RADCOM

White paper

Insights into Encrypted Adaptive Video Streaming:

Challenges, Trends, and Solutions

Contents

Introduction

Video services continue to dominate global internet traffic, now accounting for approximately 80% of all mobile data traffic. At the same time, most of that video traffic is encrypted, and this trend is expected to continue. This presents a significant challenge for telecom operators striving to extract meaningful insights from their networks and ensure a high Quality of Experience (QoE) for subscribers.

Encrypted adaptive video streaming protocols, such as HTTPS and QUIC, render traditional traffic inspection methods ineffective, making it challenging for operators to monitor service performance and user experience. Yet, video streaming remains a vital area for growth and monetization, with opportunities to bundle and upsell services. To succeed, operators need visibility into encrypted traffic without violating privacy standards.

RADCOM ACE addresses this challenge by leveraging advanced Artificial Intelligence (AI), machine learning (ML), and heuristic models to deliver accurate, real-time insights into perceived Quality of Experience (QoE), even across fully encrypted video streams. By analyzing massive volumes of encrypted streaming data at scale, RADCOM generates key quality indicators (KQIs) such as:

- Video start delay
- Rebuffering events and durations
- Streaming resolution levels and duration per level
- Effective video throughput (in bytes)

These KQIs are normalized into a QoE score, ranging from 0 to 5, which offers operators a clear and consistent measure of user experience across the network.

This white paper outlines RADCOM's AI-driven methodology for assessing the quality of experience (QoE) of encrypted video. It details how encrypted data is processed, how metrics are derived, and how operators can use these insights to ensure service excellence for leading streaming platforms, including YouTube, Netflix, Amazon Prime Video, and Facebook.

Understanding adaptive video streaming

Streaming applications like YouTube utilize Dynamic Adaptive Streaming over HTTP (DASH)—also known as MPEG-DASH—to encrypt and deliver video content efficiently. DASH is a multi-bitrate (MBR) streaming method designed to maintain a high Quality of Experience (QoE), even under fluctuating network conditions. It works by breaking video content into small segments (typically a few seconds each), with each segment encoded at multiple bitrates and resolutions.

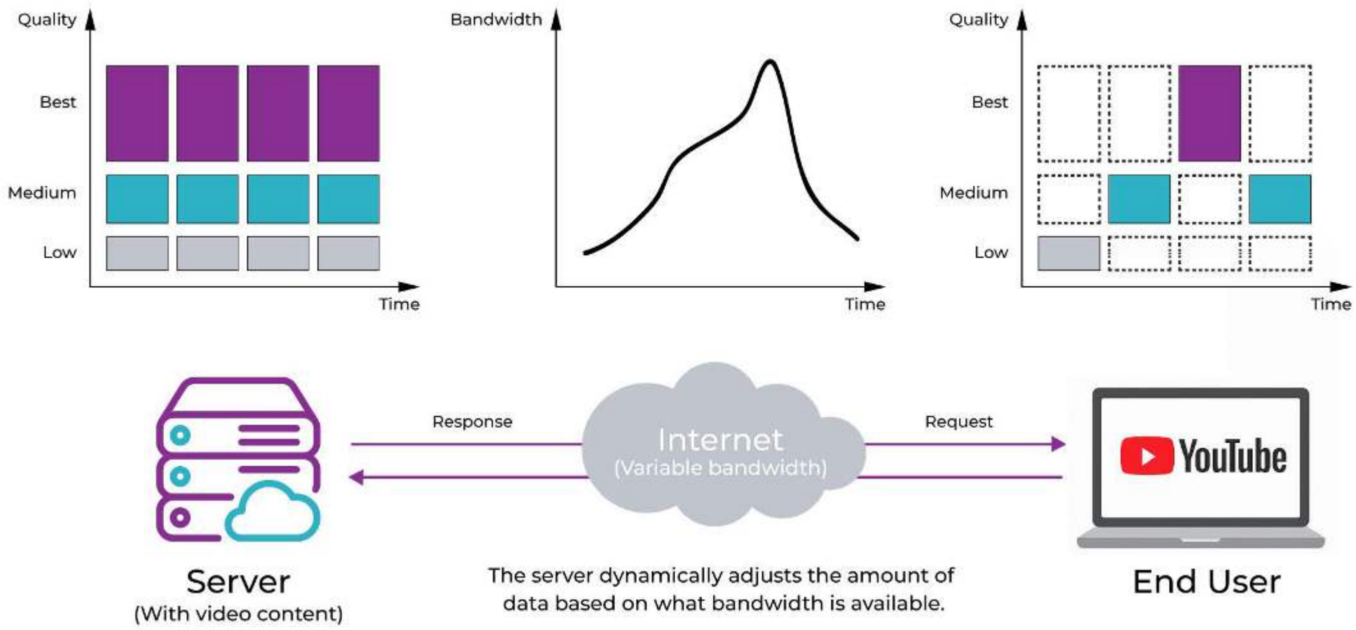

Figure 1: Dynamic Adaptive Streaming over HTTP (DASH)

