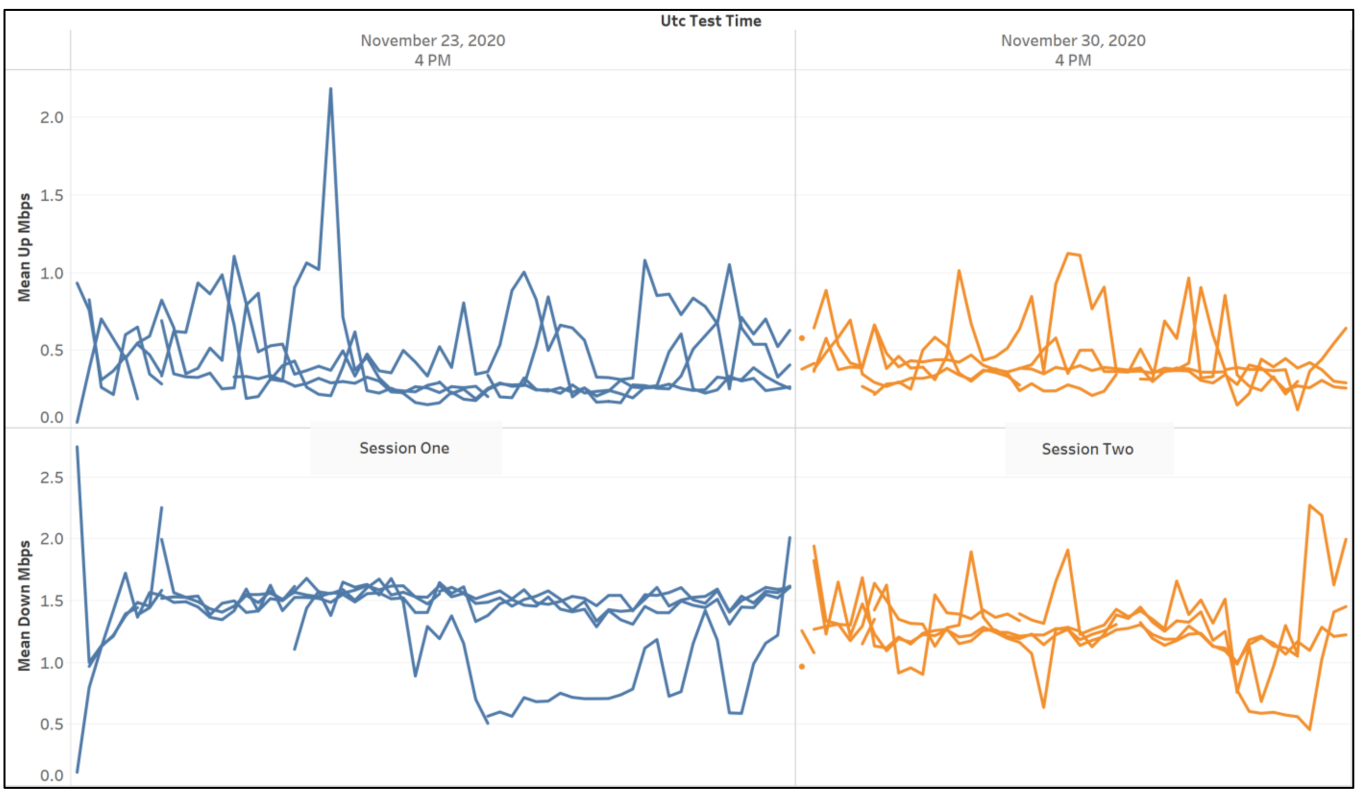

Figure 2 – Mean Down and Up Consumption for Zoom

Slack

Slack is significantly more than a video conferencing tool. Slack supports individual messaging as well as message groups for business collaboration. One advantage to Slack is the ease of starting calls and adding users to calls. Volunteers with lower bandwidth connections consistently report issues with long connection times (30 to 90 seconds) and reported more disconnects.

