Authors: J. Holland   A. Begen
         Akamai Technologies, Inc.   Networked Media
         S. Dawkins
         Tencent America LLC

Operational Considerations for Streaming Media

Abstract

   This document provides an overview of operational networking issues
   that pertain to quality of experience in streaming of video and
   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 12 November 2021.

Table of Contents

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]).

A substantial part of this growth is due to increased use of
streaming video, although the amount of video traffic in real-time
communications (for example, online videoconferencing) has also
grown significantly. While both streaming video and
videoconferencing have real-time delivery and latency requirements,
these requirements vary from one application to another. For
example, videoconferencing demands an end-to-end (one-way) latency
of a few hundreds of milliseconds whereas live streaming can
tolerate latencies of several seconds.

This document specifically focuses on the streaming applications and defines streaming as follows: Streaming is transmission of a continuous media from a server to a client and its simultaneous consumption by the client. Here, continous media refers to media and associated streams such as video, audio, metadata, etc. In this definition, the critical term is "simultaneous", as it is not considered streaming if one downloads a video file and plays it after the download is completed, which would be called download-and-play. This has two implications. First, server's transmission rate must (loosely or tightly) match to client's consumption rate for an uninterrupted playback. That is, the client must not run out of data (buffer underrun) or take more than it can keep (buffer overrun) as any excess media is simply discarded. Second, client's consumption rate is limited by not only bandwidth availability but also the real-time constraints. That is, the client cannot fetch media that is not available yet.

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 are 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.

1.1. Notes for Contributors and Reviewers

Note to RFC Editor: Please remove this section and its subsections before publication.

This section is to provide references to make it easier to review the development and discussion on the draft so far.

1.1.1. Venues for Contribution and Discussion

This document is in the Github repository at:

https://github.com/ietf-wg-mops/draft-ietf-mops-streaming-opcons

Readers are welcome to open issues and send pull requests for this document.

Substantial discussion of this document should take place on the MOPS working group mailing list (mops@ietf.org).

   *Join: https://www.ietf.org/mailman/listinfo/mops

   *Search: https://mailarchive.ietf.org/arch/browse/mops/

1.1.2. Template for Contributions

Contributions are solicited regarding issues and considerations that have an impact on media streaming operations.

Please note that contributions may be merged and substantially edited, and as a reminder, please carefully consider the Note Well before contributing: https://datatracker.ietf.org/submit/note-well/

Contributions can be emailed to mops@ietf.org, submitted as issues to the issue tracker of the repository in Section 1.1.1, or emailed to the document authors at draft-ietf-mops-streaming-opcons@ietf.org.

Contributors describing an issue not yet addressed in the draft are requested to provide the following information, where applicable:

   *a suggested title or name for the issue

   *a long-term pointer to the best reference describing the issue

   *a short description of the nature of the issue and its impact on
    media quality of service, including:

      -where in the network this issue has root causes

      -who can detect this issue when it occurs

   *an overview of the issue's known prevalence in practice. pointers
    to write-ups of high-profile incidents are a plus.

   *a list of known mitigation techniques, with (for each known
    mitigation):

      -a name for the mitigation technique

      -a long-term pointer to the best reference describing it

-a short description of the technique:

        owhat it does

        owhere in the network it operates

        oan overview of the tradeoffs involved-how and why it's
         helpful, what it costs.

    -supplemental information about the technique's deployment
     prevalence and status

1.1.3. History of Public Discussion

   Presentations:

    *IETF 105 BOF:

     https://www.youtube.com/watch?v=4G3YBVmn9Eo&t=47m21s

    *IETF 106 meeting:

     https://www.youtube.com/watch?v=4_k340xT2jM&t=7m23s

    *MOPS Interim Meeting 2020-04-15:

     https://www.youtube.com/watch?v=QExiajdC0IY&t=10m25s

    *IETF 108 meeting:

     https://www.youtube.com/watch?v=ZaRsk0y3O9k&t=2m48s

    *MOPS 2020-10-30 Interim meeting:

     https://www.youtube.com/watch?v=vDZKspv4LXw&t=17m15s

2. Bandwidth Provisioning

2.1. Scaling Requirements for Media Delivery

2.1.1. Video Bitrates

   Video bitrate 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]:

Kbps = (HEIGHT * WIDTH * FRAME_RATE) / (MOTION_FACTOR * 1024)

   Height and width are in pixels, frame rate is in frames per second,
   and the motion factor is a value that ranges from 20 for a low-

motion talking heads video to 7 for sports, and content with a lot
of screen changes.

The motion factor captures the variability in bitrate due to the
amount and frequency of high-detail motion, which generally
influences the compressability of the content.

The exact bitrate required for a particular video also depends on a
number of specifics about the codec used and how the codec-specific
tuning parameters are matched to the content, but this equation
provides a rough estimate that approximates the usual bitrate
characteristics using the most common codecs and settings for
production traffic.