During playback, the DASH client uses an adaptive bitrate (ABR) algorithm to dynamically select the highest-quality segment that current network conditions and buffer levels can support. This adaptive approach minimizes stalls and rebuffering, ensuring smooth and consistent video delivery.

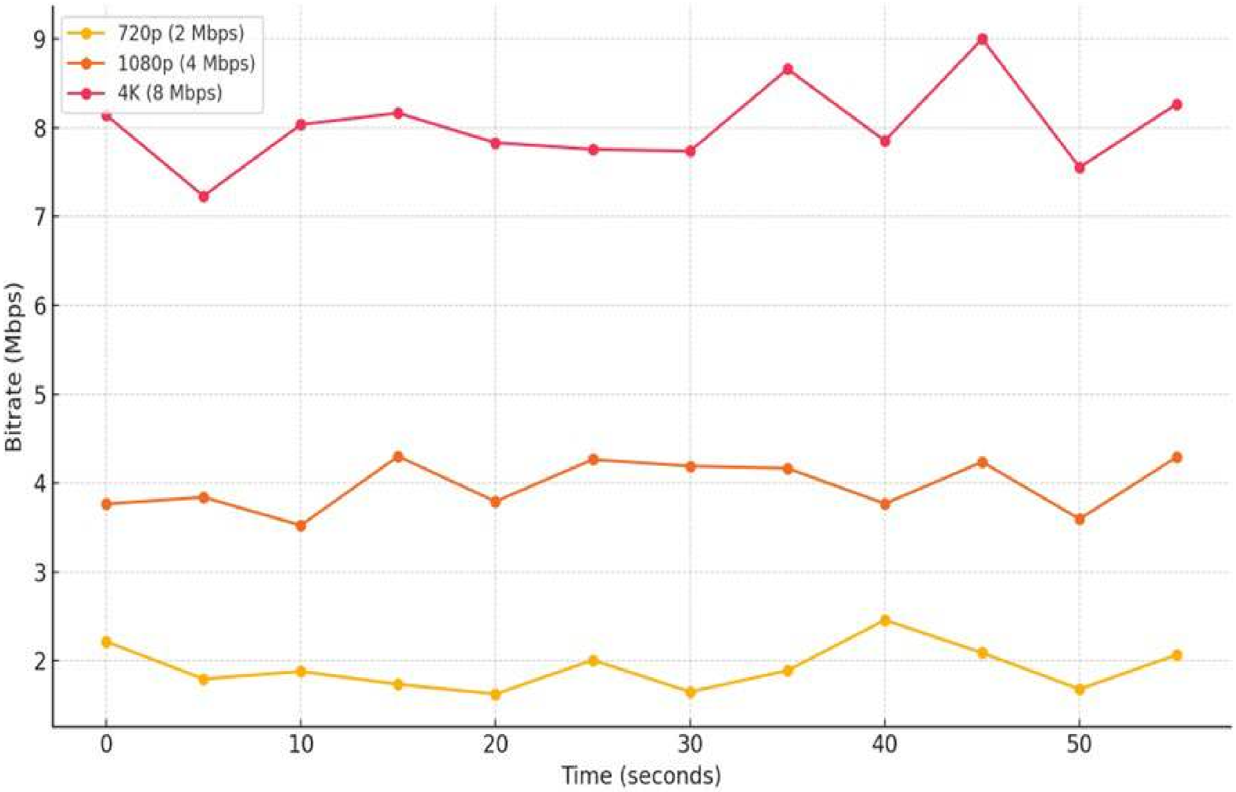

Figure 2: Wireshark analysis of video streaming

To accurately assess QoE for encrypted video streaming, it is essential for content classification algorithms to automatically detect the quality level of each video segment directly from the network traffic, despite the encryption. By classifying the video quality of encrypted streams, operators can regain visibility and estimate key video streaming KQIs such as initial playout delay, video resolution, and rebuffering frequency and duration.