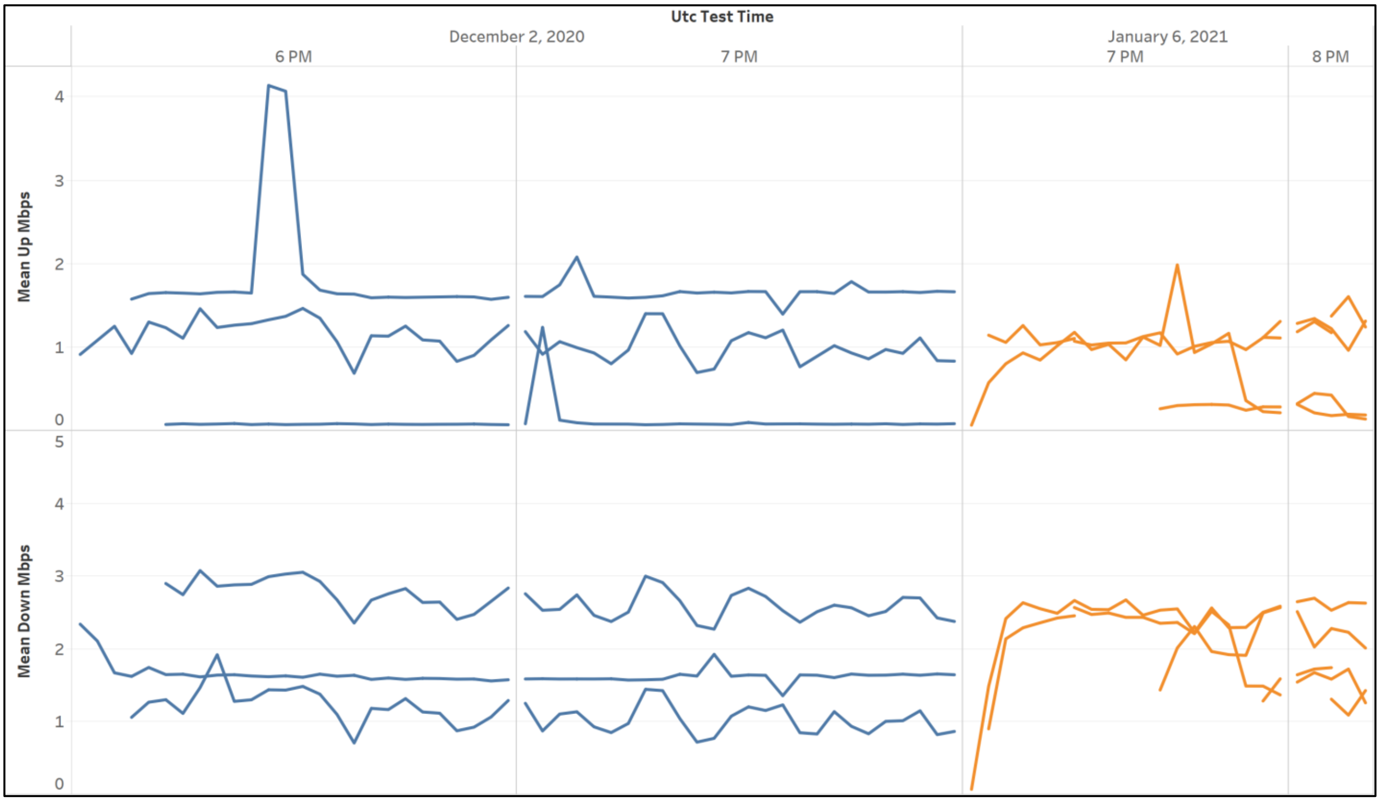

Figure 3 – Mean Down and Up Consumption for Slack

Google Meet

Google Meet is part of the popular G-Suite business package. The major differentiating factor of Google Meet is consistently low send bandwidth and less variation in send and receive bandwidth usage.

Note that all of the applications, including Google Meet, operated consistently below the bandwidth limitations of the DSL volunteer during video conferencing sessions.