Here are a few common resolutions used for video content, with their
typical and peak per-user bandwidth requirements for 60 frames per
second (FPS):

Name	Width x Height	Typical	Peak
DVD	720 x 480	1.3 Mbps	3 Mbps
720p (1K)	1280 x 720	3.6 Mbps	5 Mbps
1080p (2K)	1920 x 1080	8.1 Mbps	18 Mbps
2160p (4k)	3840 x 2160	32 Mbps	70 Mbps

Table 1

2.1.2. Virtual Reality Bitrates

Even the basic virtual reality (360-degree) videos (that allow users
to look around freely, referred to as three degrees of freedom -
3DoF) require substantially larger bitrates when they are captured
and encoded as such videos require multiple fields of view of the
scene. The typical multiplication factor is 8 to 10. Yet, due to
smart delivery methods such as viewport-based or tiled-based
streaming, we do not need to send the whole scene to the user.
Instead, the user needs only the portion corresponding to its
viewpoint at any given time.

In more immersive applications, where basic user movement (3DoF+) or
full user movement (6DoF) is allowed, the required bitrate grows
even further. In this case, the immersive content is typically
referred to as volumetric media. One way to represent the volumetric
media is to use point clouds, where streaming a single object may
easily require a bitrate of 30 Mbps or higher. Refer to [MPEGI] and
[PCC] for more details.

2.2. Path Requirements

The bitrate requirements in Section 2.1 are per end-user actively
consuming a media feed, so in the worst case, the bitrate 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 80 Gbps would be necessary in order for all of them to get typical content at 1080p resolution at 60 fps, or up to 180 Gbps to get sustained high-motion content such as sports, while maintaining the same resolution.

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 be reduced to as low as a per-feed requirement instead of a per-user requirement.

2.3. Caching Systems

When demand for content is relatively predictable, and especially when that content is relatively static, caching content close to requesters, and pre-loading caches to respond quickly to initial requests, is often useful (for example, HTTP/1.1 caching is described in [RFC7234]). This is subject to the usual considerations for caching - for example, how much data must be cached to make a significant difference to the requester, and how the benefits of caching and pre-loading caches balances against the costs of tracking "stale" content in caches and refreshing that content.

It is worth noting that not all high-demand content is also "live" content. One popular example is when popular streaming content can be staged close to a significant number of requesters, as can happen when a new episode of a popular show is released. This content may be largely stable, so low-cost to maintain in multiple places throughout the Internet. This can reduce demands for high end-to-end bandwidth without having to use mechanisms like multicast.

Caching and pre-loading can also reduce exposure to peering point congestion, since less traffic crosses the peering point exchanges if the caches are placed in peer networks, and could be pre-loaded during off-peak hours, using "Lower-Effort Per-Hop Behavior (LE PHB) for Differentiated Services" [RFC8622], "Low Extra Delay Background Transport (LEDBAT)" [RFC6817], or similar mechanisms.

All of this depends, of course, on the ability of a content provider to predict usage and provision bandwidth, caching, and other

mechanisms to meet the needs of users. In some cases (Section 2.4), this is relatively routine, but in other cases, it is more difficult (Section 2.5, Section 2.6).

2.4. Predictable Usage Profiles

Historical data shows that users consume more video and videos at higher bitrates than they did in the past on their connected devices. Improvements in the codecs that help with reducing the encoding bitrates with better compression algorithms could not have offset the increase in the demand for the higher quality video (higher resolution, higher frame rate, better color gamut, better dynamic range, etc.). In particular, mobile data usage has shown a large jump over the years due to increased consumption of entertainement as well as conversational video.

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.

2.5. Unpredictable Usage Profiles

Although TCP/IP has been used with a number of widely used applications that have symmetric bandwidth requirements (similar bandwidth requirements in each direction between endpoints), many widely-used Internet applications operate in client-server roles, with asymmetric bandwidth requirements. A common example might be an HTTP GET operation, where a client sends a relatively small HTTP GET request for a resource to an HTTP server, and often receives a significantly larger response carrying the requested resource. When HTTP is commonly used to stream movie-length video, the ratio between response size and request size can become quite large.

For this reason, operators may pay more attention to downstream bandwidth utilization when planning and managing capacity. In addition, operators have been able to deploy access networks for end users using underlying technologies that are inherently asymetric, favoring downstream bandwidth (e.g. ADSL, cellular technologies, most IEEE 802.11 variants), assuming that users will need less upstream bandwidth than downstream bandwidth. This strategy usually works, except when it does not, because application bandwidth usage patterns have changed.