Protocol	Description	Advantages	Disadvantages
MPEG-DASH	Open standard for adaptive streaming; codec-agnostic	▪ Open and standardized (ISO/IEC) ▪ Codec-agnostic (supports H.264, H.265, AV1, etc.) ▪ Widely supported in modern players and devices	▪ Not natively supported by Apple Safari ▪ Requires additional effort for DRM and compatibility
HLS (HTTP Live Streaming)	Developed by Apple; widely used for iOS and Safari	▪ Strong support on Apple devices ▪ Built-in encryption and DRM ▪ Simple to implement and widely adopted	▪ Higher latency ▪ Initially limited to MPEG-TS (less efficient than newer formats like fMP4)
HDS (HTTP Dynamic Streaming)	Adobe's streaming solution (Flash-based)	▪ Integrated well with Flash ecosystem ▪ Provided analytics and DRM via Adobe tools	▪ Flash is obsolete ▪ No longer supported in most browsers ▪ Not suitable for modern streaming
Smooth Streaming	Developed by Microsoft for Silverlight	▪ Good integration with Microsoft platforms ▪ Adaptive bitrate streaming supported	▪ Obsolete due to Silverlight deprecation ▪ Limited cross-platform/browser support
CMAF (Common Media Application Format)	A format rather than a protocol, used with HLS and DASH	▪ Reduces storage and encoding overhead by using a common format ▪ Enables low-latency streaming ▪ Compatible with both HLS and DASH	▪ Still requires use with a protocol (e.g., DASH or HLS) ▪ May require more complex setup

Table 1: Main types of dynamic adaptive streaming over HTTP and comparable protocols

MPEG-DASH is ideal for cross-platform, codec-flexible streaming, though it lacks full support from Apple. HLS remains the best choice for Apple-centric delivery, although it can experience higher latency. In contrast, HDS and Smooth Streaming are now considered obsolete. CMAF has emerged as a modern enhancement that bridges HLS and DASH, enabling more efficient and unified content delivery.

Challenges in understanding video quality of experience

Encrypted video streaming applications—like YouTube—often generate multiple parallel flows, each responsible for downloading different video segments. These flows may use protocols such as HTTPS, HTTP/2, QUIC, or a combination of them, which makes it increasingly difficult to monitor and accurately assess the video's Quality of Experience (QoE).

To address this, RADCOM has developed a heuristic-based approach that constructs a real-time, aggregated view of encrypted video traffic during the pre-processing stage. This process leverages RADCOM's Deep Packet Inspection (DPI) module alongside its flow aggregator components to generate a unified, structured time series representing video session behavior.

This intelligent aggregation is optimized explicitly for RADCOM's AI engine, enabling the system to generate accurate, real-time video quality metrics for each encrypted streaming session. These metrics serve as the foundation for assessing perceived Quality of Experience (QoE), even in complex, multi-flow, and encrypted environments.

The AI/ML learning process

RADCOM has developed an innovative 'video crawler' designed to capture and analyze network-level video streaming traffic, user experience feedback, and application-level data. This tool enables the collection of rich datasets that reflect real-world streaming behavior.

Using this data, RADCOM's AI engine is trained via supervised machine learning, where user experience feedback and application metrics serve as labeled data. By correlating these labels with patterns in the underlying network traffic, RADCOM's solution learns to accurately infer video Key Quality Indicators (KQIs)—such as buffering, resolution, and startup delay—based solely on network behavior, even in encrypted environments.

