Media OPerationS
Internet-Draft
Intended status: Informational
Expires: 23 April 2026

G. Deen
Comcast-NBCUniversal
S. Mishra
Verizon

20 October 2025

Network Overlay Impacts to Streaming Video
draft-ietf-mops-network-overlay-impacts-03

Abstract

This document examines the operational impacts on streaming video applications resulting from network policy changes introduced by network overlays. Such overlays may alter IP address assignment, transport protocols, routing behavior, or DNS resolution. These changes can, in turn, affect critical aspects of content delivery, including latency, CDN cache selection, delivery path optimization, traffic classification, and content access controls.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at https://ietf-wg-mops.github.io/draft-ietf-mops-network-overlay-impacts/draft-ietf--mops-network-overlay-impacts.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-ietf-mops-network-overlay-impacts/.

Discussion of this document takes place on the Media OPerationS Working Group mailing list (mailto:mops@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/mops/. Subscribe at https://www.ietf.org/mailman/listinfo/mops/.

Source for this draft and an issue tracker can be found at https://github.com/ietf-wg-mops/draft-ietf-mops-network-overlay-impacts.

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/.

Copyright Notice

Table of Contents

1. Introduction

Enhancing the privacy of Internet users has been a significant focus of the IETF since the Snowden disclosures and the publication of [RFC7258]. [RFC7264] explored in greater detail the technical threats identified in RFC7258, and described high-level mitigation approaches. Since then, IETF working groups have endeavored to address these specific threats, producing a wide range of new standards with built-in privacy enhancements. Protocol such as QUIC [RFC9000] is an examples of this new generation: always-enabled encryption and other protections embedded directly into the design.

Meanwhile, Internet video streaming has become part of daily life for billions of viewers. For many, streaming is the primary way to watch sports, entertainment, user-generated content (UGC) and news. It has grown to dominate Internet volume: an hour of HD video can consume approximately 1.5 2.5 GB, and streaming is estimated to account for 8085% of global Internet traffic. The operational considerations for this growth are documented in [RFC9317].

Early streaming efforts were focused simply on making video available on a device. Today' s ecosystem demands high-scale, low-latency delivery, including live events and 4K/8K streams. For prerecorded content such as Video on Demand (VOD) and UGC, distributing encoded content via Content Delivery Network (CDN) caches is a common technique for meeting scale. The IETF's CDNI working group and the Streaming Video Technology Alliance (SVTA at svta.org) have extended these architectures with services like Open Caching.

The newest frontier is live streaming— primarily around major sports events and other high-interest broadcasts. These can involve tens to hundreds of millions of viewers simultaneously and impose strict latency and scale requirements. Live delivery pipelines are highly optimized and sensitive to changes in underlying network behavior.

However, as consumer devices and services increasingly incorporate privacy-enhancing features (in response to [RFC7258] and [RFC7264]), they sometimes introduce unexpected or hard-to-detect changes in network behavior. These changes can interfere with— or even undermine— the efficiency, scaling, and low-latency architectures that streaming platforms have invested heavily to build.

The authors acknowledge the many challenges of improving Internet privacy while supporting the vast variety of Internet applications and use cases. This is difficult work, and it is natural that operational considerations must be carefully incorporated into architectural designs.

This document is intended to highlight the negative impacts that have been observed by streaming platforms, from an operational perspective, when privacy-enhancing overlays or other network-policy changes are deployed. It aims to provide insights to application developers, platform architects, and network operators into how such overlays may affect streaming video delivery.

2. Internet Privacy Enhancements

The IETF\’s efforts to strengthen Internet privacy and mitigate pervasive monitoring, as described in [RFC7258], have driven a series of architectural and protocol-level developments. The initial focus was on encrypting network data flows, most commonly through the wider adoption of Transport Layer Security (TLS). Over time, these efforts have expanded to include policy- and design-level changes— such as modifying routing paths, selecting privacy-preserving DNS resolvers, and introducing encrypted transport protocols— to better obscure and isolate user traffic from observation within the underlying network infrastructure.

