

BITMOVIN








COMICS

RACE INTO THE FUTURE OF STREAMING

IN CELEBRATION OF THE 10TH EDITION OF THE BITMOVIN VIDEO DEVELOPER REPORT







Welcome

I am proud to introduce the 10th edition of the Video Developer Report, built from insights shared by 486 respondents across 94 countries. Respondents span every corner of the industry, from OTT and broadcast to eLearning and publishing, and range from developers and engineers to C-suite executives at organizations from early-stage startups to Fortune 100 companies.

The honest framing for this year's report is that even calling it a year in review undersells how fast the ground has shifted. The pace has accelerated month on month, and the last three months alone have reshaped what it is possible to build. AI has stopped being a feature teams

add to the stack and started being a participant in it. Advertising has become the primary commercial growth engine in streaming. The QoE signal has been promoted from dashboard report to load-bearing infrastructure. And the first serious proposal to rebuild video delivery at the transport layer is moving through the standards process with real industry momentum behind it.

Ten years of this report means ten years of watching the same questions track the industry's direction. Latency, quality, monetization, and tooling have been the persistent themes. What changes year to year is how far each has moved from the frontier toward the expectation. This year, AI crossed that line. Low latency crossed it earlier and is now the infrastructure problem behind MOQ. The report this year is about what comes next at each layer.

I hope this year's findings help you benchmark your own workflows and understand where the market is heading.

Igor Oreper
Chief Strategy Officer @ Bitmovin

CONTENTS

Welcome	5
Key findings	8
The State of Video in 2026	10

AI ACROSS THE VIDEO STACK

Foreword by Jan Ozer	16
AI in the workflow: From feature to primitive	17
The deeper shift: Who is operating the workflow	20
MCP: The connective tissue	24
AI in the viewing experience: Where disruption is starting	25
Trust in the content itself	26

GROWING THE AD BUSINESS

Foreword by Dan Rayburn	30
Where the revenue is	31
CTV: The most valuable screen is also the hardest	32
New formats, new inventory	33
The technology that unlocks scale	34

QOE, PLAYBACK AND OBSERVABILITY

Foreword by Dan Rayburn	40
Detection time is the whole game	41
The player sees what nothing else can	43
The device landscape and the player codebase	45
A dashboard is not data	48
Detection time and the recovery window	50

LOW LATENCY, MOQ AND THE FUTURE OF VIDEO DELIVERY

Foreword by Jan Ozer	54
Twenty years of borrowed infrastructure	56
MOQ: What it is, and the honest state of it	58
Low Latency is the catalyst, not the point	60
The Contribution Layer: Proof That UDP Works in Production	61
The business case: Cost and efficiency	62
CDN: The delivery infrastructure layer	63

Methodology	66
-------------	----

KEY FINDINGS

AI Across the Video Stack

- AI is no longer a feature you add. It's what the stack runs on, shifting the question from what it can do to what you let it operate.
- 46% of developers use AI coding tools daily, only 4% don't at all. Cost isn't the limiter anymore. 35% cite accuracy on domain-specific work, a different complaint than two years ago.
- Agents now operate the workflow. Operators describe outcomes; agents call the platform.
- Transcription leads AI use cases at 48%. Scene and ad placement detection at 28% leads straight into Chapter II.
- Provenance is the next trust problem. Better generation tools also make fabricated footage harder to catch, pushing C2PA and watermarking into infrastructure.

Growing the Ad Business

- Subscriber growth matured; advertising took over. AVOD leads at 69%, FAST and SVOD level at 54%. FAST caught SVOD in two survey cycles.
- CTV is the most valuable screen and the hardest to execute on. Devices weren't built for advertising, so the Player carries the load.
- Pause ads (45%) and squeezebacks (36%) are shipping inventory now, not experiments, with an unforgiving quality bar.
- SGAI at 12% two years post-launch is small but significant. One shared manifest per audience means addressability stops scaling with size.
- Fill rate is now real-time, putting ad performance data in front of sales, not just engineers.

QoE, Playback and Observability

- You can instrument every server and still not know if the viewer had a good time. Backend telemetry reports intent, not experience.
- Detection time is the recovery time objective. No recovery promise survives a detection gap.
- The Player is the only place a session is visible in real time. Everything upstream reports intent, not outcome.
- Dashboards aren't data. Aggregated views suit reporting, not automated response, which needs raw per-session events.
- The next reader of QoE data is a machine. MCP and AI assistants debut as access methods this year (9%, 7%). Automated response inherits the signal's flaws.
- Mobile isn't an edge case. Native iOS/Android (49%/47%) and mobile web (47%/45%) mean the Player needs one consistent signal across native, browser, and CTV.

