345 Encinal Street
Santa Cruz, CA 95060
+1 (831) 426-5858

Preparing Your IP Network for Video Conferencing

AND WORKAROUNDS FOR LESS THAN IDEAL CONDITIONS

WHITEPAPER

Poly Infrastructure Engineering

August 30, 2019

 3725-86432-001A

Abstract

Carrying real-time audio and video traffic over a network brings a unique set of concerns that require specialized considerations when compared to “normal” data traffic (i.e. web pages, e-mail, file downloading, etc.). Attention must be paid to designing and configuring IP networks for real-time audio/video (A/V) traffic in order to deliver a good quality experience to all users. Following a set of “best practices” for network design and configuration, as well as in deployment/configuration of individual products, helps ensure a high-quality encounter for all users. If a network is less than ideal and/or “best-practices” cannot be strictly followed, you can employ a variety of remedies to mitigate, if not eliminate, perceived quality issues. Please contact a Poly support representative for further details and assistance.

Call Quality Issues

Audio and video participants may experience call quality problems such as choppy audio/video, missed words in speaking, blurry video, jumbled audio, dropped calls, video artifacts (rainbows, pixilation, lines, shadows, …), “talk-over,” and many others. These problems usually stem from one or more of the following conditions, which particularly affect the real-time nature of A/V communications:

- **Latency/Delay** – Network latency/delay is the time it takes a packet to travel between two points in the network. For A/V communications, this equates to the time between when a person makes a sound and when the person on the other end of a call hears that sound. Or it is the time between when a gesture is made on camera and when the other participants see it happen on their screens. There is always some delay in a network as packets are routed and processed by one device to another on their network journey. When the delays begin to be “too” long, users may begin to experience talk-over, where one person starts talking over another, because they think it is their “turn.” As delays increase, other quality issues are perceived such as lip-sync issues (audio and video are not synchronized and people look as if they are being “dubbed”), jumbled audio/video, or even packet loss (as network services may begin discarding outdated packets). For reference, **delays of just 200ms are human-perceptible** and over-talk can start happening.

- **Packet Loss** – Packets on a network may be discarded or dropped for a variety of reasons: dropping of outdated packets, data corruption from interference, poor signal strength (wireless), various network configurations, etc. As packet loss increases, users experience audio/video “cut out,” jumbling, and various on-screen artifacts. The goal of every network design and operation is to have 0% average packet loss. On a call without any corrective features (such as Poly’s LPR), **packet losses going over 1% may start to become noticeable** by participants, and at 5% audio quality degrades significantly and video (people or content) is generally unusable.
  Packet loss is determined in UC devices by a jump in the UDP sequence number. UC devices usually provide debugging tools to capture network traffic data for analysis. This jump can also be observed by standard network tools.

- **Jitter** – Is the inconsistency in delays/latency of packets going between two network points. For A/V calls, packets not travelling smoothly/consistently between two endpoints can result in packets arriving out-of-order and/or getting dropped before processing, which leads to packet loss and the associated call quality issues. **Average jitter is typically desired to be 40ms or less for good call quality**. Jitter is typically the result of slow network speeds, congested networks, and/or improper network routing configuration.

- **Bandwidth Overuse (network congestion)** – Networks, or individual network segments, being overutilized, or having more packets than they can handle at any point in time, can lead to many follow-on problems. Like a road over congested with too many cars on it, packets may fill up a network pipe and become bottlenecked. The result is that these packets may be delayed (increasing latency and jitter), rerouted (increasing latency and jitter), or outright discarded (increasing packet loss) depending on network configurations and conditions. It is generally recommended to **keep peak utilization (the most network bandwidth used at any point in time) at 75%-80% or less of the total available bandwidth**. For example, for a 1 Gbps network link, no more than 750 Mbps of traffic should ever be placed over it at any one time.

Fixing the Network and Best-Practices

To address the conditions that lead to poor call quality, addressing the network itself is usually the most effective place to start.

