Delivering the Highest IP Video Quality Efficiently While Improving Customer Experience

A Technical Paper prepared for SCTE•ISBE by

Garey Hassler
Distinguished Engineer Software Architect
Comcast
1515 Wynkoop St. Ste. 200, Denver, CO 80202
720-502-3717
Garey_Hassler@comcast.com

Table of Contents

List of Figures

List of Tables

Introduction

Television audiences expect the best video quality their devices can render and to have it available for viewing in seconds. Years of video quality branding, like high-definition (HD) and 4K, and other consumer educational efforts, have fashioned discernible viewers capable of distinguishing sharpness, brightness and resolution and equating this to quality. Furthermore, with decades of broadcast consumption, viewers have been observing diminishing tune times, or time-to-first-frame (TTFF), creating an expectation, which can be challenging to achieve in a redundant IP video service.

Multichannel video program distributors (MVPDs) can dramatically improve their ability to achieve viewers' expectations by enhancing a video origin to support Hypertext Transfer Protocol Version 2.0 (HTTP/2) [1], synchronizing video transcoding and packaging, applying video quality measuring, and automating origin selection.

A proof-of-concept (POC) was constructed to independently measure the video quality of a single video source, relative to the synchronized transcoding and packaging of two geographically dispersed sites. Applying automated decision logic on the quality measurements to select an origin for distribution on a per-fragment interval ensures the highest quality is distributed to viewers. Utilizing (HTTP) version 2.0 PUSH method to immediately distribute the IP content through a content delivery network (CDN), reduces network traversal from a round-trip-time to an end-to-end latency.

Evaluations of the POC were conducted in simultaneous test executions compared with existing client/server production components, in a non-isolated network, thus any network congestion or impairments applied equally to the test measurements. Preliminary test results over multiple executions indicate a 45% – 56% improvement in TTFF, while delivering the highest quality video fragment to all test players.

IP Video Service Resilience

1. Redundant and Resilient Linear IP Video Service

Delivering a linear IP video service redundantly and resiliently is a challenge for MVPDs. With the introduction of duplicate versions of a channel and quality monitoring tools, new latencies can impact the service. The following subsections look at some of the drivers leading to the proof-of-concepts covered in this document.

1.1. Synchronized Transcoding and Packaging

With the development of coordinated processing by transcode vendors, it's possible to configure a pair of transcoders into a master/slave relationship, allowing a MVPD to improve the reliability and resilience of an IP video service. When transcoding from two origins which are synchronized and operating at two independent datacenters, devices are capable of seamlessly retrieving and playing the video content distributed from either video origin. This feature significantly reduces the likelihood a customer will experience the dreaded buffering symbol when there is a network impairment or hardware disruption. However, this introduces new complexities and questions for an MVPD to consider:

- What methodology should the distributor employ when operating the origin: even distribution across channel lineup, load balancing by viewer, primary versus secondary model, or some other approach entirely?
- How does a distributor ensure customers receive the best available video quality when there is a video provider source or network issue occurring at only one origin?
- How does a distributor prevent the introduction of additional latencies into the delivery of the video?

Serving the broadest set of devices is every MVPD's desire. To facility this objective, a standard from SCTE affectionately called Common Intermediate Format [2] is the output produced by the packaging system. The media files utilized are a transport stream and this format is maintained until processing occurs at a Just-In-Time-Packager, which transforms the manifest and media files into a format requested by a client.

1.2. Hypertext Transfer Protocol Version 2.0 (HTTP/2)

The advanced features provided by the HTTP/2 standard afford some significant benefits for the distribution of video content. The standard defines an upgrade mechanism for client devices that support both HTTP/1.1 and HTTP/2 to propose use and support of these advances in requests it initiates. Furthermore, the additional features allow compatible servers and clients to support a new feature called Server Push, which essentially allows a server to preemptively deliver associated files to the client based upon the previously initiated request. Common examples of the types of additional files might be images or Cascading Style Sheets (CSS) that are associated with the content requested by the client.