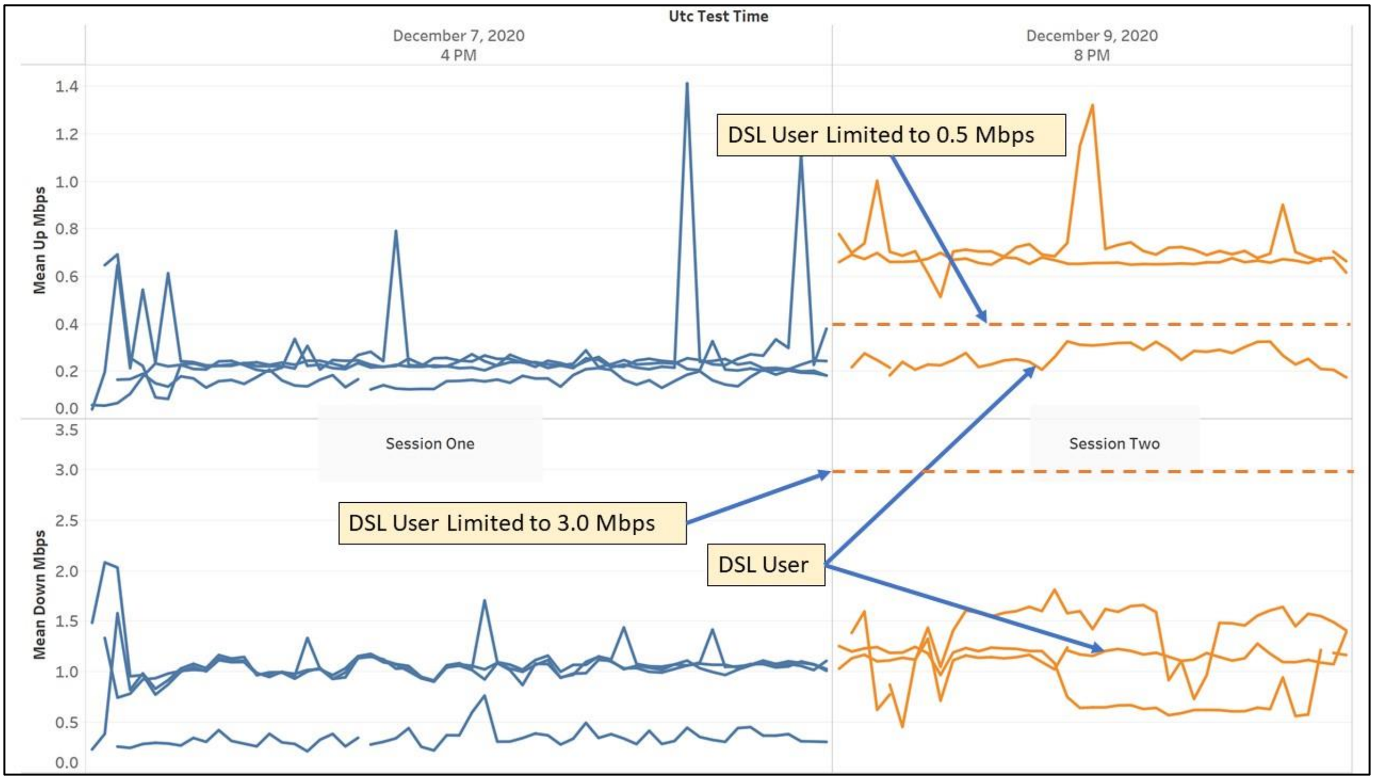

Figure 4 – Mean Down and Up Consumption for Google Meet

WebEx

WebEx is a business video conferencing tool from Cisco. Introduced in 1995, it was among the first widely available video conferencing products. The graph in Figure 6 shows two sessions with four participants with active video. We also measured bandwidth usage for three large-scale WebEx meetings with approximately 50 participants on audio and a single presenter sharing a screen. The bandwidth usage for these sessions remained at or below the bandwidth consumed by sessions with three to eight participants all on video.