1. Most importantly, monitor the network and the attached infrastructure devices (switches, routers, cables, wireless access points, firewalls, etc.…). Most network devices have built in monitoring facilities to record the packet loss, jitter, latency, and utilization values, among others. Monitoring devices/services/tools can also be attached at various network points to provide greater details. Without monitoring, network issues will only be known once outages or poor call qualities are reported. Even then, diagnosing the specific causes may be difficult in any medium to large network.

2. Make sure all network hardware is in proper functioning order. Many network issues can be traced back to a faulty cable, loose connection, malfunctioning switch, or the like. The distributed and self-healing nature of networks may not make failing hardware obvious, as degradation rather than full outage, is often the only symptom.

3. Invest in bandwidth and hardware where needed. Since bandwidth overuse can lead to many networking problems, it is important to keep peak utilization at 75%-80% or less between any two points at any time. Upgrading link speed (i.e. 100Mbps -> 1Gpbs -> 10 Gbps), adding new links, installing faster higher-capacity/faster hardware, etc. can help to alleviate both bandwidth and latency concerns. This is generally the #1 thing that can be done to help resolve call quality problems.

4. Ensure your call QoS (Quality of Service) for A/V calling by changing your network configuration (layer 2 and/or layer 3) to prefer real time traffic (RTP) over other types of traffic and real-time devices (endpoints, MCUs, etc.) over other types of devices (i.e. web servers). Configure this in the switches, routers, and firewalls that form the network. There are a variety of schemes that can be employed to accomplish this: DiffServ DSCP packet marking, TOS packet marking, and separate VLANs with high priority routing for A/V devices; to name a few. These types of configurations help ensure that the latency, jitter, and packet-loss are minimized automatically by your network if/when bandwidth becomes overutilized. It also helps insulate A/V traffic from other non-real-time traffic (i.e. call quality won't suffer just because someone else is transfering a large file). Note: In many individual products, particularly Poly devices (phones, video endpoints, and

infrastructure), many of these QoS features are also supported. Together with the rest of the network infrastructure configuration, these features can provide for true end-to-end QoS and higher quality calls. Consult individual product manuals for QoS setting capabilities.

5. Fix routing and/or other related configuration/issues. Packets not taking the most efficient and/or direct path between two points can lead to increased latency and jitter (and possible packet loss as a result). Often an improperly configured route can be sending packets in an inefficient direction, or even unnecessarily over a congested network path. It is important to properly configure and periodically test network routes, as incorrect routing can creep in, but this may not be immediately obvious except in the possible loss of call quality.

6. Deactivate firewall ALG and RTP gateway features. Even though many popular firewalls offer A/V ALG (Application Layer Gateway) features, they often lack the nuances, expertise, and general feature sets needed to properly implement A/V calling services correctly. As a result, they may end up erroneously dropping packets, increasing latency, and/or generally interfering with A/V calls in undesirable ways. Where security is needed, a purpose-built A/V security SBC (Session Border Controller) device, such as the Polycom DMA-Edge system, should be employed. With the proper ports forwarded to the them, these SBC devices take over the security duties for A/V calls from the firewall and still ensure the full range of A/V calling features, while at the same time minimizing latency, jitter, and packet loss.

7. Minimize the use of VPNs for A/V calls. VPNs introduce latency/delays to the A/V packets and often do not support the QoS settings/features needed to prioritize real-time-traffic over other types of traffic through them. It is desirable to avoid VPNs for all A/V packets wherever possible. With proper firewall configuration, and use of security SBCs, VPNs are often unnecessary to obtain equal or greater security levels.

8. Ensure that the network core devices are configured to allow many and rapid UDP-RTP packets from single IP address sources. There are security configurations in many switches/routers/firewalls that can interfere with both the signaling and UDP-RTP packets needed for A/V calling. While in a call, an A/V device may send several UDP-RTP packets every few milliseconds. Network devices should be configured to not see this volume of traffic as an “attack” and incorrectly filter (i.e. drop, AKA: packet-loss) these packets.

9. Deploy MCUs and other critical A/V infrastructure near the center of the network where bandwidth is usually higher and the average distances, and thus latency, between endpoints and these core components are minimized.