The IAB’s [RFC7258] identifies pervasive monitoring as an attack on privacy, while [RFC7624] outlines potential technical and operational responses to mitigate its impact. The development of the QUIC transport protocol, defined in [RFC9000], exemplifies the application of these principles: QUIC integrates confidentiality, integrity, and authentication into the transport layer itself, ensuring that user data and most protocol metadata remain encrypted by default.

Collectively, these privacy-enhancing measures have reshaped how networks and applications interact. However, they also introduce new considerations for operational visibility, traffic management, and performance optimization, which are particularly relevant to streaming video applications.

2.1. Network Overlays

The IETF\’s privacy-enhancement efforts in response to [RFC7258] have driven a range of architectural and policy design choices, including the adoption of “always-on” encryption, as exemplified by QUIC [RFC9000]. While many such developments have minimal impact on video streaming, some introduce new behaviors that can be described as creating network overlays— logical networks that operate on top of the

underlying native network, but apply different routing, transport, or policy decisions than either the native network or the streaming application would independently choose.

Network overlays that alter policies or paths in ways not directly visible, selectable, or detectable by the streaming application or platform can have significant operational effects. These overlays may transparently modify network properties— such as source IP addresses, DNS resolver choices, or routing behavior— without the knowledge of the streaming service or end user. Such hidden policy changes can inadvertently disrupt the assumptions underlying adaptive streaming architectures, content delivery path optimization, or CDN selection mechanisms.

When a network overlay modifies connection properties in ways that differ from application expectations, the result can be mismatched assumptions between the application and the actual transport environment. This disconnect may cause degraded performance, misclassification of network paths, or unexpected latency and throughput characteristics, all of which affect streaming quality and operational predictability.

Protocols such as MASQUE [RFC9484] and services built on it such as Apple’s iCloud Private Relay (https://www.apple.com/privacy/docs/iCloud_Private_Relay_Overview_Dec2021.PDF) illustrate privacy-enhancing network overlays that deliberately alter connection policies relative to the open Internet. While beneficial for user privacy, such mechanisms can also obscure the visibility and control that streaming services rely on for consistent content delivery and Quality of Experience (QoE) management.

2.1.1. Emerging Operational Issues with Network Overlay Policy Changes

Streaming video applications and content delivery platforms are increasingly encountering operational challenges associated with network overlays. These challenges arise when overlays introduce policy changes that are unexpected, inconsistently applied, or difficult— or even impossible— for the streaming platform to detect or adapt to in real time. While the specific impacts vary depending on the overlay’s design and implementation, several common classes of operational issues have been observed across deployments. These include mismatches in routing and cache selection, unexpected transport-layer behavior, and inconsistencies in latency or throughput reporting that affect Quality of Experience (QoE) monitoring and optimization.

2.2. Policy Changes

Changes to network policies introduced by overlays can alter the expected behavior of streaming applications in several ways.

For example, an overlay that modifies encryption policies— such as transforming HTTP URLs in manifests into HTTPS connections— can disrupt architectures that rely on the network' s ability to identify or classify video flows. In such cases, the visibility of traffic used for caching, optimization, or QoS treatment may be reduced or lost entirely.

Similarly, overlays that alter routing policies can interfere with the Content Delivery Network (CDN) cache selection logic used by streaming platforms. A change in routing path may cause the application to connect to a more distant cache, resulting in higher latency, lower throughput, and degraded video quality, even when a closer cache would otherwise have been selected.

An example of a routing policy change is illustrated in Figure 1, showing how a network overlay can apply a routing policy that diverges from that of the underlying base network, resulting in a modified traffic path and different delivery characteristics.