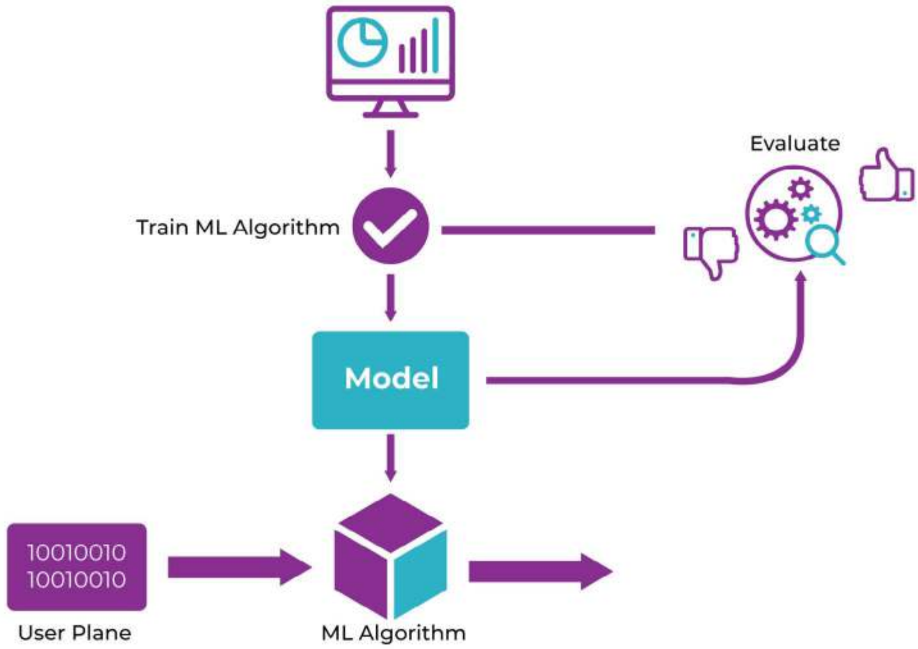

Figure 3: RADCOM's innovative ML learning process

The training process is conducted offline, leveraging a combination of real traffic from the operator's live network and controlled data from the RADCOM lab. The resulting model is then validated against historical traffic traces and test user data to ensure robustness and accuracy across a wide range of network conditions and usage scenarios.

Resulting video metrics for user experience analysis

RADCOM's video streaming quality metrics are entirely network-based, leveraging the characteristics of the underlying traffic to estimate user experience without requiring access to application-layer data. All metrics are derived directly from encrypted network traffic using advanced analytics and protocol-aware inspection.

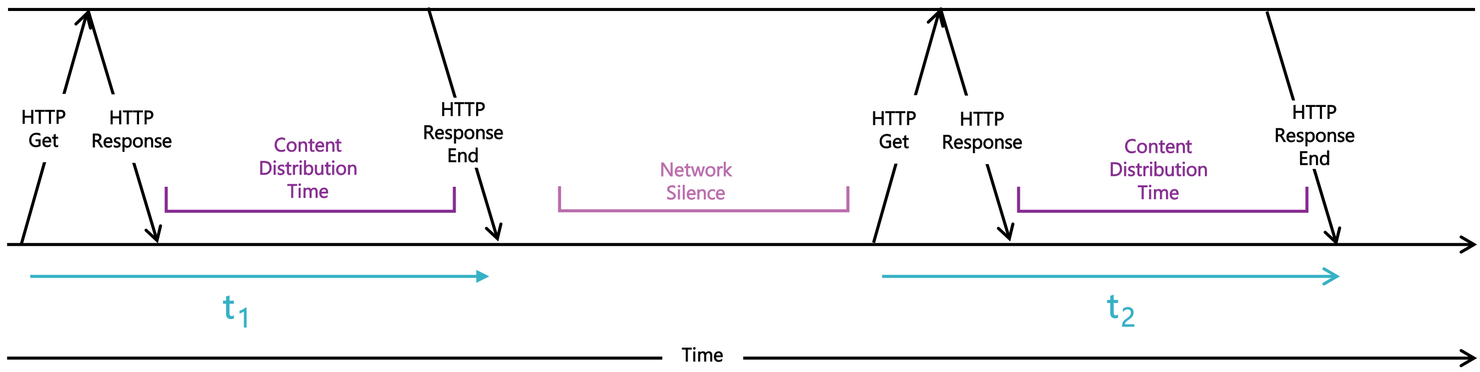

Figure 4: Video segment download pattern

RADCOM's technology identifies and analyzes individual video segments downloaded over the network. The download duration is calculated by detecting the start of the first video segment request and the arrival of the final packet in the corresponding video response. It is important to note that this download duration is typically shorter than the actual playback length of the video, as segments are downloaded ahead of time to maintain a buffer and ensure smooth playback.

Figure 5 illustrates a single encrypted video flow where traditional HTTP-level information is unavailable. In such cases, all relevant download activity—from the initial video request to the final response—is inferred and summarized using heuristic algorithms developed by RADCOM's research team.

Time t_1 represents the duration of the download for the first video segment. The interval between t_1 and t_2 indicates an application-driven pause, often referred to as an application timeout. During this period, the video player evaluates the current buffer and playback conditions to determine whether additional segments need to be requested. This helps optimize the selected video resolution and avoids unnecessary data downloads when buffer levels are sufficient.

At time t_2, the application resumes downloading by requesting a new video segment to refill the playback buffer. RADCOM's solution aggregates these parallel data flows in real time, using proprietary algorithms to construct a unified structure that represents multi-flow request-response patterns. This aggregated data forms the foundation for calculating video quality metrics (KQIs).

It is essential to note that the total video download duration typically concludes before the complete video playback ends, as segments are downloaded in advance. RADCOM's AI-based video quality metrics are calculated based on these valid, aggregated segment-level download durations.