One example of this type of change was when peer-to-peer file sharing applications gained popularity in the early 2000s. To take

one well-documented case ([RFC5594]), the Bittorrent application
created "swarms" of hosts, uploading and downloading files to each
other, rather than communicating with a server. Bittorrent favored
peers who uploaded as much as they downloaded, so that new
Bittorrent users had an incentive to significantly increase their
upstream bandwidth utilization.

The combination of the large volume of "torrents" and the peer-to-
peer characteristic of swarm transfers meant that end user hosts
were suddenly uploading higher volumes of traffic to more
destinations than was the case before Bittorrent. This caused at
least one large ISP to attempt to "throttle" these transfers, to
mitigate the load that these hosts placed on their network. These
efforts were met by increased use of encryption in Bittorrent,
similar to an arms race, and set off discussions about "Net
Neutrality" and calls for regulatory action.

Especially as end users increase use of video-based social
networking applications, it will be helpful for access network
providers to watch for increasing numbers of end users uploading
significant amounts of content.

2.6. Extremely Unpredictable Usage Profiles

The causes of unpredictable usage described in Section 2.5 were more
or less the result of human choices, but we were reminded during a
post-IETF 107 meeting that humans are not always in control, and
forces of nature can cause enormous fluctuations in traffic
patterns.

In his talk, Sanjay Mishra [Mishra] reported that after the CoViD-19
pandemic broke out in early 2020,

  *Comcast's streaming and web video consumption rose by 38%, with
   their reported peak traffic up 32% overall between March 1 to
   March 30,

  *AT&T reported a 28% jump in core network traffic (single day in
   April, as compared to pre stay-at-home daily average traffic),
   with video accounting for nearly half of all mobile network
   traffic, while social networking and web browsing remained the
   highest percentage (almost a quarter each) of overall mobility
   traffic, and

  *Verizon reported similar trends with video traffic up 36% over an
   average day (pre COVID-19)}.

We note that other operators saw similar spikes during this time period. Craig Labowitz [Labovitz] reported

  *Weekday peak traffic increases over 45%-50% from pre-lockdown
   levels,

  *A 30% increase in upstream traffic over their pre-pandemic
   levels, and

  *A steady increase in the overall volume of DDoS traffic, with
   amounts exceeding the pre-pandemic levels by 40%. (He attributed
   this increase to the significant rise in gaming-related DDoS
   attacks ([LabovitzDDoS]), as gaming usage also increased.)

Subsequently, the Inernet Architecture Board (IAB) held a COVID-19 Network Impacts Workshop [IABcovid] in November 2020. Given a larger number of reports and more time to reflect, the following observations from the draft workshop report are worth considering.

  *Participants describing different types of networks reported
   different kinds of impacts, but all types of networks saw
   impacts.

  *Mobile networks saw traffic reductions and residential networks
   saw significant increases.

  *Reported traffic increases from ISPs and IXPs over just a few
   weeks were as big as the traffic growth over the course of a
   typical year, representing a 15-20% surge in growth to land at a
   new normal that was much higher than anticipated.

  *At DE-CIX Frankfurt, the world's largest Internet Exchange Point
   in terms of data throughput, the year 2020 has seen the largest
   increase in peak traffic within a single year since the IXP was
   founded in 1995.

  *The usage pattern changed significantly as work-from-home and
   videoconferencing usage peaked during normal work hours, which
   would have typically been off-peak hours with adults at work and
   children at school. One might expect that the peak would have had
   more impact on networks if it had happened during typical evening
   peak hours for video streaming applications.

  *The increase in daytime bandwidth consumption reflected both
   significant increases in "essential" applications such as
   videoconferencing and VPNs, and entertainment applications as
   people watched videos or played games.

*At the IXP-level, it was observed that port utilization
       increased. This phenomenon is mostly explained by a higher
       traffic demand from residential users.

3. Adaptive Bitrate

3.1. Overview

   Adaptive BitRate (ABR) is a sort of application-level response
   strategy in which the streaming client attempts to detect the
   available bandwidth of the network path by observing the successful
   application-layer download speed, then chooses a bitrate for each of
   the video, audio, subtitles and metadata (among the limited number
   of available options) that fits within that bandwidth, typically
   adjusting as changes in available bandwidth occur in the network or
   changes in capabilities occur during the playback (such as available
   memory, CPU, display size, etc.).

   The choice of bitrate occurs within the context of optimizing for
   some metric monitored by the client, such as highest achievable
   video quality or lowest chances for a rebuffering (playback stall).

3.2. Segmented Delivery

   ABR playback is commonly implemented by streaming clients using HLS
   [RFC8216] or DASH [DASH] to perform a reliable segmented delivery of
   media over HTTP. Different implementations use different strategies
   [ABRSurvey], often proprietary algorithms (called rate adaptation or
   bitrate selection algorithms) to perform available bandwidth
   estimation/prediction and the bitrate selection. Most clients only
   use passive observations, i.e., they do not generate probe traffic
   to measure the available bandwidth.