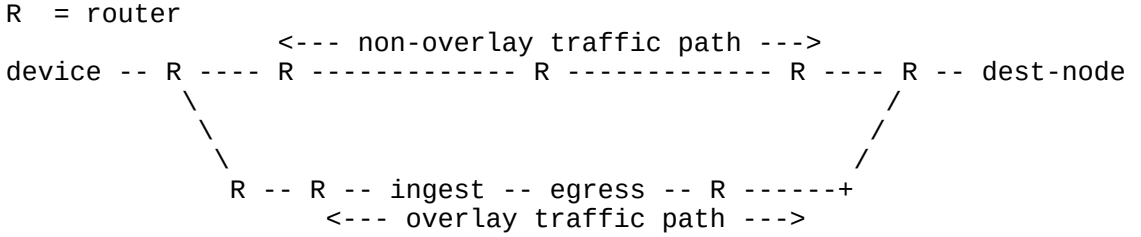

Figure 1: Network Overlay routing selects traffic via an alternate path

2.2.1. Partitioning

Network Overlay policy changes often include the use of alternate routing policies, as a core element of their design involves tunneling connections through different network paths to enhance user privacy and reduce tracking. This architectural concept— partitioning— is further discussed in the IAB document Partitioning as an Architecture for Privacy (https://datatracker.ietf.org/doc/draft-iab-privacy-partitioning/). By isolating traffic and obscuring its correlation with the underlying native network, partitioning helps defend against pervasive monitoring and traffic analysis.

While effective for privacy protection, these routing partitions can also alter network visibility and path selection in ways that affect streaming video performance, such as cache selection accuracy, latency, and adaptive bitrate (ABR) responsiveness.

2.2.2. Protocol Policy Changes

Network overlays have been observed to alter application and transport protocol changes from those originally selected by the streaming application. In some cases, privacy-enhancing or optimization mechanisms automatically translate connections — for example, converting HTTP/2 over TCP into HTTP/3 over QUIC, or upgrading HTTP/2 sessions to HTTPS with TLS encryption. Such conversions are typically performed to enforce stronger privacy, security, or efficiency policies, but they may occur without visibility or control by the streaming application.

A key operational impact arises when protocol substitution changes the network characteristics perceived by the video application. For example, a video application may perform a preliminary fetch to measure network conditions before selecting an appropriate bitrate for content delivery. If the application's test probe uses HTTP/2 over TCP, but the subsequent content request is transparently converted by the overlay to HTTP/3 over QUIC, the measured results no longer reflect the actual transport path. This mismatch can lead to inaccurate bandwidth estimation, causing the adaptive bitrate (ABR) algorithm to select non-optimal streaming parameters and degrade user experience.

2.2.3. Encryption Policy

Changes to the encryption policy applied to video streams — whether by adding encryption where it was not originally used, or by removing or terminating encryption where it was expected — can introduce significant operational challenges for streaming applications and delivery networks.

In some cases, network overlays or privacy-enhancing systems may automatically enforce encryption, converting plaintext HTTP video traffic into HTTPS or encapsulating transport flows within encrypted tunnels. While this improves confidentiality, it can also obscure traffic classification and disable optimizations that rely on visibility into flow metadata, such as CDN cache selection, adaptive bitrate tuning, or Quality-of-Service (QoS) marking.

Conversely, if encryption is removed or terminated prematurely, such as through a proxy that decrypts and re-encrypts video traffic, it can violate end-to-end security assumptions made by the application or CDN, potentially exposing content or user data to unauthorized inspection.

In both cases, mismatched encryption policies between the streaming application, CDN, and the underlying network can lead to reduced performance, incorrect cache usage, or inconsistent delivery behavior.

2.2.3.1. Forced Encryption Upgrade

Enforcing encryption upgrades — for example, converting unencrypted HTTP/2 traffic into HTTP/2 over TLS (HTTPS) — can disrupt streaming workflows that rely on the network's ability to inspect or classify content as part of the delivery process. When network visibility into streaming flows is removed, content-aware optimizations such as CDN cache selection, multicast distribution, or traffic prioritization may fail to function as designed. As a result, video traffic may be misclassified as generic encrypted data, leading to incorrect policy enforcement or suboptimal delivery behavior.

This issue is particularly significant in mobile and multicast-based environments, where network-assisted detection of video streams is often required to achieve efficient bandwidth utilization and maintain quality of experience. In such cases, forced encryption upgrades may prevent the network from applying appropriate delivery optimizations, resulting in degraded performance or increased operational complexity.