Low Latency, MOQ and the Future of Video Delivery

- Low latency was never the prize, a transport built for video is. Video's run on HTTP/TCP for two decades for lack of alternatives.
- MOQ planning intent: 28%. Production use: 2%. That gap is the story.
- The business case is cost and efficiency at scale, not speed.
- Latency is the catalyst, not the destination, new transport only pays off in a narrow band of live formats. SRT's 29% lead shows the shift already happened upstream.
- Honest counter: no published RFC, no proven use case at scale. MOQ is direction, not default.

The State of Video in 2026

Four pressures are shaping what video teams build in 2026. None of them are new problems. They are long-standing ones that finally got expensive enough to reorganize around.


Latency is the clearest case. The industry got it good enough to stop arguing about it, then discovered that good enough at scale costs more than anyone budgeted for. Cost control is that same lesson applied more broadly: once a capability is good enough to stop worrying about, what is left is the bill for running it at scale.

Monitoring and root cause sit underneath both, because nothing gets controlled that cannot be seen in time to act on. And advertising has


become the thing the rest of the stack is in service of, because that is where the revenue moved.

The survey tracks all four challenges directly. Live latency led the list at 36%, ahead of cost control at 35%, consistent monitoring and analytics at 28%, finding root cause at 27%, and ad insertion at 27%.

The bottom of the same list tells the other half of the story. Viewer engagement sits at 11%, quality of service under bandwidth restrictions at 8%, and content provenance and authenticity at 6%. Delivery problems outrank audience problems by roughly four to one.




What are the top three challenges you are experiencing?



Where do you see the most opportunity for innovation?

On the innovation side, advertising led the question at 22%, nearly double any other category, with AI-generated recommendations at 13% in second place. Below that, it's a genuine

four-way tie: content protection, enhanced color, low- and ultra-low-latency, and live streaming at scale all land at 8%, so no single technical pain point stands out from the rest.



The figure worth pausing on is low latency's fall from the top spot: it led the innovation chart at 43% in 2021 and sits at 8% today. Teams stopped chasing it because low latency did the job it needed to do for most content. What's left is narrower and more commercial: live sports, betting, auctions, and other cases where sub-second delivery pays for itself, which is exactly the audience MOQ is built for.

That is the shape of the year, and it is how this report is organized. Chapter I covers AI, because it has become the layer teams build everything else with. Chapter II covers advertising, because that is where the growth is. Chapter III covers what viewers actually experience, and how long it takes teams to find out when that breaks. Chapter IV covers the transport underneath all of it, and why replacing HTTP with something built for video went from a research question to a roadmap question.

AI Across the Video Stack

AUTHORS



Jacob Arends
Senior Product Manager / AI Scene Analysis & Playback @ Bitmovin



Patrick Weber
Director of Engineering / Player Web SDK @ Bitmovin

AI at scale in video is now possible in ways it simply was not six months ago. The next six months will look just as different again.

Two years ago the open question was whether any of this would work in production. It does, and it got there faster than most roadmaps assumed. The questions now are architectural. Which decisions do you hand to an agent, what does a platform have to expose for that to be safe, and who is accountable when the agent gets it wrong.



This chapter's data paints multiple pictures of how AI is being used today in production, development, operations, and user experience, and how it will likely be used in the near-term future. The first is where AI is already paying for itself. Transcription tops the list, followed by recommendations, scene detection, and video quality optimization. These are tasks where the return on investment is easy to measure.

The second story is the rapid adoption of AI coding tools. Nearly half of respondents report using them every day, suggesting they've become part of the standard development toolkit. Only 4% report not using them at all, a remarkably small percentage given the scope of roles the respondents perform.

The third story is about what's next, specifically the Model Context Protocol (MCP), which enables AI to work directly with APIs and software services to automate a range of operational tasks. It's early days, but if MCP takes hold, the savings could be far greater than from production capabilities like transcription or coding assistants.

The chapter then shifts to the viewer experience, where AI is increasingly being applied to content discovery, recommendations, and personalization. These technologies extend viewing time, increasing ad revenue and reducing churn, and are faster and cheaper than traditional alternatives. The final section examines the impact of AI-driven media creation. As AI-generated media becomes more common, so does the need to establish authenticity. Eighteen percent of respondents report creating media with AI, while 6% cite content provenance as one of their top three challenges. That percentage seems certain to grow as generative AI becomes more capable and more widely deployed.

Whether you're evaluating AI for the first time or looking for the next opportunity to improve productivity, reduce costs, or enhance the viewer experience, this chapter provides a useful snapshot of where organizations are finding value today and where the next opportunities for productivity and cost savings may emerge.