In network conditions where congestion or high packet loss is present, the period of Network Silence (between t_1 and t_2) may shrink or disappear entirely. RADCOM's AI engine learns from these traffic patterns. It uses them to accurately detect re-buffering events as they occur, contributing to a more precise assessment of the viewer's Quality of Experience.

RADCOM provides a comprehensive set of metrics for encrypted adaptive video streaming. These metrics are available through multiple analytics dashboards, providing deep visibility into user experience, even across encrypted traffic.

The following metrics are supported:

1. Minimum Time to Play

   An estimated metric, based on a heuristic model, that measures the time elapsed from the video request to the start of playback. Due to variability in video codecs and encoding efficiency, this metric is derived through trained estimations, validated against thousands of video downloads and user feedback to ensure accuracy.

2. Video Re-buffering Count and Duration

   Leveraging AI-driven heuristic modeling, RADCOM identifies genuine re-buffering events using time series analysis of segment download behavior.

   a. Re-buffering count: Number of detected re-buffering events during a session.
   b. Re-buffering duration: Total time during which playback was stalled due to buffering.

3. Video Resolution Duration

   This metric captures the duration of the video played at each resolution level, measured at the segment level in milliseconds (ms). The resolution is inferred using time series modeling that incorporates user-specific download history for enhanced precision.

   Resolution levels:

   - Level 1: 144p–240p
   - Level 2: 360p–480p
   - Level 3: 720p–1080p
   - Level 4: 1440p–2160p+ (including 4320p)

   These values can be used to calculate the relative proportion of each resolution during the entire video playback session.

4. Video Duration and Byte Count

   Indicates the total video playback duration along with the cumulative number of downloaded bytes associated with that video session.

5. Effective Throughput

   Represents the download throughput for the video session, calculated by aggregating video segment flows and excluding idle or silent network periods to provide a true reflection of content delivery performance.

Here is a summary of the video streaming metrics:

#	CDR Field Name	Units	Data type	Description	Sample Data	Expected Values
1	Estimate_ minimal_ playout_time	mSec	Numeric (11,0)	Minimum elapsed time until video starts playing on the client	3,000	0-40,000 (up to 40 seconds)
2	Video_total_ duration	mSec	Numeric (11,0)	Effective video duration as observed by the client	30,000	0-3,600,000 (up to 1 hour)
3	Video_ rebuffer_ count		Numeric (11,0)	Total number of re-buffering occurrences	3	0-10
4	Video_ rebuffer_ duration	mSec	Numeric (11,0)	Sum of re-buffering duration for all re-buffering events	2,000	0-3,600,000 (up to 1 hour)
5	Video_ resolution_ duration_1	mSec	Numeric (11,0)	Duration of video in resolution 1	5,000	0-3,600,000 (up to 1 hour)
6	Video_ resolution_ duration_2	mSec	Numeric (11,0)	Duration of video in resolution 2	5,000	0-3,600,000 (up to 1 hour)
7	Video_ resolution_ duration_3	mSec	Numeric (11,0)	Duration of video in resolution 3	5,000	0-3,600,000 (up to 1 hour)
8	Video_ resolution_ duration_4	mSec	Numeric (11,0)	Duration of video in resolution 4	5,000	0-3,600,000 (up to 1 hour)
9	Mean_ effective_ throughput_ received	'Bytes/ sec	Numeric (11,4)	Total DL effective throughput	10,000	0-10,000,000 (up to 10MB)
10	Total DL video bytes	Bytes	Numeric (11,0)	Total user bytes for video not including SSL , TCP headers or retransmissions i.e. sum of DL bytes for all data windows	1,000,000	0-100,000,000 (up to 100MB)

Video streaming insights

RADCOM provides advanced video Key Quality Indicators (KQIs) and actionable insights across multiple operational dimensions, including APN, device types, geographic regions, and charging characteristics. These KQIs help operators monitor and optimize video streaming experiences at scale.

- Video vMOS (Video Mean Opinion Score)

This metric represents a weighted composite score based on key quality indicators such as:

- Time to Start
- Re-buffering Proportion
- Resolution Level

The weights and thresholds used in this calculation are fully configurable to align with operator-specific performance and experience goals.
While there is currently no universally accepted industry standard for video vMOS, RADCOM's methodology is tailored to reflect real user experience, offering granular visibility and supporting data-driven decision-making across the network.

vMOS component	vMOS Score
Minimum time to play score	2
Rebuffering Score	2
Resolution Score	1

vMOS Score (0-5) = Minimum time to play score + Rebuffering Score + Resolution Score