2.2.3.2. Forced Encryption Downgrade

Conversely, removal or termination of encryption originally applied by a streaming platform can introduce serious operational and security concerns. In many streaming architectures, transport-level encryption (e.g., HTTPS or QUIC) is not only used to ensure confidentiality but also forms an integral part of the content protection and integrity assurance mechanisms.

When an intermediate network overlay or proxy terminates TLS sessions or otherwise downgrades an encrypted connection to plaintext, it can invalidate end-to-end trust assumptions between the client, CDN, and content provider. Such behavior may expose sensitive metadata, enable unauthorized content inspection or modification, and violate Digital Rights Management (DRM).

In effect, a forced encryption downgrade undermines both security and operational reliability, leading to potential playback failures, content delivery errors, or loss of user trust.

2.2.4. Address Policy Changes

Network overlays that modify IP addressing policies— such as converting IPv4 to IPv6, IPv6 to IPv4, or reassigning source IP addresses— can introduce a range of operational challenges for streaming platforms, particularly when these changes occur unexpectedly or are invisible to the application. Such address changes can disrupt routing decisions, CDN cache selection, and traffic localization processes that depend on stable endpoint addresses. They also complicate diagnostic and troubleshooting efforts, as engineers analyzing logs, performing test probes, or correlating session data may inadvertently use incorrect or outdated IP information.

Source IP Address assignment changes, again when done invisibly to the application can cause significant disruption. Platform authentication gateways that associate session authorizations with the session’s device’s IP address can result in service access denial when associated addresses change unexpectedly. For example, when the device address as seen by the video application is different from the device address seen by the associated streaming platform, this can result in the platform rejecting logins, content access and other service functions from the device.

A related issue arises when the source IP address observed by the streaming platform differs from that seen by the client application or device. Because many streaming architectures use IP-based session binding— such as platform authentication gateways that associate user or device authorization with a specific IP address— unannounced address translation can result in service access failures, login rejections, or denied content delivery. For example, when an overlay reassigns or masks the client’ s IP address, the streaming platform may interpret this as a new or unauthorized connection, even though the client session remains active. This mismatch can lead to intermittent playback interruptions, degraded user experience, or increased operational complexity for both service providers and network operators.

2.2.5. DNS Policy Changes

Network overlays that modify DNS resolver settings or redirect DNS queries can have significant implications for Content Delivery Networks (CDNs) that rely on DNS-based load balancing for cache selection and traffic localization.

Many CDN architectures determine the “best” cache for a client by observing the source IP address of the DNS resolver making the request. When an overlay substitutes or masks the resolver—either intentionally or as part of privacy-enhancing policies—the CDN may incorrectly infer the client’s location, resulting in non-optimal cache selection, increased latency, or reduced video quality.

2.2.5.1. EDNS0

The EDNS(0) (Extension Mechanisms for DNS, [RFC6891]) extension was introduced to allow resolvers to include additional client subnet information in DNS queries, improving CDN cache selection accuracy. If a network overlay redirects DNS queries to a resolver that does not support EDNS(0) or deliberately strips this information, the CDN loses critical context for determining the most appropriate edge cache. This can lead to the selection of a distant or overloaded cache, negatively impacting video startup time, buffering, and overall user experience.

2.2.6. Log Data Changes

Accurate and consistent logging is essential for diagnosing streaming performance and operational issues. Network overlays that alter connection properties—such as DNS resolvers, IP addresses, or transport protocols—can cause log entries to differ between the client device and the streaming platform. When such discrepancies occur, engineers attempting to correlate logs for troubleshooting may misinterpret session behavior or fail to identify the true source of a problem. Unexpected or misleading log data therefore undermines both problem determination and root-cause analysis, complicating operational monitoring and incident response workflows.

2.2.7. Geo Location & Identification

Network overlays that alter the apparent source location of user devices can interfere with streaming platforms’ ability to accurately determine geospatial attributes such as country, region, or network domain.

Many CDNs and content providers rely on IP addressbased geolocation to enforce regional content licensing, apply local regulations, or select nearby caches for optimal performance. When an overlay substitutes or masks the client’s IP address—presenting it as originating from a different region or outside of known geolocation mappings—the platform may be unable to correctly associate the user with their actual location.