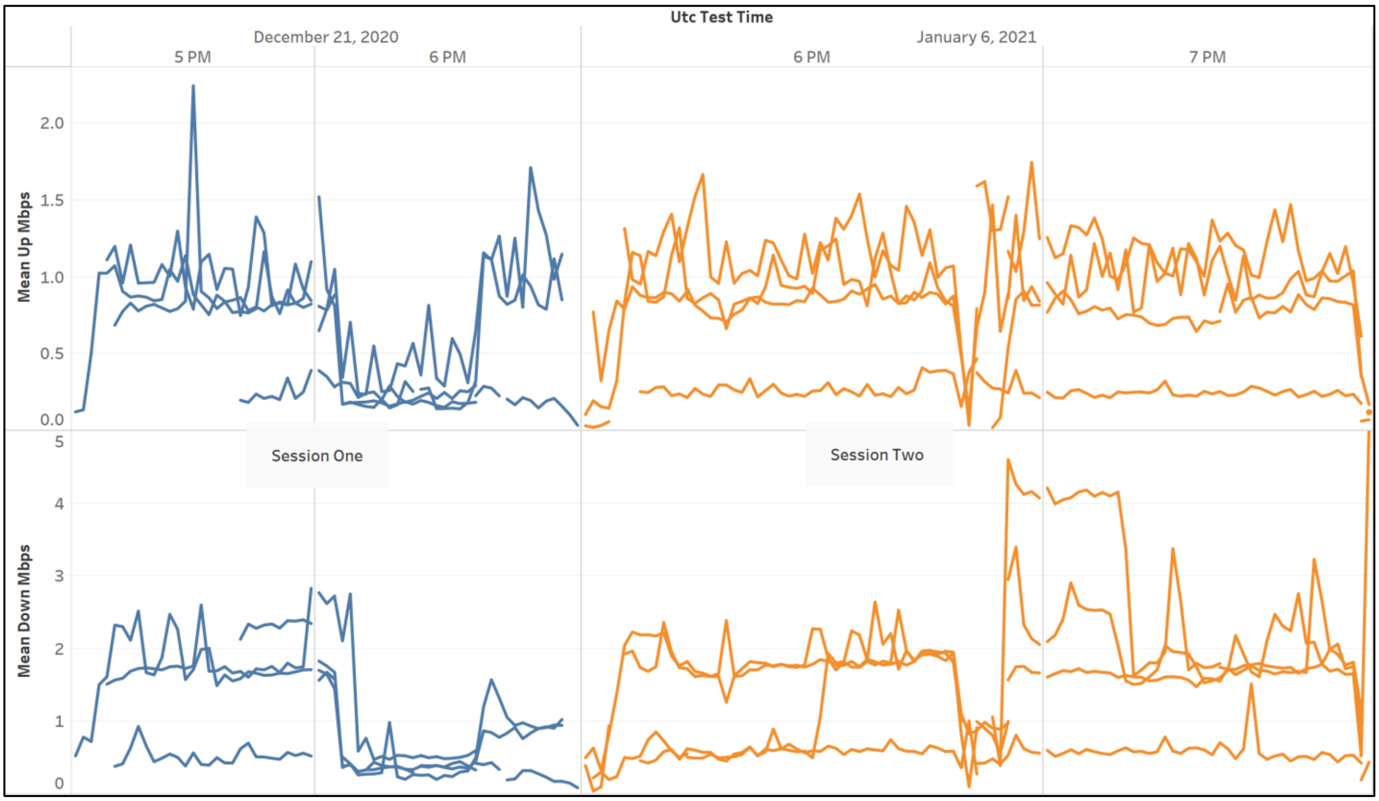

Figure 5 – Mean Down and Up Consumption for WebEx

Microsoft Teams

As part of the Microsoft Office 365 suite, Teams is a ubiquitous video conferencing application for business. Study participants reported few issues with Teams, and several users reported the best initial connection speed and reliability.

Figure 6 documents results from sessions with the same four users on two consecutive days.