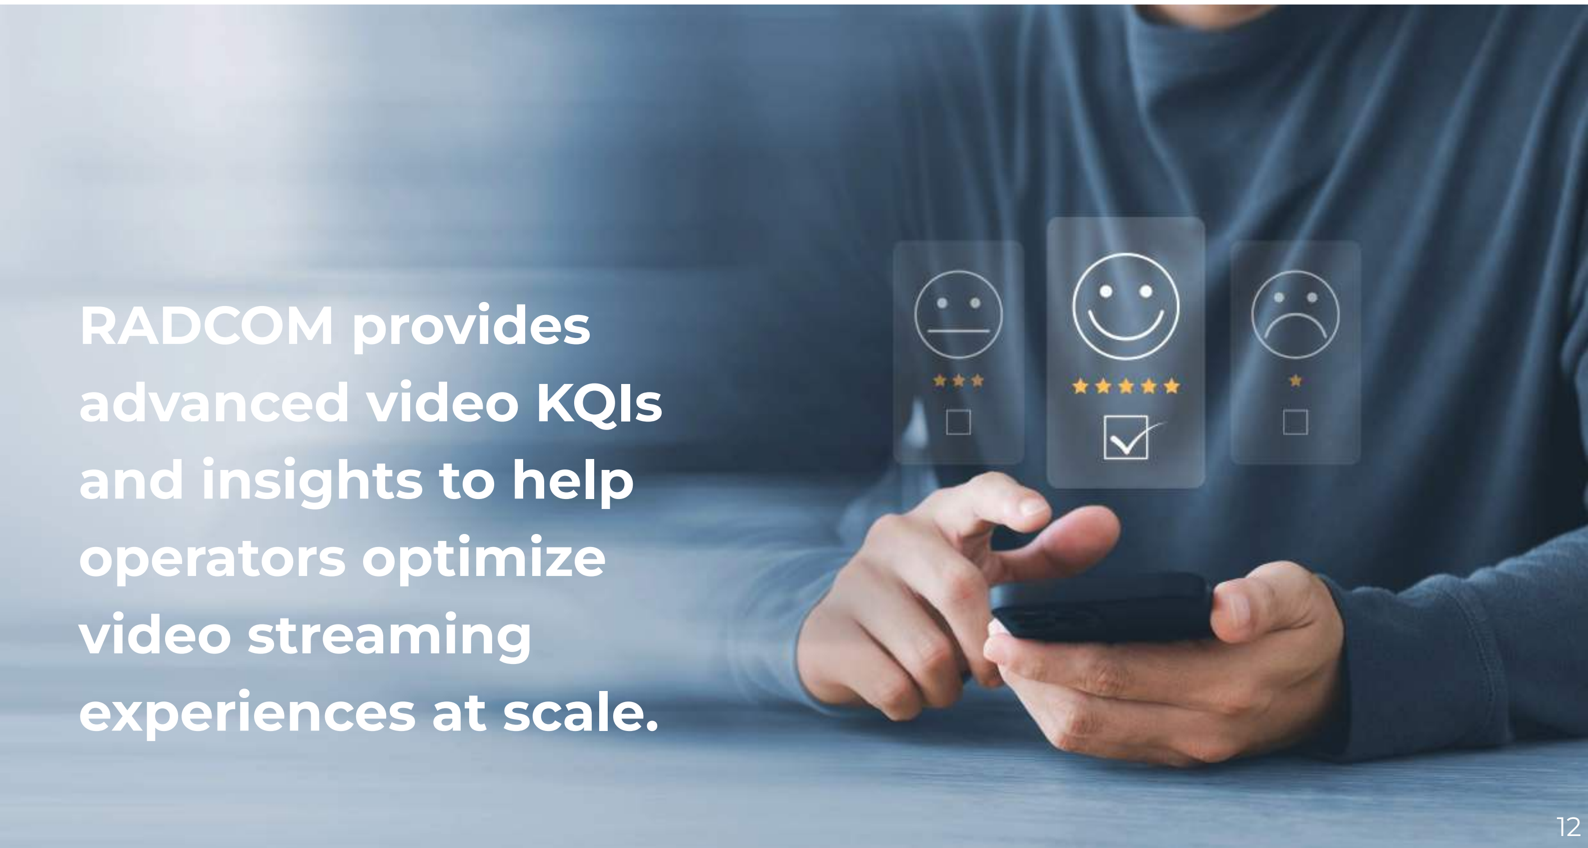

Minimum time to play score:

Playout score (2 = best, 0=worst)	**Minimum playout time (sec)**
2	less or equal to 1
1.7	greater than 1 and less or equal to 3
1.4	greater than 3 and less or equal to 6
1.2	greater than 6 and less or equal to 8
1	greater than 8 and less or equal to 10
0.8	greater than 10 and less or equal to 12
0.6	greater than 12 and less or equal to 14
0.4	greater than 14 and less or equal to 16
0.2	greater than 16 and less or equal to 18
0	greater than 18

Rebuffering score =

Rebuffering score (2 = best, 0=worst)	**Rebuffering Ratio = rebuffering duration / duration of video**
2.0	0
1.5	less or equal to 0.1
1	Greater than 0.1 and less than or equal to 0.2
0.8	Greater than 0.2 and less than or equal to 0.4
0.6	Greater than 0.4 and less than or equal to 0.6
0.3	Greater than 0.6 and less than or equal to 0.8
0	Greater than 0.8

Resolution score (0-1) = [(Video_resolution_duration_2 +2 × Video_resolution_duration_3 + 3X Video_resolution_duration_4) / (sum of video resolution duration 1-4) X 3]

Additional video streaming KQIs (filtered on dimensions such as APN, device, and region):

- Average Rebuffering Ratio: avg(rebuffering duration / duration of the video) where rebuffering ratio > 0
- Average Video Rebuffering Count: average rebuffering count for videos where the rebuffering count > 0
- Video Resolution: Ratio of video duration for each of the four resolution levels
- Average Minimal Playout Time
- Average Video Duration: Average video duration where vMOS > 0
- Total Video Data: Total video bytes for all videos where vMOS > 0
- Average Download Throughput: Average effective video throughput where vMOS > 0

Examples of video streaming insights

Figure 6 illustrates the relative distribution of video playback duration across four resolution levels, segmented by subscriber attributes such as APNs, device types, frequency bands, and geographic locations.