This can result in users being denied access to region-restricted content that they would otherwise be authorized to view, or being directed to distant CDN caches, causing degraded video quality and higher latency.

In addition, such location ambiguity complicates analytics, fraud detection, and rights management processes that depend on consistent geographic identifiers.

2.2.8. CDN interconnection troubleshooting

In CDN interconnection scenarios, when two CDN domains collaborate to localize a point of failure, they typically begin by identifying the delivery path and selecting observation points along that path to take diagnostic measurements. Through iterative testing. they narrow down the problem domain to isolate the failure's location.

However, when network overlays alter routing behavior, this process becomes unreliable. Since CDNs depend on their request routing information to determine where along the delivery path measurements should be taken, the presence of an overlay that reroutes or tunnels traffic means that the expected observation point no longer lies on the actual traffic path. As a result, the flow cannot be observed where the CDN expects it to be, making fault localization and coordination between interconnecting CDNs significantly more difficult.

2.2.9. Routing Changes

Routing changes introduced by network overlays can alter the expected path between video applications and the infrastructure services they rely on. Such changes may cause a wide range of operational problems — including degraded performance, inconsistent latency, or failures in CDN cache selection and session persistence.

When routing behavior differs from what the video platform or application expects, content delivery optimizations such as proximity-based cache selection, adaptive bitrate decisions, and transport-layer congestion management can become ineffective. These effects can be difficult to detect, as the overlay's routing policy is often not visible to the streaming application or operators monitoring network performance.

2.2.9.1. End to End Problem Discovery

A common issue in video delivery is locating where along the delivery path the video transport is encountering problems. Often such problems are more complex than the connection not working at but instead involve identifying bottlenecks, lost packets, and congestion issues. When the routing changes from what is expected or visible to support tools it becomes an operational trouble spot for users and platform support to locate and determine the source of the problems.

2.2.9.2. CDN Edge Cache Selection due to Routing

A significant and often overlooked problem is the addition of network latency compared to edge CDN caches or access network peering connections. Routing changes which cause bypassing edge CDN caches and instead choosing less optimal caches

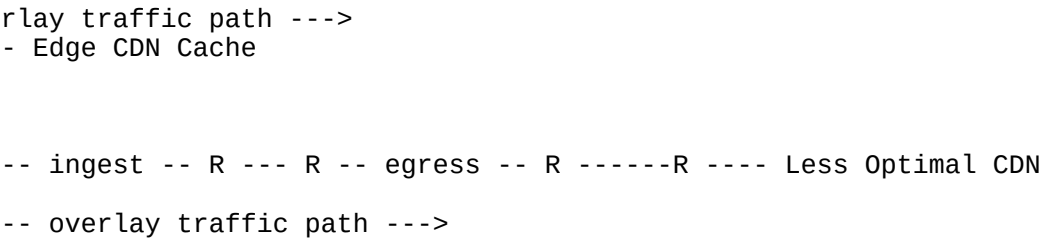

Figure: Routing Changes altering CDN Cache selection

2.2.9.3. Performance and Problem determination

Network overlays often interfere with the tools used in performance and problem determination. This is due to either the tool and protocols not being able to traverse the alternative route tunnel impacting service's ability to diagnose connection and performance problems, or the network overlay itself not supporting the tool and not supporting or carrying the tools functions.

2.2.9.4. Impact of Changing Network Routing and other Policies

The problem for streaming applications occurs when the underlying network properties and policies change from what is expected by the streaming application, especially when such changes are either hidden or not visible to the streaming application.

While the open Internet is a dynamic environment, changing of basic network behavior and policies from what is expected as seen from the streaming application, deviates unexpectedly from what the streaming application expects. This behavior disrupts the optimized streaming

delivery architecture for the end-user device. Changes to Network Policies such as routing, source IP address assigned to the streaming application traffic, DNS resolver choice etc. influences this behavior.