By expanding the interpretation of a request for a linear video channel from a single request for a fragment[1] of some common duration into a broader definition to one of requesting to consume a channel, it becomes possible to theorize and conceptualize a single request for all future variants of a video manifest and media fragments. In its simplest form of a client and server, shown below, a client would initiate a request to "tune" to a channel along with an indication it supports upgrading to the HTTP/2 standard. The server will respond by accepting to upgrade, notifying the client of a future content via a Server Push using the PUSH_PROMISE of a new stream to be initiated by the Server, and then deliver the requested video manifest file.

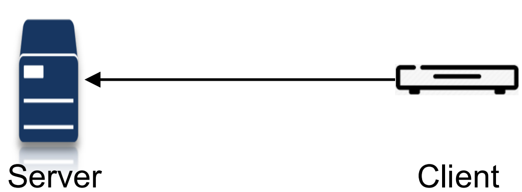

Figure 1 – Simple HTTP Service

[1] A small discrete video file typically of a common size. In this POC, fragment durations are 2 seconds in length.

In this configuration, a Linear IP Video Origin Server can repeatedly initiate a new stream with the client, send a new PUSH_PROMISE for a subsequent stream and deliver the most recent updated changes to the manifest or media file as they are produced on its normal cadence of processing. The server is capable of detecting or inferring when the client no longer needs the media through the simple termination of the Transmission Control Protocol (TCP) communication link. Alternatively, the server can force a reestablishment of the process at any time by closing the TCP connection with the client.

1.3. Content Delivery Network

For an MVPD providing thousands of streams servicing millions or even tens of millions of video customers, the network is far more complicated than the one shown in Figure 1 – Simple HTTP Service. They will frequently include multiple layers of a Content Delivery Network (CDN) to support the caching that reduces latencies and improves resiliency when delivering video over large areas. The diagram below illustrates an example of a mid-level or mid-tier of a CDN cache and an edge cache, the latter often being in close proximity to the consumer or client device.

Figure 2 – HTTP Service with CDN

In a typical video provider network there will be multiple mid-tier caches and a greater number of edge caches distributed throughout the service providers network to serve the millions of client devices. Delivery of the video across these growing number of network hops is a common network form or pattern called a “fan-out distribution.”