Jan Ozer

Founder @ Streaming Learning Center; Contributing Editor @ Streaming Media

AI in the workflow: From feature to primitive

The first place AI has earned its keep is in the workflow, because the ROI is measurable and the inputs are already digital. The pattern is consistent: AI replaces a step that used to be manual, expensive, or simply impossible at catalog scale. Advertising optimization and ad placement, search and recommendation, content discoverability, automated highlight and clip generation, contextual metadata for adtech: these are the jobs AI is doing in production today.

The survey shows how broad that adoption has become. Across all 486 respondents, 98% report at least one AI or ML use case for video, and the use cases span from the immediately operational to the emerging. The pattern is the same across all of them, AI earns adoption where the ROI is demonstrable today.



For which of the following video use cases do you expect to use machine learning (ML) or artificial intelligence (AI) to improve the video experience for your viewers?

Transcription at 48% is the only use case above 40%. Recommendations follow at 34% and video quality optimization at 30%, then


seven use cases cluster between 23% and 30%. Generating video or image content is at 18%. Only 2% report no plans for AI at all.

Speech and audio remain the industry's most dependable entry point into AI for video. Transcription and speech-to-text has ranked first in every edition since 2021 and sits at 48% this year. The input was already digital, the output is cheap to verify, and the manual alternative was expensive, so nobody needed convincing.

The rest of the list divides along a clear line. Recommendations at 34% and video quality optimization at 30% make existing work faster. Scene, shot boundary, and ad placement detection at 28% enables work that was previously too expensive to attempt at catalog scale, and it is the same capability the advertising chapter depends on.



AI/ML use cases for video: key trends over time (2020-2026)



Which AI coding tools do you use for development?

The deeper shift: Who is operating the workflow

The more consequential change is not which AI tools teams are using. It is who, or what, is operating the workflow. The dashboard is no longer the only entry point. Agents are. Developers and operators no longer click through configuration screens. They describe what they want to a coding agent, and the agent calls the platform. The video stack is becoming a tool that other tools use.

The developer toolchain changed faster than any other part of the video stack, and it changed without anybody deciding to.


Nobody ran a procurement cycle for AI coding assistants. Engineers picked them up, the tools worked, and now they are load-bearing.

The distribution shows how that happened. OpenAI Codex and ChatGPT lead at 56%, Claude Code at 50%, Google Gemini Code Assist at 42%, and GitHub Copilot at 37%. Those shares total well past 100%, which is the interesting part. Teams are running several assistants side by side, not standardizing on one. 46% use them daily as a core part of their workflow and only 4% do not use them at all.


"AI is becoming a core workflow primitive across the video industry. Beyond established automation, LLM-powered search, editing and analytics capabilities over MCP are creating measurable value across internal operations and customer-facing experiences."



Anastasia Kundush
Head of Product @ Alpic




How often do you use AI coding tools?



What do you primarily use AI coding tools for?

What developers hand over says more than how many of them use the tools. The work being delegated is the work where being mostly right quickly beats being entirely right slowly. Debugging leads at 68%, writing new code at 60%, and prototyping at 59%, all tasks with a

tight feedback loop where a wrong answer is cheap to catch. Documentation at 42% and code review at 52% show the tools have moved past the enjoyable parts of the lifecycle into the parts nobody volunteers for.



What limits your use of AI coding tools the most?

That same logic explains the complaint. Accuracy and reliability of output is the main barrier at 35%, while cost sits at 10%. Two years ago the blockers were access and budget. Now that the tools are everywhere and inexpensive, what remains is output quality on domain-specific work. That is a harder problem and a better one to have.


MCP: The connective tissue

MCP (Model Context Protocol) is what makes the agent-addressable stack real. It lets a coding agent in one place call a media platform in another, and it is what keeps agentic experiences production-grade. An API an agent can drive has to be callable from natural language and composable across services.

The survey reflects exactly what developers want from AI when it meets their platforms. 58% say the most useful thing an AI tool can do with a third-party SDK is provide pre-built integrations to the vendor API. 51% point to AI that can read and search vendor documentation. The AI tools that understand the platform are worth more than generic code assistants.

In a single year, developer expectations moved from AI that writes code to AI that operates systems. That is a materially different ask of a platform vendor, because it means the API surface has to be legible to something that is not a human reading documentation.

Pre-built MCP server integrations barely registered a year ago. Teams have stopped asking for help writing the integration. They want the platform to already be callable. That is what agent-addressable means in practice, and it is now a procurement criterion.



When using AI tools with third-party SDKs or APIs (like a video player or encoding service), what do you find most useful?