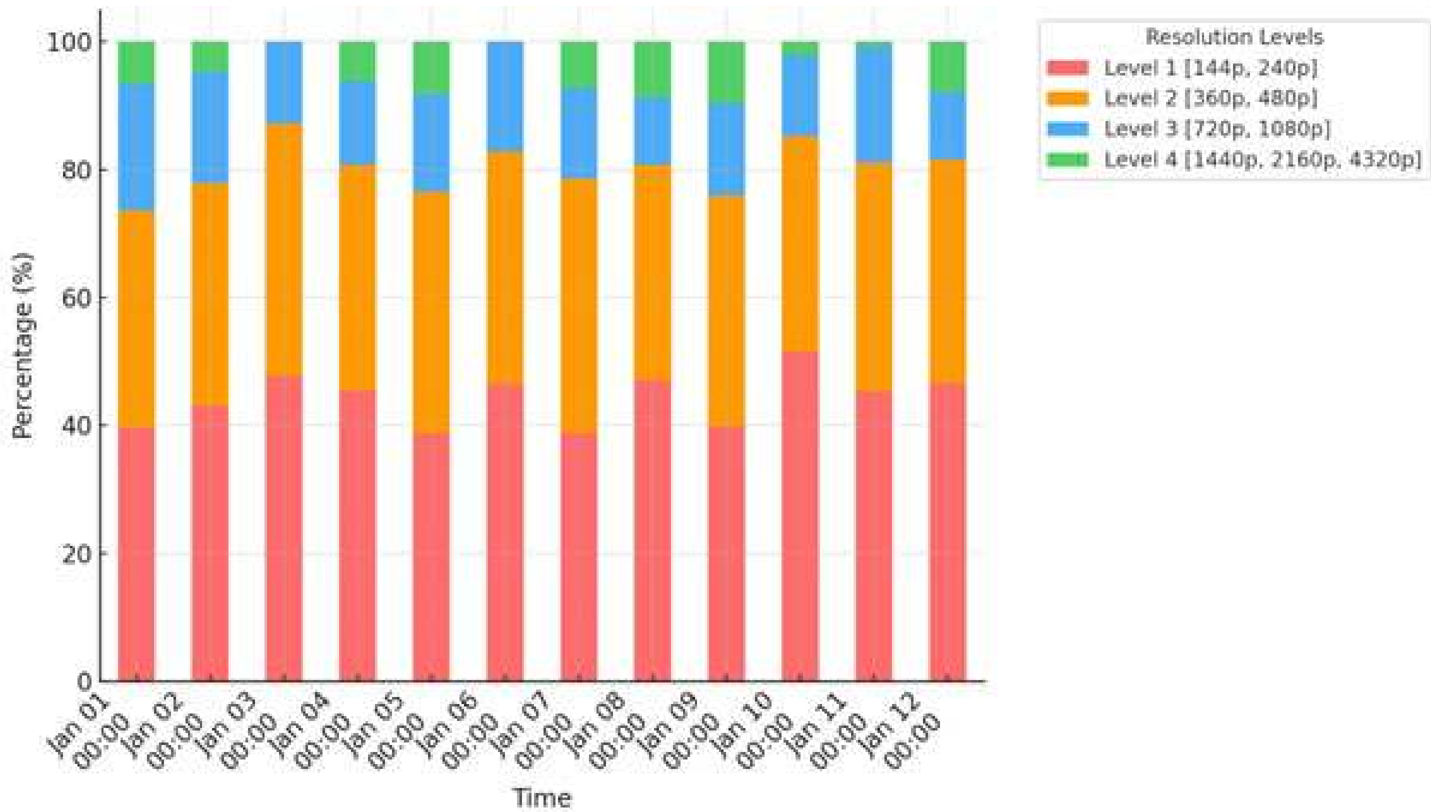

Figure 5: Relative video duration distribution

The figure below shows that Postpaid subscribers have a 25% higher vMOS score compared to Prepaid. The other diagram represents a premium Internet APN, which has a 30% higher vMOS than a standard Internet APN.

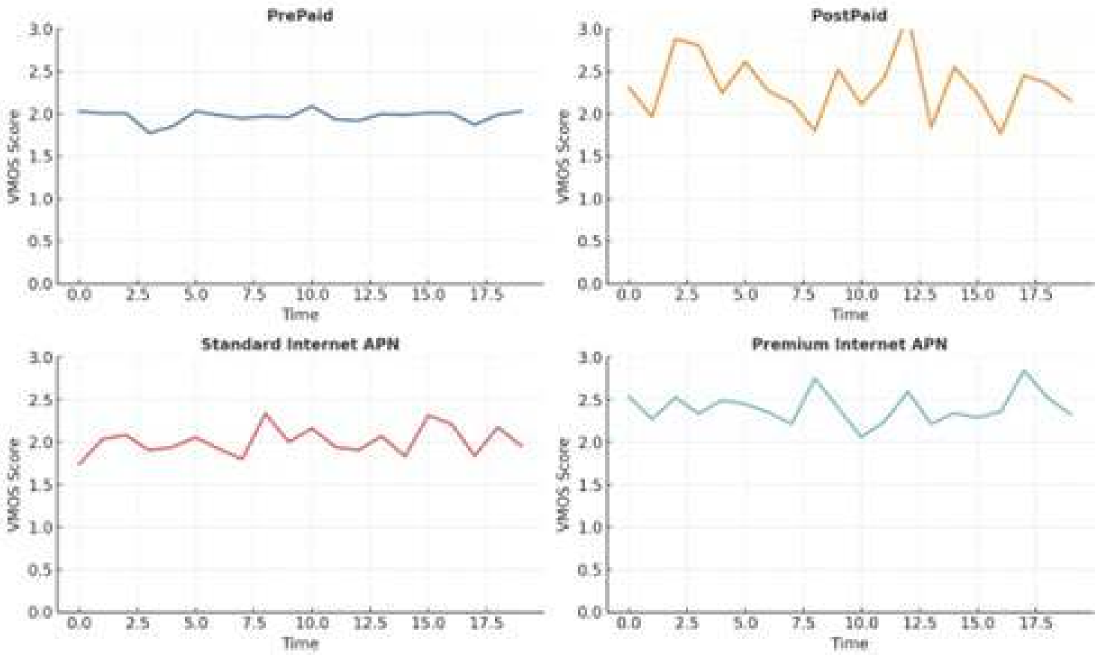

Figure 6: Comparing different vMOS scores

Conclusion

As video streaming becomes a dominant source of data traffic, operators must ensure the delivery of high-quality video services to protect revenue streams and meet growing customer expectations. With subscriber video consumption on the rise, the cost of carrying this traffic has also increased. To remain competitive, operators must evolve their business models to monetize video services effectively while maintaining superior service quality.

Delivering a consistently high-quality video experience is essential for differentiation and customer retention. To support this, RADCOM ACE offers advanced solutions specifically designed to address the unique challenges of encrypted video streaming.

Leveraging AI-powered analytics, RADCOM offers end-to-end visibility into subscriber Quality of Experience (QoE) and network performance for leading video platforms. Operators gain accurate insights into streaming quality and usage across multiple dimensions—including devices, roaming partners, locations, and APNs—enabling data-driven decisions and proactive service assurance. RADCOM offers powerful AI-driven intelligence that empowers you to understand and optimize customer experience, even in fully encrypted environments.