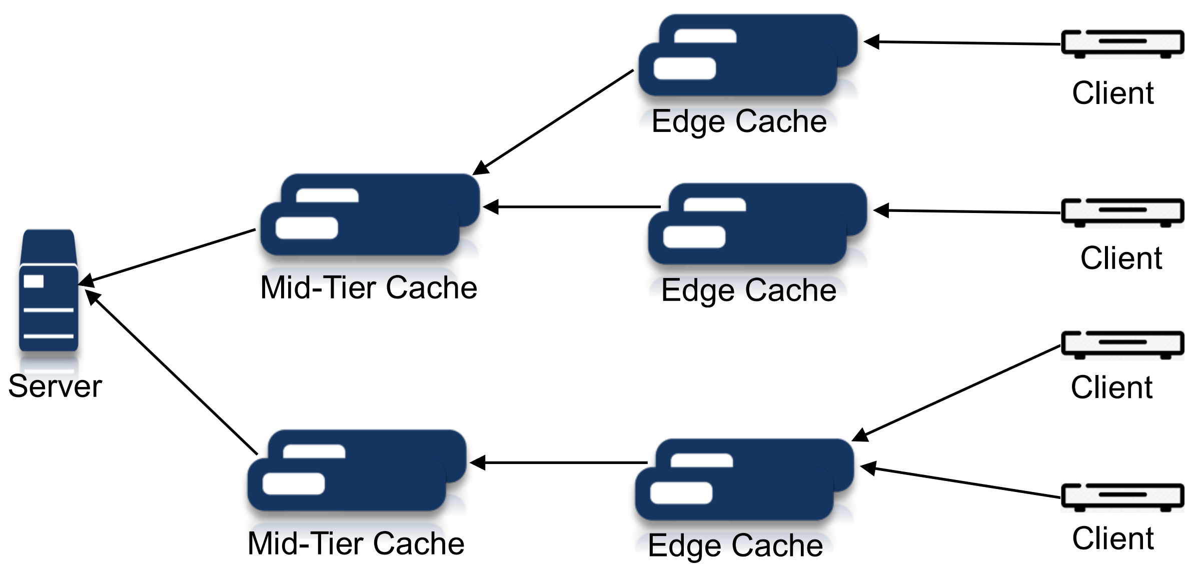

Figure 3 – HTTP Fan-Out Example

The introduction of a CDN results in the need to propagate the HTTP/2 through one or more intermediary systems, disrupting the server and clients' ability to infer when either system no longer needs a video channel. With the introduction of a Session Management module into the CDN service, the "proxying" of the TCP communication link closures and terminations can be propagated through the entire network. The Session Management processing is responsible for maintaining a mapping or association of upstream links (communication links towards the server) and with communication on its downstream links (those communication links that originate from the client side of the diagram.) Once all of the downstream links in a component for a channel terminate, the Session Management closes its upstream link. With this concept, propagation of the inference of channel viewing can be conveyed through the video distribution network without any new messaging beyond standard TCP messages.

1.4. Video Quality and Automated Origin Process

As mentioned previously, a video distributor wants to be able to provide the best available video quality equally to all of its subscribers. To achieve this goal, the distributor monitors the video quality it is receiving from a source provider as well as the quality being produced by the transforms conducted by transcoding and packaging the media. When augmenting this distribution with redundant transcode and packaging, there are greater opportunities for the quality to vary between those two independent video origins. The MVPD can use a variety of vendors or open source tools to perform the quality measurement, but how does it go about ensuring all viewers are receiving the best quality at a particular moment in time?

The introduction of a new component dubbed in this document as the Video Quality Agent (VQA) can apply processing and decision logic based upon the scores of each transcode and packaging origin pair of a channel. Based upon some threshold of deviation in quality between the two origins, the VQA controls which origin is the source of content for all viewers, through notification to one origin to serve the content, and the other or secondary origin to redirect content to the origin currently serving the channel. This control can consider multiple factors in its decision processing and can be as granular as a single fragment of media.

In the situation where a distributor wants to operate in an even distribution of requests to two redundant origins, such as a simple round-robin approach from Client devices, the redirecting of traffic to a single origin introduces additional latencies and increased network load for some clients. When combined with the HTTP/2 Server Push functionality and the expanded Session Management functionality with the two new features, this mechanism can propagate the highest quality in an Automated Origin process. To facilitate the indication to downstream systems that a media file is duplicative from multiple origins, the active Origin Server includes a new custom HTTP header, such as X-Video-Origins, indicating it is the authoritative origin for both redundant servers. The new HTTP header contains the host information of the peer Origin Server of the paired servers and is provided by the Video Quality Agent or through configuration. When the media file is "pushed" or published to the downstream systems connected in the fan-out, the Session Manager of the system creates a new link or index mapping that references the cached file as an item to serve for any request for either origin, thus a request for the secondary origin can be served with the file sent by the primary origin. In a similar manner, the cache can publish to the next downstream systems that have upgraded to use the HTTP/2 mechanism for delivery. This pattern is repeatable in each system's Session Management processing.

2. Proof-of-Concepts

2.1. HTTP/2 and Session Management

Evaluating the effectiveness of this conceptual set of hypotheses was conducted using software derived from the preexisting production software systems, emulating the HTTP/2 and Session Management functionality. The diagram below illustrates the test environment created to conduct the POC experiments.