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

3.2.1. Idle Time between Segments

   When the bitrate selection is successfully chosen below the
   available capacity of the network path, the response to a segment
   request will typically complete in less absolute time than the
   duration of the requested segment. The resulting idle time within
   the connection carrying the segments has a few surprising
   consequences:

      *Mobile flow-bandwidth spectrum and timing mapping.

      *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.

A detailed investigation of this phenomenon is available in
[NOSSDAV12].

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 videoconferencing, 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. Acknowledgements

Thanks to Mark Nottingham, Glenn Deen, Dave Oran, Aaron Falk, Kyle Rose, Leslie Daigle, Lucas Pardue, Matt Stock, Alexandre Gouaillard, and Mike English for their very helpful reviews and comments.

7. Informative References

[ABRSurvey] Abdelhak Bentaleb et al, ., "A Survey on Bitrate Adaptation Schemes for Streaming Media Over HTTP", 2019, <https://ieeexplore.ieee.org/abstract/document/8424813>.

[CVNI]   Cisco Systems, Inc, ., "Cisco Visual Networking Index: Forecast and Trends, 2017-2022 White Paper", 27 February 2019, <https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.html>.

[DASH]   "Information technology -- Dynamic adaptive streaming over HTTP (DASH) -- Part 1: Media presentation description and segment formats", ISO/IEC 23009-1:2019, 2019, <https://www.iso.org/standard/79329.html>.

[IABcovid] Jari Arkko / Stephen Farrel / Mirja Kühlewind / Colin Perkins, ., "Report from the IAB COVID-19 Network Impacts Workshop 2020", November 2020, <https://datatracker.ietf.org/doc/draft-iab-covid19-workshop/>.

[Labovitz] Labovitz, C. and Nokia Deepfield, "Network traffic insights in the time of COVID-19: April 9 update", April 2020, <https://www.nokia.com/blog/network-traffic-insights-time-covid-19-april-9-update/>.

[LabovitzDDoS] Takahashi, D. and Venture Beat, "Why the game industry is still vulnerable to DDoS attacks", May 2018, <https://venturebeat.com/2018/05/13/why-the-game-industry-is-still-vulnerable-to-distributed-denial-of-service-attacks/>.

[Mishra]   Mishra, S. and J. Thibeault, "An update on Streaming Video Alliance", 2020, <https://datatracker.ietf.org/meeting/interim-2020-mops-01/materials/slides-

interim-2020-mops-01-sessa-april-15-2020-mops-interim-an-update-on-streaming-video-alliance>.

[MPEGI]   Boyce et al, J.M., "MPEG Immersive Video Coding
            Standard", n.d., <https://ieeexplore.ieee.org/document/
            9374648>.

[MSOD]   Akamai Technologies, Inc, ., "Media Services On Demand:
            Encoder Best Practices", 2019, <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>.

[NOSSDAV12] Saamer Akhshabi et al, ., "What Happens When HTTP
            Adaptive Streaming Players Compete for Bandwidth?", June
            2012, <https://dl.acm.org/doi/10.1145/2229087.2229092>.

[PCC]   Sebastian Schwarz et al, ., "Emerging MPEG Standards for
            Point Cloud Compression", March 2019, <https://
            ieeexplore.ieee.org/document/8571288>.

[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>.

[RFC5594]   Peterson, J. and A. Cooper, "Report from the IETF
            Workshop on Peer-to-Peer (P2P) Infrastructure, May 28,
            2008", RFC 5594, DOI 10.17487/RFC5594, July 2009,
            <https://www.rfc-editor.org/info/rfc5594>.

[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.-K., 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>.

   [RFC6817]   Shalunov, S., Hazel, G., Iyengar, J., and M. Kuehlewind,
                 "Low Extra Delay Background Transport (LEDBAT)", RFC
                 6817, DOI 10.17487/RFC6817, December 2012, <https://
                 www.rfc-editor.org/info/rfc6817>.

   [RFC7234]   Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke,
                 Ed., "Hypertext Transfer Protocol (HTTP/1.1): Caching",
                 RFC 7234, DOI 10.17487/RFC7234, June 2014, <https://
                 www.rfc-editor.org/info/rfc7234>.

   [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>.

   [RFC8622]   Bless, R., "A Lower-Effort Per-Hop Behavior (LE PHB) for
                 Differentiated Services", RFC 8622, DOI 10.17487/RFC8622,
                 June 2019, <https://www.rfc-editor.org/info/rfc8622>.

Authors' Addresses

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

   Email: jakeholland.net@gmail.com

   Ali Begen
   Networked Media
   Turkey

   Email: ali.begen@networked.media

   Spencer Dawkins
   Tencent America LLC
   United States of America

   Email: spencerdawkins.ietf@gmail.com