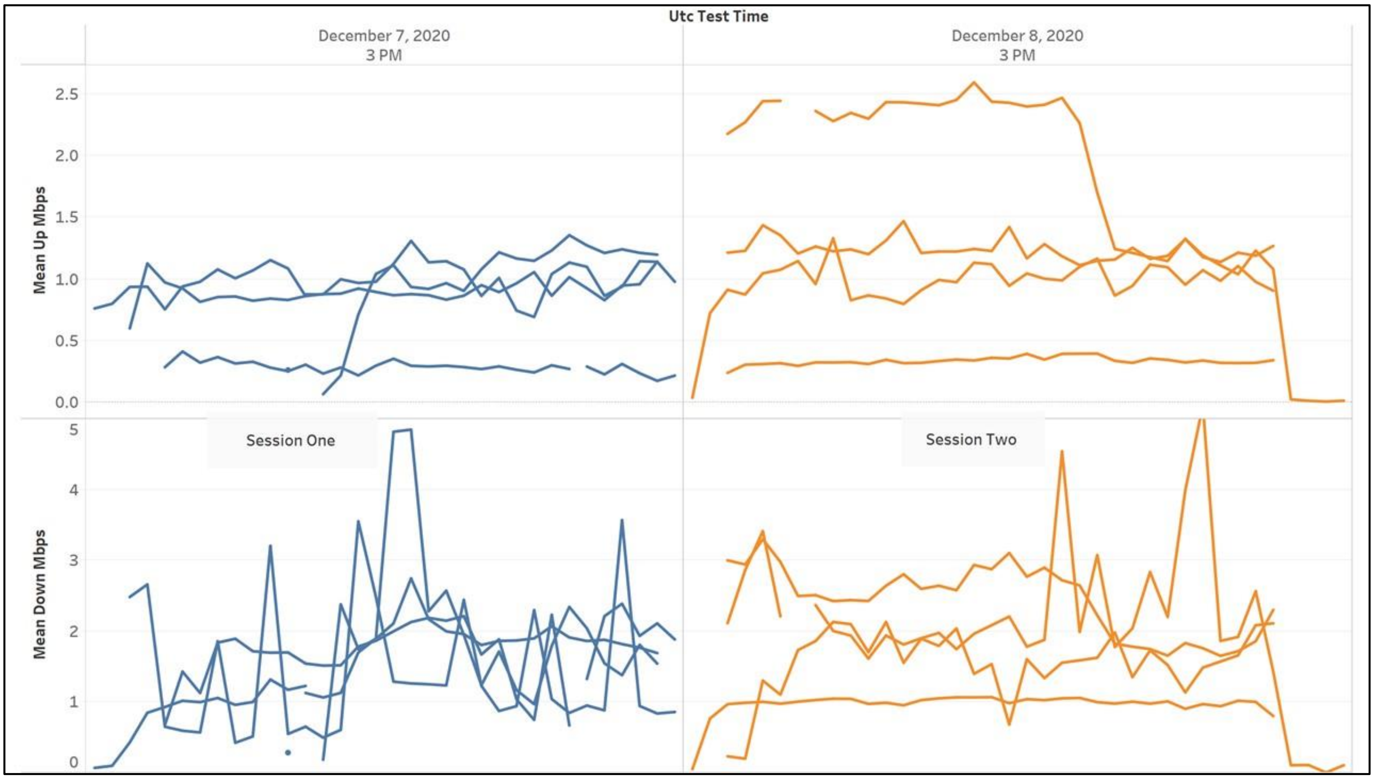

Figure 6 – Mean Down and Up Consumption for Teams

In general latency impacts audio quality, while loss impacts video quality.

Bandwidth is not the only factor contributing to a satisfactory video conferencing experience. Increased latency and/or loss also influence the experience of a video conferencing participant. For this reason, NetForecast also measured video conferencing application performance under controlled conditions to assess the outcome when latency and loss were intentionally introduced into the network connection. Note that all latency values discussed in this section are Round Trip Time (RTT) values.

Three-person video conferencing sessions were conducted using each of the five video conferencing applications. The principal investigator connected through a Wide Area Network (WAN) emulator that enabled them to artificially increase latency and inject packet loss. During video conferencing sessions the investigator introduced loss and latency in increasing combinations until the sessions degraded. Performance was considered to be degraded when there was obvious loss of video resolution, i.e., video became pixelated, choppy, or frozen, and/or audio was corrupted or lost.

Effect of Latency on Audio Quality

Human sensitivity to audio response time in conversations has been extensively researched. The ITU G.114 standard recommendation regarding mouth-to-ear delay states that the user quality of experience (QoE) boundary between satisfied and dissatisfied is at 290 msec one-way latency. This means the RTT must stay below 580msec. Although this latency target seems easy to achieve, we must account for all latency contributions to the audio session.