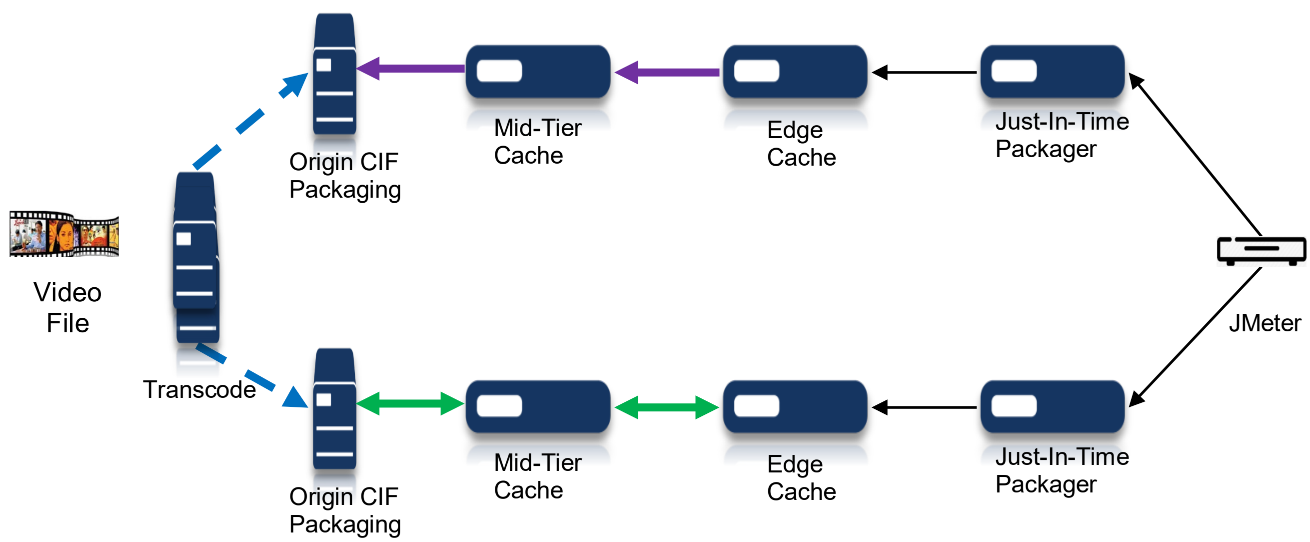

Figure 4 – HTTP/2 and Session Management Test Environment

Beginning from the left-hand side of the diagram, an HD mezzanine quality video file is the input into a transcoder, which is processed into 5 MPEG-4 audio and video profiles of two seconds in duration with an average bandwidth of 3.5 Mbps for the highest resolution, 2.0 Mbps for the middle resolution, and 0.5 Mbps for the lowest resolution profile. The output of the transcoder is a multicast UDP packets sent to both Origin CIF Packaging servers, as indicated by the dotted blue lines. The upper Origin serves the media files using the HTTP/1.1 client/server methodology over the systems, indicated by the purple lines, and the lower Origin delivers the media files to the Mid-Tier and Edge-Cache, depicted by the green lines, with arrows on both ends, in reference to communications beginning by a client and then propagating downstream, using HTTP/2 and employing the Session Management functionality. The POC limits the use of HTTP/2 to the Origin, Mid-Tier and Edge Caches for testing for two reasons: 1) to limit the software updates prior to obtaining demonstrative evidence of actual improvements and 2) to validate that the optimization could occur, without the need for all systems to implement simultaneously, before benefits could be realized by the service and the MVPD.

In order to reduce the complexity of the Session Management algorithm, separate communication links were employed for the manifests from those of the fragment files. This simplified the code development and management of the synchronization, in cases when publishing multiple fragment files prior to publishing the manifest to the next downstream systems.

To minimize external variables, such as network latencies skewing the measurements across each test run or between the two paths, the systems were all deployed in the same datacenter, on virtual machines with the same resource configurations and similar network traversal paths. To further mitigate any outside bandwidth impairments, data collection occurred simultaneously over both methods. Each test execution was conducted by a single JMeter instance, operating within two mutually exclusive threads with the same script to allow for common clock and time collection as they emulated a client's behavior.