Having a reliable understanding of the delivery path is essential for streaming operators and the introduction of network overlays like those based on technologies such as MASQUE especially when designed to be undetectable by the applications using them has introduced new technical challenges for streaming operators and network operators as well as for their viewers.

The core problem occurs when changes to network policies are made often without notification or visibility to applications and without clear methods of probing to determine and test changed behaviors that affect the streaming application’s content delivery path resulting in increased latency, changes of IP address for the application as seen by either the application or the streaming service connection, changes to DNS resolvers being queried and the results returned by DNS, and changes to application transports such as adding or removing outer layer encryption are all problems that have been observed in production streaming platforms.

2.2.10. Unintended Content Blocking

A strongly undesirable unintended side-effect of network policy changes is the blocking of content to the viewer. This may be the primary content URLs access which are blocked, or possibly advertising fetched from a second URL from the main video content. This can be due to policy changes altering device IP addresses, or changes to routing that run afoul of enforced traffic routing policies.

Such blocking may be connected to restrictions built upon data feeds used for geofiltering and georestrictions, for example restriction which block delivery to networks identified as either commercial data centers or other CDNs service network addresses. Essentially, running afoul of configurations possibly used to combat security threats that expect streaming viewers to be on home or possibly mobile networks, but not in commercial data centers or CDN content networks and so block delivery to IP addresses in those unexpected network blocks. This is more likely to occur in network overlays that shift egress traffic to commerical or CDN blocks.

This is a particularly troublesome problem to determine as it may appear inconsistently from one streaming session to another. Small changes in URLs in manifests from on session to another, especially on streaming platforms that make use of multi-CDN delivery and may encounter different delivery and security protection policies from the different multi-CDN operators involed.

3. Policy Changes Hidden from Applications

One of the central recurring issues with streaming applications running on devices or networks with changed policies due to network overlays is that the changes are often hidden from the applications.

Applications often find it difficult or even impossible to detect when network policy changes will be active and what they are changing. For example, a device may have a designated default DNS resolver for the device but may have a different resolver selected depending on how the streaming application queries the DNS.

Likewise, a streaming application might find that one application transport protocol such as HTTP queries will have one set of routing policies applied to it but a different application transport like HTTPS may have a different set of routing policies applied.

Streaming applications that cannot determine the exact behavior to be expected can prevent the streaming application from making good content source decisions and can prevent applications from being able to provide reliable feedback and logs when problems are encountered.

4. Making It Easy (for Users) by Working Under the Covers

Historically, incorporating privacy features into consumer-facing products has been complex. This challenge arises from the need to address a wide range of use cases while also offering users easy access to advanced privacy frameworks and taxonomies. Many attempts have been made and very few have achieved finding success with end users.

Perhaps learning from the lessons of offering too many options, the recent trend in privacy enhancements has steered toward either a very simple "Privacy On or Off" switch or in other cases automatically enabling or "upgrading" to enhance privacy. Apple’s iCloud Private Relay can be easily turned on with a single settings switch, while privacy features such as Encrypted DNS over HTTP and upgrade from HTTP to HTTPS connections have had a several deployments that automatically enable them for users when possible.

Keeping with the motto of "Keep It Simple", users are generally not provided with granular Network Overlay controls permitting the user to select what applications, or what network connections the Network Overlay policies can apply to.

Adhering to the "Keep It Simple" approach the application itself has very little connection to privacy enhancing Network Overlays. Applications generally do not have a means to detect when networking policy changes are active. Applications generally do not have a means to access policy change settings or to interact to change them.

5. Streaming Video

Streaming Video, while just one of the many different Internet applications does standout from other uses in several significant ways that perhaps merit some amount of special consideration in understanding and addressing the impacts caused by particular privacy enhancing design and service offering choices.

Firstly, Streaming video operates at a hard to imagine scale - streaming video is served globally to more than 2 billion users daily and continues to grow in leaps and bounds.

Secondly, the content types delivered through streaming has evolved from the pre-recorded low-resolution, low-bit rate, latency tolerant video-on-demand movies, live or pre-recorded TV shows, and user generated videos delivered by pioneering streaming platforms to now including low-latency 4K and 8K live sports events, while also evolving the pre-recorded content with high-bit rate such as 4K and 8K cinema quality and High Dynamic Range (HDR) lighting.