Figure 7 shows the typical path between two users of a video conference service. The service operates an application on the users' device (PC, tablet, phone) and a conference server on the internet which, among other things, operates audio mixing software. The mixing software operates a store-and-forward function for each packet from speaker to all listeners.

All conferencing services add delay to the audio at the red processing points in Figure 7. This processing-induced delay is the end-to-end latency floor to which the latency of all the networks along the data path is added. The sum of all delay contributions between the two users farthest from the conferencing server determines QoE for all session users.

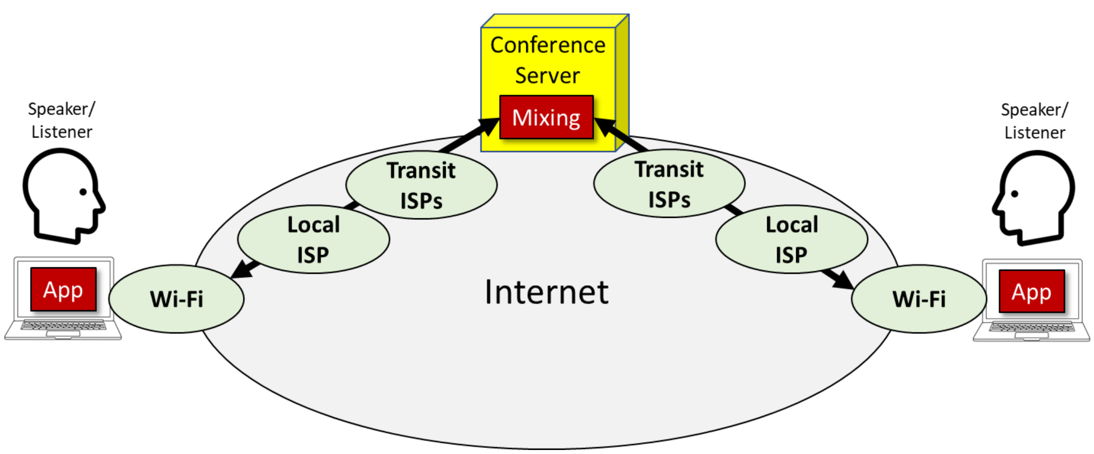

Figure 7 – Elements Along an End-to-End Conference Session

NetForecast measured the inherent internal latency of each video conferencing application and found that the latency they add is fixed but different for each application. The internal latency is a product of service design tradeoffs balancing and synchronizing voice and video quality. The values in the audio processing delay column in Figure 8 are the starting point towards the 580 msec limit defined by the ITU. As the figure shows, some applications are at a clear disadvantage compared to others. The network RTT allowance column shows how much network delay (the cumulative delay from all the green ovals in figure 7) can be added without exceeding the 580 msec limit.

Application	Audio Processing Delay (msec)	Network RTT Allowance (msec)
Zoom	360	220
Slack	300	280
Google Meet	450	130
WebEx	240	340
Teams	340	240

Figure 8 – Total Network End-to-End RTT Required for User Satisfied QoE

Remaining within the ITU latency limit for transcontinental video conferences can be challenging because packet routing and server location can add delay. A conferencing service must select a single server from a pool of available servers. This selection should be based on the optimal distance among all conference participants, but factors such a server utilization may route the call to a suboptimal location. Also, the internet path to the server selected by a local ISP and/or intermediate transit ISPs may be sub-optimal. The same challenges apply in the opposite direction.

Nearly all user/server paths we measured included approximately 20 internet hops. Paths for a user in central Virginia went south to Atlanta, then to northern Virginia, and terminated at a server in New York City. Poor routing such as this adds significant delay and degrades session quality.

Low latency is actually more critical than bandwidth or loss for a successful business or eLearning conference call.

Long distances—such as a video conference call involving participants on both US coasts—increase the risk of poor audio quality. If multiple participants are spread over long distances, a call can easily degrade into meaningless noise when multiple speakers try to speak simultaneously. This is one reason video conferencing applications enable a conference moderator to mute all talkers.