Test results of 50 executions of the environment, each 5 minutes in duration on the three different video resolution levels are illustrated below. Table 1 – Average Range of 50 Test Runs of 5 Minute Duration shows the summary of the span of average times over all of the test runs. The test collected the complete time necessary to request and receive a complete copy of the manifest and the next two-second fragment, beginning with a simulated "tune" to the channel. In addition, the average improvement of each test run was calculated with the range of results shown in the final column.

Table 1 – Average Range of 50 Test Runs of 5 Minute Duration

Test Case Complexity	Avg. HTTP/1.1 Delivery Range	Avg. HTTP/2.2 Delivery Range	Improvement Range
Highest 3.5 Mbps Profile	13.316393 – 17.230584 (ms)	7.595753 – 10.039682 (ms)	45% - 56%
Middle 2.0 Mbps Profile	10.332659 – 14.64361 (ms)	7.121569 – 8.585287 (ms)	34% - 41%
Lowest 0.5 Mbps Profile	9.518596 – 11.713008 (ms)	6.899321 – 8.29262 (ms)	28% - 35%

In most circumstances, when a viewer tunes to a channel, another viewer or service, like a cloud recording service, will have already requested the media files. This is valid for all except for the very first system needing a file. A technique was defined (though not evaluated in the POC) to improve a client's time-to-first-frame by enhancing the operations of the Origin Server, upon any initial communication link, to request a manifest from a downstream system. This request indicates the support of the upgrade process. Furthermore, it can conceivably be considered to be attempting to consume the corresponding media. The concept is to automatically, preemptively transmit a predefined number of the most recently created media files produced by the server to the requesting system, in anticipation of subsequent requests for the associated media files, upon a client's processing of the manifest file returned in the initial request. Even with modern browser improvements that initiate multiple concurrent requests to fill their video buffers, so as to expedite the retrieval of multiple fragments, the reduced latency and playback initiation occurred in a manner that was perceptively sooner when using the HTTP/2 operations.

2.1.1. Dispelling the Single Viewer Misconception

Upon initial reflection, most people are under a misconception the improvements demonstrated by the POC don't justify the additional overhead expense of the new processing, because only the first client actually retrieves the files from the origin. However, this does not take into consideration all of the devices actually waiting upon the cache to obtain a copy of the content to service subsequent requests. Most CDNs and caching software today will actually hold or queue up multiple requests received while an outstanding request for the same file is actively being requested from the origin. The feature has different names in CDNs. As examples, Apache Traffic Server calls it "reader-while-writer" and Varnish calls it 'coalesced'. Regardless of the name, all requests received by the CDN, while awaiting the file from the upstream system, experience some latency and delay. In a widely distributed network with a large viewership, it is quite common for tens of thousands of requests to experience some level of queueing in various systems, as the content traverses the fan-out. This is alleviated with the HTTP/2 method, as the manifest is the final file delivered. And, because the Origin initiates the operation, it is only bound by its' computational ability and network transfer capacities to deliver to the next downstream systems.

2.2. Video Quality

A discrete POC was created to evaluate the implications of the Video Quality and Automated Origin Process when employing the delivery of content using the HTTP/2 methodology. The Video Quality Agent (VQA) evaluates the quality of the transcoding and packaging systems, determining a quality score for the combined transforms. This scoring can be performed on a referenced basis, through comparison with the original video file or source input, or non-referenced basis, using no comparison to the input. When the quality produced in one origin path deviates by some margin, perhaps 5%, the VQA directs an Origin Server change. The VQA coordinates and controls all HTTP/2 publishing by directing the Origin CIF Packaging Server with the best quality to enable publishing, and directs the lower quality or scoring Origin to disable publishing.

