

Mops
Internet-Draft
Intended status: Informational
Expires: January 8, 2020

J. Holland
Akamai Technologies, Inc.
July 07, 2019

Taxonomy of Issues in Internet Media
draft-jholland-mops-taxonomy-00

Abstract

This document provides a taxonomy of networking issues that pertain to quality of experience in delivery of video or other high-bitrate media over the internet.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 8, 2020.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Bandwidth Provisioning	3
2.1. Scaling Requirements for Media Delivery	3
2.1.1. Video Bitrates	3
2.1.2. Virtual Reality Bitrates	3
2.2. Path Requirements	4
2.3. Caching Systems	4
2.4. Predictable Usage Profiles	4
3. Adaptive Bit Rate	5
3.1. Overview	5
3.2. Segmented Delivery	5
3.2.1. Idle Time Between Segments	5
3.2.2. Head of Line Blocking	6
3.3. Unreliable Transport	6
4. IANA Considerations	6
5. Security Considerations	7
6. References	7
6.1. Normative References	7
6.2. Informative References	7
Author's Address	8

1. Introduction

As the internet has grown, an increasingly large share of the traffic delivered to end users has become video. Estimates put the total share of internet video traffic at 75% in 2019, expected to grow to 82% by 2022. What's more, this estimate projects the gross volume of video traffic will more than double during this time, based on a compound annual growth rate continuing at 34% (from Appendix D of [CVNI]).

In many contexts, video traffic can be handled transparently as generic application-level traffic. However, as the volume of video traffic continues to grow, it's becoming increasingly important to consider the effects of network design decisions on application-level performance, with considerations for the impact on video delivery.

This document aims to provide a taxonomy of networking issues as they relate to quality of experience in internet video delivery. The focus is on capturing characteristics of video delivery that have surprised network designers or transport experts without specific video expertise, since these highlight key differences between common assumptions in existing networking documents and observations of video delivery issues in practice.

Making specific recommendations for mitigating these issues is out of scope (though some existing mitigations might be mentioned in passing). The intent is to provide a point of reference for future solution proposals to use in describing how new technologies address or avoid these existing observed problems.

2. Bandwidth Provisioning

2.1. Scaling Requirements for Media Delivery

2.1.1. Video Bitrates

Video bit-rate selection depends on many variables. Different providers give different guidelines, but an equation that approximately matches the bandwidth requirement estimates from several video providers is given in [MSOD]:

$$\text{Kbps} = (\text{HEIGHT} * \text{WIDTH} * \text{FRAME_RATE}) / (7 * 1024)$$

Height and width are in pixels, and frame rate in frames per second. The actual bit-rate required for a specific video will also depend on the codec used and some other characteristics of the video itself, such as the frequency of high-detail motion, which may influence the compressability of the content, but this equation provides a rough estimate.

Here are a few common resolutions used for video content, with their minimum per-user bandwidth requirements according to this formula:

Name	Width x Height	Approximate Bit-rate for 60fps
DVD	720 x 480	3 Mbps
720p	1280 x 720	8 Mbps
1080p	1920 x 1080	18 Mbps
2160p (4k)	3840 x 2160	70 Mbps

2.1.2. Virtual Reality Bitrates

TBD: Reference and/or adapt content from expired work-in-progress [I-D.han-icrg-arvr-transport-problem].

The punchline is that it starts at a bare minimum of 22 Mbps mean with a 130 Mbps peak rate, up to 3.3 Gbps mean with 38 Gbps peak for high-end technology.

2.2. Path Requirements

The bit-rate requirements in Section 2.1 are per end-user actively consuming a media feed, so in the worst case, the bit-rate demands can be multiplied by the number of simultaneous users to find the bandwidth requirements for a router on the delivery path with that number of users downstream. For example, at a node with 10,000 downstream users simultaneously consuming video streams, approximately up to 180 Gbps would be necessary in order for all of them to get 1080p resolution at 60 fps.

However, when there is some overlap in the feeds being consumed by end users, it is sometimes possible to reduce the bandwidth provisioning requirements for the network by performing some kind of replication within the network. This can be achieved via object caching with delivery of replicated objects over individual connections, and/or by packet-level replication using multicast.

To the extent that replication of popular content can be performed, bandwidth requirements at peering or ingest points can sometimes be reduced to a per-feed requirement instead of a per-user requirement.

2.3. Caching Systems

TBD: pros and cons of caching decisions at different locations within the network?

Peak vs. average provisioning, and effects on peering point congestion under peak load?

Provisioning issues for caching systems?

2.4. Predictable Usage Profiles

TBD: insert charts showing historical relative data usage patterns with error bars by time of day in consumer networks?

Cross-ref vs. video quality by time of day in practice for some case study? Not sure if there's a good way to capture a generalized insight here, but it seems worth making the point that demand projections can be used to help with e.g. power consumption with routing architectures that provide for modular scalability.

3. Adaptive Bit Rate

3.1. Overview

Adaptive Bit-Rate (ABR) is a sort of application-level congestion response strategy in which the receiving media player attempts to detect the available bandwidth of the network path by experiment or by observing the successful application-layer download speed, then chooses a video bitrate that fits within that bandwidth, typically adjusting as changes in available bandwidth occur in the network.

The choice of bit-rate occurs within the context of optimizing for some metric monitored by the video player, such as highest achievable video quality, or lowest rate of expected rebuffering events.