An effective conference call involving dialog among multiple participants requires a low latency starting point. For this reason, NetForecast concludes that low latency is actually more critical than bandwidth or loss for a successful business or eLearning video conference session.

As latency increases, participants experience increasingly poor synchronization between a speaker's audio and video image. As the principal investigator intentionally increased latency, lip synchronization deteriorated because the video arrived before the audio. This indicates that conferencing applications take more time processing audio than video, despite the fact that audio requires moving fewer bits than video.

Effect of Loss on Video Quality

Loss during the calls was increased based on the percentage of packets dropped by the emulator. Once loss reached a level above which the application was unable to compensate, issues like video freezing or pixilation and dropped or garbled audio emerged. Loss degraded user video experience quality more than latency. Figure 9 shows the percent loss at which performance degraded to the point that it interfered with the user experience.

As you can see, Zoom, WebEx and Teams are less sensitive to packet loss, experiencing degradation at 10% packet loss, compared to Slack, which experienced degradation at 5% loss, and Google Meet, which experienced degradation at 1% loss.

Application	Packet Loss at Which User Experience Degrades (%)
Zoom	10
Slack	5
Google Meet	1
WebEx	10
Teams	10

Figure 9 – Packet Loss Effect on a User's Video Experience

Although fiber or coax internet services would typically not see these types of latency and loss events, wireless connections, either through cellular networks or via Wi-Fi, are susceptible to coverage, signal strength, and signal interference issues that can result in loss and latency values that exceed these thresholds.

CONCLUSIONS

The effort expended by video conferencing application providers to improve reliability and reduce bandwidth has paid dividends. All users, even those with the slowest connection speeds, reported reasonable call quality on all the platforms studied. A single user of any internet connection at 5Mbps downstream speed and 3Mbps upstream speed should expect to conduct video conference calls with no bandwidth-related issues. Experiments conducted in a controlled environment show latency and loss can significantly affect the user's video conference experience. Wired internet connections should provide adequate bandwidth: however, latency and loss can occur on connections of all speeds. Other factors such as the quality of a cellular connection or Wi-Fi signal could adversely impact video conferencing. Users should be aware that signal strength and integrity can cause video conferencing issues.

Based on NetForecast's bandwidth usage measurements, observations by volunteers, and verification in a controlled environment, even low bandwidth internet connections can provide an adequate video conferencing experience as long as the connection is free from latency and loss issues. Even marginal latency increases, when combined with audio processing times of the video conferencing applications, can degrade a user's audio experience. Although we identified circumstances that degrade user QoE, based on our data and user observations, we conclude that the applications tested performed well most of the time.

ABOUT THE AUTHORS

Alan Jones is NetForecast's Director of Software Development. He has lead teams in developing products and internal infrastructure for some of the largest telecom companies in the world. After eight years in cellular handset design and testing, he spent over a decade working on test systems for mobile networks. He currently works with mobile and cloud-based product development.

Peter Sevcik is the founder and CTO of NetForecast and is a leading network performance expert. An internet pioneer, Peter was among the first to measure and develop internet performance improvement techniques. He helped design more than 100 corporate and commercial networks. In addition, Peter invented the Apdex performance reporting methodology, and has co-patented application response-time prediction and network congestion management algorithms.

Rebecca Wetzel is a principal at NetForecast, and an internet industry veteran. She helped realize the commercialization of the internet in its early days, and worked to design and market some of the internet's first value-added services such as IP-based VPNs, web hosting, and managed firewall services, as well as internet protocol testing services. She also spent many years as an internet industry analyst and consultant.

REFERENCES

Vendor-Recommended Bandwidth Requirements

Zoom: https://support.zoom.us/hc/en-us/articles/201362023-System-requirements-for-Windows-macOS-and-Linux

Slack: https://slack.com/help/articles/115003538426-Troubleshoot-Slack-calls

Google Meet: https://support.google.com/a/answer/1279090?hl=en#zippy=%2Cstep-review-bandwidth-requirements

WebEx: https://www.cisco.com/c/en/us/products/collateral/conferencing/webex-meetings/white_paper_c11-691351.html

Teams: https://docs.microsoft.com/en-us/microsoftteams/prepare-network