Finally, the expectations of streaming video viewers have significantly evolved from the days of settling for being able to watch a movie in a PC browser. Viewers expect to watch on any device type they want ranging from low-end-streaming sticks that plug into a USB port, to 4K and HDR capable laptops, 4K and 8K HDR TV screens, gaming consoles, smart phones and many more choices. Viewers also expect to have the same great viewing experience while at home connected via high-speed wired Internet, high-speed WiFi, or mobile cellular 5G and even satellite Internet connections.

To meet the growth to billions of users, the growth in content type, quality and speed expectations, and on-any-device anywhere that I am over any-network-connection expectations of users the Streaming Video technology infrastructure has had to itself evolve significantly. This video streaming evolution work is being done in the IETF and in the Streaming Video Technology Alliance (SVTA) (https://www.svta.org/), and in a number of other technical and industry groups.

It’s hard to overstate just how much the growth of streaming video has contributed to the growth of the Internet. Internet connections of multiples of hundreds of megabits and gigabits speeds today are because of the needs of video streaming, the ongoing work on low-latency networking and ultra-low-latency video delivery are both driven by the use of streaming video.

5.1. Advances in Streaming Video Architecture

Internet streaming has greatly matured and diversified from its early days of viewers watching pre-recorded 320x240, 640x480 standard definition 480p movies to wired PCs connected to the Internet via high-latency, low-bandwidth DSL as early DOCSIS modems.

Streaming has grown to the extent that it has become a daily go-to video source worldwide for billions of viewers and has expanded from pre-recorded movies to encompass every type of video content imaginable. This growth to billions of viewers and the addition of low latency sensitive content and new connectivity options like WiFi, Cellular and Satellite in addition to high-speed DOCSIS and fiber is the world streaming platforms now provide service in.

With the large user base and its usage, the Streaming platforms also have significant technical challenges to meet viewer expectations:

* (1) Delivery scales that commonly range from hundreds of thousands to many millions of viewers simultaneously, with billions of daily global views.

* (2) Low latency demands from live sports, live events and live streamed content.

* (3) Content resolutions and corresponding formats which have jumped from the days of SD-480p to 4K (3840x2160) and 8K (7680x4320) along with bit rates which can had data needs of 10-24+ Mbps for 4K with 8K demanding 40 Mbps under extreme compression and 150-300 Mbps for high quality such as cinema.

* (4) Devices with very diverse capabilities low-cost streaming sticks, to Smart TVs, tablets, phones, and game consoles.

* (5) Broad range of connectivity choices including WiFi, Gig speed-low latency DOCSIS, Fiber, satellite, and 5G cellular networks.

* (6) Application transport protocols including MPEG DASH, HLS, HTTP2/TCP, HTTP3/QUIC, WebRTC, Media over QUIC (MoQ) and specialty application transports such as SRT, HESP etc.

To meet these challenges streaming platforms have significantly invested in developing delivery architectures that are built with detailed understandings of each element in the content delivery pathway, starting from the content capture all the way through to the screen of the viewer.

Streaming applications are part of an end-to-end architecture that is optimized around achieving the best experience including low latency video delivery to viewing devices. The open Internet can be unpredictable with temporary issues like packet loss, congestion and other conditions. However, streaming architecture is designed to handle these momentary problems as effectively as possible often through use of dynamic adaptive approaches designed into streaming protocols and platform components.

6. Middleboxes and learning from the past

The IETF has discussed this situation in the past, more than 20 years ago in 2002 Middleboxes: Taxonomy and Issues [RFC3234] was published capturing the issues with Middleboxes in the network and the effects of hidden changes occurring on the network between the sender and receiver.

7. Appendix A: Network Overlays are different than VPNs

While conceptually similar in many ways to VPN (Virtual Private Network) technology, the various network overlay technologies currently being deployed as well as new ones currently being designed by the IETF differ quite significantly from the older VPN approach they are replacing in a number of ways.