3.2. Segmented Delivery

ABR strategies are commonly implemented by video players using HLS [RFC8216] or DASH [DASH] to perform a reliable segment delivery of video data over HTTP. Different player implementations and receiving devices use different strategies, often proprietary algorithms, to perform the bit-rate selection and available bandwidth estimation.

This kind of bandwidth-detection system can experience trouble in several ways that can be affected by networking design choices.

3.2.1. Idle Time Between Segments

When the bit-rate selection is successfully chosen below the available capacity of the network path, the response to a segment request will complete in less absolute time than the video bit-rate speed.

The resulting idle time within the connection carrying the segments has a few surprising consequences:

- o Mobile flow-bandwidth spectrum and timing mapping.
- o TCP Slow-start when restarting after idle requires multiple RTTs to re-establish a throughput at the network's available capacity. On high-RTT paths or with small enough segments, this can produce a falsely low application-visible measurement of the available network capacity.

3.2.2. Head of Line Blocking

In the event of a lost packet on a TCP connection with SACK support (a common case for segmented delivery in practice), loss of a packet can provide a confusing bandwidth signal to the receiving application. Because of the sliding window in TCP, many packets may be accepted by the receiver without being available to the application until the missing packet arrives. Upon arrival of the one missing packet after retransmit, the receiver will suddenly get access to a lot of data at the same time.

To a receiver measuring bytes received per unit time at the application layer, and interpreting it as an estimate of the available network bandwidth, this appears as a high jitter in the goodput measurement.

Active Queue Management (AQM) systems such as PIE [RFC8033] or variants of RED [RFC2309] that induce early random loss under congestion can mitigate this by using ECN [RFC3168] where available. ECN provides a congestion signal and induce a similar backoff in flows that use Explicit Congestion Notification-capable transport, but by avoiding loss avoids inducing head-of-line blocking effects in TCP connections.

3.3. Unreliable Transport

In contrast to segmented delivery, several applications use UDP or unreliable SCTP to deliver RTP or raw TS-formatted video.

Under congestion and loss, this approach generally experiences more video artifacts with fewer delay or head of line blocking effects. Often one of the key goals is to reduce latency, to better support applications like video conferencing, or for other live-action video with interactive components, such as some sporting events.

Congestion avoidance strategies for this kind of deployment vary widely in practice, ranging from some streams that are entirely unresponsive to using feedback signaling to change encoder settings (as in [RFC5762]), or to use fewer enhancement layers (as in [RFC6190]), to proprietary methods for detecting quality of experience issues and cutting off video.

4. IANA Considerations

This document requires no actions from IANA.

5. Security Considerations

This document introduces no new security issues.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

6.2. Informative References

- [CVNI] Cisco Systems, Inc., "Cisco Visual Networking Index: Forecast and Trends, 2017-2022 White Paper", February 2019, <<https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-cl1-741490.html>>.
- [DASH] "Information technology -- Dynamic adaptive streaming over HTTP (DASH) -- Part 1: Media presentation description and segment formats", ISO/IEC 23009-1:2014, n.d..
- [I-D.han-iccr-g-arvr-transport-problem] Han, L. and K. Smith, "Problem Statement: Transport Support for Augmented and Virtual Reality Applications", draft-han-iccr-g-arvr-transport-problem-01 (work in progress), March 2017.
- [MSOD] Akamai Technologies, Inc., "Media Services On Demand: Encoder Best Practices", n.d., <<https://learn.akamai.com/en-us/webhelp/media-services-on-demand/media-services-on-demand-encoder-best-practices/GUID-7448548A-A96F-4D03-9E2D-4A4BBB6EC071.html>>.
- [RFC2309] Braden, B., Clark, D., Crowcroft, J., Davie, B., Deering, S., Estrin, D., Floyd, S., Jacobson, V., Minshall, G., Partridge, C., Peterson, L., Ramakrishnan, K., Shenker, S., Wroclawski, J., and L. Zhang, "Recommendations on Queue Management and Congestion Avoidance in the Internet", RFC 2309, DOI 10.17487/RFC2309, April 1998, <<https://www.rfc-editor.org/info/rfc2309>>.

- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <<https://www.rfc-editor.org/info/rfc3168>>.
- [RFC5762] Perkins, C., "RTP and the Datagram Congestion Control Protocol (DCCP)", RFC 5762, DOI 10.17487/RFC5762, April 2010, <<https://www.rfc-editor.org/info/rfc5762>>.
- [RFC6190] Wenger, S., Wang, Y., Schierl, T., and A. Eleftheriadis, "RTP Payload Format for Scalable Video Coding", RFC 6190, DOI 10.17487/RFC6190, May 2011, <<https://www.rfc-editor.org/info/rfc6190>>.
- [RFC8033] Pan, R., Natarajan, P., Baker, F., and G. White, "Proportional Integral Controller Enhanced (PIE): A Lightweight Control Scheme to Address the Bufferbloat Problem", RFC 8033, DOI 10.17487/RFC8033, February 2017, <<https://www.rfc-editor.org/info/rfc8033>>.
- [RFC8216] Pantos, R., Ed. and W. May, "HTTP Live Streaming", RFC 8216, DOI 10.17487/RFC8216, August 2017, <<https://www.rfc-editor.org/info/rfc8216>>.

Author's Address

Jake Holland
Akamai Technologies, Inc.
150 Broadway
Cambridge, MA 02144
United States of America

Email: jakeholland.net@gmail.com