10. Remember that the best-practices used to ensure quality of a network within a site should also be employed between network sites (i.e. between offices via MPLS links) and up to the internet if A/V calls will traverse these. Cross site and internet uplink providers offer a variety of service levels and options. These Service Level Agreements (SLAs) contractually ensure network characteristics such as bandwidth, maximum latency, and maximum jitter across their links. It is important to size these links correctly for the expected traffic bandwidth.

11. Time synchronization is very important in modern networks and many voice and video applications are uniquely susceptible to incorrect time synchronization. Reliable NTP (Network Time Protocol) servers with

very high stratum (1,2, or 3) should be used for all network connected devices, and multiple servers should be configured wherever possible (3 or more recommended).

Fixes at the Device and in the Product

Aside from general network changes to help support A/V calls, individual A/V devices may have features and configurations to help improve overall call quality.

1. QoS settings – Many devices, especially Poly devices, support features such as DiffServ and IP Precedence packet marking, and VLANs. These QoS settings, in conjunction with like configurations in the network itself, can greatly improve call quality.

2. Call Quality Selection – When placing an A/V call, the desired maximum call quality (in bit-rate) can often be selected (either configured and/or selected at call time). It is important to properly configure endpoint devices not to place calls with bit-rates higher than desired or for which there is insufficient network capacity. On the network a call will consume its bit-rate + roughly a 20% overhead of bandwidth. For example, a 100 Mbps link can only handle 10 6Mbps A/V calls, with no other traffic (e.g. (100Mbps * .75) / (6Mbps + (.2 * 6Mbps))). If more calling capacity is desired, or if other network services use these links, these endpoints should be configured to use a reduced video quality to accommodate. For context and reference, with a static background (i.e. not a movie or sporting event video, but rather just a person talking in front of a camera), basic 720p HD video calls can happen with less than 2Mbps and good SD calls can happen with as little as 384 -768Kbps.

3. LPR (Lost Packet Recovery) – This is a Poly proprietary feature that attempts to recover lost packets at the cost of network bandwidth and some extra packet latency. If both Poly A/V endpoints (audio phones, video phones, room systems, Polycom RMX MCUs, etc.) in a call support the LPR feature (and have it configured active), the LPR feature detects any packet loss during a call and automatically attempt to recover the lost packets. In most environments LPR adds significant improvements;, however, in some cases LPR may cause an increase in overhead. For this reason, LPR may be disabled for troubleshooting purposes in any Poly device. See individual product documentation or contact Poly support for details. For more information: https://www.polycom.com/collaboration/innovations/lost-packet-recovery.html.

4. DBA (Dynamic Bandwidth Allocation) – This is a Poly proprietary feature that automatically and gradually decreases the video bit-rate (video quality) if packet loss is detected between two Poly endpoints (including the Polycom RMX MCU). The feature also automatically increases the video bit-rate, to configured call levels, if the packet loss condition later subsides. DBA may be disabled in any Poly device. See individual product documentation or contact Poly Support for details.

5. LPR + DBA + PVEC – Between Poly products with all three features configured active (LPR = recovers lost packets, DBA = prevents packet loss, PVEC = hides problems from loss packets), audio or video network packet losses of up to 10% can still provide acceptable call quality and packet losses of 3% or less are usually unnoticeable on most networks (per independent testing, see https://www.ivci.com/wp-content/uploads/2016/07/whitepaper-video-packet-loss-recovery-evaluation-1.pdf).

6. Employ Polycom DMA system Bandwidth Limitation Features – The Polycom DMA Callserver product (SIP registrar + proxy, H323 gatekeeper, virtual meeting room conference call manager) contains a feature that models a network topology and controls the maximum bandwidth of calls over segments of that network. Using this feature, the network administrator can ensure no single call overloads and congests a network segment. This feature can help prevent network congestion from occurring. See the Polycom DMA administration guide regarding Site Topology and Bandwidth Limitation features or contact your Poly support representative for details.