It is also worth noting that one reason why the issues discussed in this document have not been concern with regard to VPNs is that largely VPNs have not been a pervasive way to stream video. First, many VPNs have not had very good or consistent throughput compared to the direct open Internet and so provide a poor viewing experience. Second, many video platforms block or deny service to VPN connections due to the very common use of VPNs to bypass geofiltering restrictions.

Whatever the reason, it is worth looking at how VPNs differ from the Network Overlays being discussed herein.

7.1. VPNs typically:

* (1) VPNs typically are detectable by both the video application and often by the streaming platform.

* (2) VPNs typically work at the network layer of a device, resulting in a wide-range (if not all) transports and protocols from the device flowing through the VPN

* (3) VPNs typically provide exception options allowing for exclusion from traversing via the VPN based on various criteria such as application, destination IP address, application protocol etc.

7.1.1. Network Overlays typically:

* (1) Network Overlays are often undetectable by video applications or by the streaming platform, when in use.

* (2) Network Overlays often only apply to specific application transports such as HTTP2/TCP or HTTP3/QUIC while not applying to HTTP2/TCP+TLS on the same device.

* (3) Network Overlays often only apply to HTTP connections and do not support ICMP, non-HTTP versions of DNS, NTP etc., and various tools used for network measurement, problem determination, and network management that are not HTTP based.

* (4) Network Overlays do not expose to applications any means for the application to discover the policy changes the overlay will apply to the applications network connections.

* (5) Network Overlays do not expose mechanisms or APIs for applications to interact with them such as getting or setting options.

8. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

9. Security Considerations

TODO Security

10. IANA Considerations

This document has no IANA actions.

11. 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/rfc/rfc2119>.

[RFC3234] Carpenter, B. and S. Brim, "Middleboxes: Taxonomy and Issues", RFC 3234, DOI 10.17487/RFC3234, February 2002, <https://www.rfc-editor.org/rfc/rfc3234>.

[RFC6891] Damas, J., Graff, M., and P. Vixie, "Extension Mechanisms for DNS (EDNS(0))", STD 75, RFC 6891, DOI 10.17487/RFC6891, April 2013, <https://www.rfc-editor.org/rfc/rfc6891>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <https://www.rfc-editor.org/rfc/rfc7258>.

[RFC7264] Zong, N., Jiang, X., Even, R., and Y. Zhang, "An Extension to the REsource LOcation And Discovery (RELOAD) Protocol to Support Relay Peer Routing", RFC 7264, DOI 10.17487/RFC7264, June 2014, <https://www.rfc-editor.org/rfc/rfc7264>.

[RFC7624] Barnes, R., Schneier, B., Jennings, C., Hardie, T., Trammell, B., Huitema, C., and D. Borkmann, "Confidentiality in the Face of Pervasive Surveillance: A Threat Model and Problem Statement", RFC 7624, DOI 10.17487/RFC7624, August 2015, <https://www.rfc-editor.org/rfc/rfc7624>.

[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/rfc/rfc8174>.

[RFC9000] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021, <https://www.rfc-editor.org/rfc/rfc9000>.

[RFC9317] Holland, J., Begen, A., and S. Dawkins, "Operational Considerations for Streaming Media", RFC 9317, DOI 10.17487/RFC9317, October 2022, <https://www.rfc-editor.org/rfc/rfc9317>.

[RFC9484] Pauly, T., Ed., Schinazi, D., Chernyakhovsky, A., Khlewind, M., and M. Westerlund, "Proxying IP in HTTP", RFC 9484, DOI 10.17487/RFC9484, October 2023, <https://www.rfc-editor.org/rfc/rfc9484>.

Acknowledgments

The authors would like to acknowledge the contributions from the Streaming Video Technology Alliance (SVTA) based on their work studying the impacts of network overlays on the streaming platforms. The contributions from Brian Paxton on observed overlay behavior and comments from Jay Robertson have been very helpful.

Authors' Addresses

Glenn Deen
Comcast-NBCUniversal
Email: glenn_deen@comcast.com

Sanjay Mishra
Verizon
Email: sanjay.mishra@verizon.com