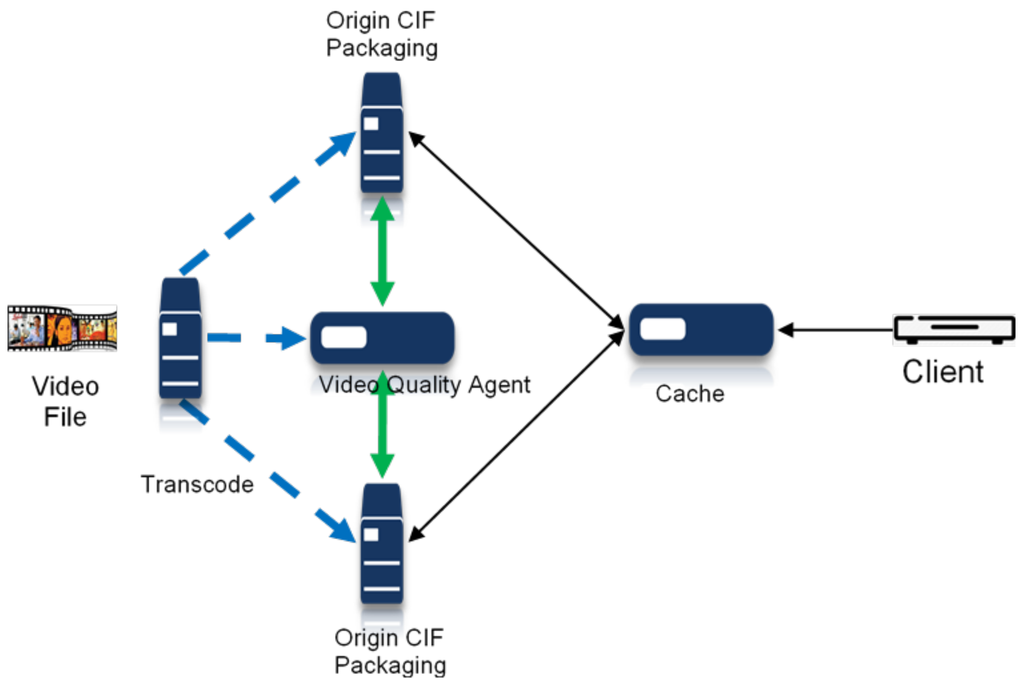

Figure 5 – Video Quality and Automated Origin Processing POC

The VQA is responsible for communicating the host information of the peer Origin CIF Packaging Server for inclusion in the X-Video-Origin. When the video quality varies beyond one of the thresholds, the VQA performs the control logic for the channel's origin switch. From the client's perspective, it continues to request content from the same origin and is completely unaware of the transition between origins.

Testing was conducted to disrupt the video by introducing packet loss which reduced the qualities created at one of the transcoders. The tests were conducted 20 times each, with varying the levels of impairment, in an attempt to cause a large enough issue or disruption to be visually impacting during playback. In all tests, the playback was uninterrupted and continued without any awareness or perception from the viewer.

Conclusion

The methods described in this document demonstrate the improvements possible for a Linear IP Video Service. Given the range of the test runs, it is safe to infer that the majority of the timing reduction in obtaining a media file from an origin is the result of the file traversing from the origin to a client without the requirement to be initiated by a preceding request. In other words, by transforming the delivery model from one of request-and-response to a client-initiated file delivery, delivery of files can be estimated to be equivalent to the end-to-end latency of the service, instead of the round-trip-time. The techniques outlined leveraging synchronized origins with video quality measurements and controls ensure all viewers receive the highest quality video available, with minimal additional complexity and no additional network communication.

Abbreviations

CDN	Content Delivery Network
CIF	Common Intermediate Format
CSS	Cascading Style Sheets
HD	High Definition
HTTP	Hypertext Transfer Protocol
ISBE	International Society of Broadband Experts
MVPD	Multichannel Video Programming Distributor
POC	Proof-of-Concept
SCTE	Society of Cable Telecommunications Engineers
TCP	Transmission Control Protocol
TTFF	Time-to-First-Frame
VQ	Video Quality
VQA	Video Quality Agent

Bibliography & References

[1] IETF RFC 7540 2015, *Hypertext Transfer Protocol Version 2* (HTTP/2)

[2] ANSI/SCTE 214-4 2018, *MPEG DASH for IP-Based Cable Services Part 4: SCTE Common Intermediate Format (CIF/TS) Manifest for ATS Streams*