7. Polycom RMX MCU, AVC Mitigation for Content Issues Over a Lossy Network – Because of how shared content video is encoded and transmitted during a conference, one bad/noisy endpoint in a call can negatively impact the content video quality for other conference participants. The Polycom RMX MCU has 3 configuration system flags that can be modified to help mitigate these effects:

    a. ***MAX_INTRA_REQUESTS_PER_INTERVAL_CONTENT -*** The maximum number of refresh (intra) requests per 10-second intervals allowed for an endpoint. Beyond that number, content sent by this participant is identified as "noisy," and it's refresh requests are suspended. *Default setting: 3*

    b. ***CONTENT_SPEAKER_INTRA_SUPPRESSION_IN_SECONDS -*** The interval, in seconds, between content refresh (intra) requests sent from the MCU to the content sender due to refresh requests initiated by other conference participants. Additional refresh requests received within that interval are deferred to the next interval. *Default setting: 5*

    c. **MAX_INTRA_SUPPRESSION_DURATION_IN_SECONDS_CONTENT -** The duration in seconds to ignore the participant's requests to refresh the Content display *Default setting: 10*

    If shared video content quality issues are experienced, an example set of settings could be:
    MAX_INTRA_REQUESTS_PER_INTERVAL_CONTENT= 2,
    CONTENT_SPEAKER_INTRA_SUPPRESSION_IN_SECONDS = 10,
    MAX_INTRA_SUPPRESSION_DURATION_IN_SECONDS_CONTENT = 10.

    For details and/or help configuring these settings please contact your Poly support representative.

8. Finally, if a video endpoint participant notices call quality issues during a call, the participant can mute their video. Poly video endpoints provide a function to deactivate/mute a participant's video stream. This action allows for the maximum available bandwidth for audio and content video and may help to alleviate temporary network issues.

Conclusion

By following general networking best-practices and employing the use of specially designed Poly product features, overall call quality can be increased and a high-quality of audio/video services can be deployed. For help with networks or Poly products specifically, please contact a Poly support representative for assistance.

References

Poly/Polycom Publications (2008) *Lost Packet Recovery* https://www.polycom.com/collaboration/innovations/lost-packet-recovery.html

Wainhouse Research (2008, January) *Polycom's Lost Packet Recovery (LPR) Capability* https://www.ivci.com/wp-content/uploads/2016/07/whitepaper-video-packet-loss-recovery-evaluation-1.pdf

Gardiner, Berni (2008) *QoS: What Is It? What Do We Need It?* https://www.ivci.com/wp-content/uploads/2016/07/whitepaper-qos-what-is-it-why-do-we-need-it-global-knowledge-1.pdf

Grech, Mat (2018, December 20) *Acceptable Jitter & Latency for VoIP: Everything You Need to Know*. https://getvoip.com/blog/2018/12/20/acceptable-jitter-latency/

Rivenes, Logan (2016, June 22) *What is Acceptable Jitter* https://datapath.io/resources/blog/what-is-acceptable-jitter/

CISCO Whitepapers (2013, June 17) *Best Practices in Core Network Capacity Planning* Document ID: 1506945661136136 https://www.cisco.com/c/en/us/products/collateral/routers/wan-automation-engine/white_paper_c11-728551.html

CISCO Whitepapers (2006, February 2) *Understanding Jitter in Packet Voice Networks (Cisco IOS Platforms)* Document ID: 18902 https://www.cisco.com/c/en/us/support/docs/voice/voice-quality/18902-jitter-packet-voice.html

Lynch, Kathryn (2009, July 29) *Is your network ready to handle videoconferencing?* https://www.techrepublic.com/blog/data-center/is-your-network-ready-to-handle-videoconferencing/

Wilson, Marc (2019, March 20) *Packet Loss – What is it, How to Diagnose and Fix It in your Network* https://www.pcwdld.com/packet-loss

Robert, Jason & Patterson Carey (2012, April) *Video Encoding Settings for H.246 Excellence* http://www.lighterra.com/papers/videoencodingh264/

Also see the Poly/Polycom product document library for specific products: https://documents.polycom.com/ It includes guides, papers, and manuals for all Poly/Polycom/Plantronics products including endpoints and infrastructure.