AI in the viewing experience: Where disruption is starting

The second shift is more aspirational and just starting. AI is moving from the back office into the viewing surface itself, in three places. Like we saw earlier in the report, video teams are already using AI to power recommendations and personalization at a meaningful scale. The next question is who owns that layer, the streamer's own app, or the AI tool the viewer opens first.

External AI tools are becoming the discovery layer. As we saw earlier, Recommendations already ranks among the top AI use cases at 34%, and viewers are increasingly asking Claude and ChatGPT for suggestions directly instead of browsing a home screen. For media companies, that means catalog metadata needs to surface into the tools people are already in, through MCP servers and agent-friendly APIs that expose the streamer's catalog into the discovery layer.

The AI tool could become the playback layer itself, not just the discovery layer. The standards underneath, MCP and open SDK interfaces, are maturing quickly. Whether they can carry rights-managed content, addressable ads, and session analytics at production

grade is the open question. If they can, the AI tool stops being a recommender that bounces viewers back into the streamer's app and becomes a viewing surface in its own right. The streamers who expose their content into the agent layer first, with rights and monetization intact, keep ownership of the viewer relationship even when the entry point is not their own app.

Personalization already ranks among the top AI use cases at 28% today, mostly running server-side. On-device LLMs are the parallel shift to watch. Models are getting smaller in data footprint and larger in parameter count simultaneously. That is the combination that makes them practical to ship inside TVs, phones, and set-top boxes. The streamer stops paying the inference bill for personalization because the compute runs on the user's hardware. Picture quality increasingly becomes the device's job. New input modalities, voice control, semantic search, live translation, arrive bundled with the device at no additional cost.

Trust in the content itself

A year ago the trust question in AI was about the orchestration layer: could you trust an agent to drive your encoding pipeline, your ad break decisions, your recommendations? In 2026 that has largely moved from blocker to managed risk. The guardrails around modern LLMs have matured fast.

The next concern is trust in the content itself. The same investment making AI-generated film and video credible is making fabricated footage easier to produce and harder to detect. For news outlets and broadcasters who validate sources at speed, content provenance standards like C2PA are quietly becoming important infrastructure. Content provenance

ranks last among the challenges respondents named this year, at 6%. In the 2026 survey, 11% are already using C2PA for content provenance and 11% are using watermarking, with a further 16% planning to adopt watermarking within 12 months. Worth building now, while it is still cheap to build.

AI in 2026 is no longer one feature on a roadmap. It is a participant in the stack: building the encodes, indexing the catalog, surfacing the content, and increasingly playing it back. Stacks that are agent-addressable, MCP-native, and content-provenance-aware are the ones that will compound in value over the next year.

“Whether you feel that video provenance is wholly related to AI or not, it’s certainly a growing concern for service delivery. Consumers are starting to demand a way to check content authenticity at the point of consumption, and C2PA offers a flexible, mature framework to address that concern.”



Steve Christian
Fractional CMO @ EZDRM

“AI’s future value in video is operational, not novelty. MCP connects agents to APIs, documentation, metadata, and observability data so customers can improve content readiness, discovery, viewer experience, and monetization while turning complexity into repeatable workflows.”



Ed Laczynski
General Manager @ Zype

WHAT THIS MEANS FOR VIDEO TEAMS

- ❑ Start measuring AI tool output quality on domain-specific video workflow tasks, not just general benchmarks. Accuracy on encoding configuration or pipeline code is what matters.
- ❑ Make your platform agent-addressable. Teams exposing their video stack to MCP-compatible agents will integrate faster with the tools developers are already using daily.
- ❑ Invest in content provenance infrastructure before it becomes a regulatory requirement. C2PA and watermarking are cheaper to build proactively than reactively.

Growing the Ad Business

AUTHORS



James Varndell

VP Product / Playback @ Bitmovin




Lukas Kröpfl

VP Engineering / Player & Support @ Bitmovin

If you want revenue growth in your streaming service in 2026, advertising is where it is coming from.

For most of the last decade, the streaming growth question was how many subscribers you could add. That question has been answered, and the answer is not many more. Acquisition costs rose, churn stayed stubborn, and the addressable pool of people willing to pay for another subscription stopped expanding. So the industry did what industries do when the primary revenue line flattens. It went looking for the second one.

That second line is advertising, and the largest operators have already reorganized around it. Morgan Stanley estimates that in 2025, over 100% of Netflix's and Disney+'s US net subscriber additions came through their ad-supported tiers, with ad-free numbers actually shrinking. The entire net growth of the two biggest streamers is advertising-led.



Respondents see it from the inside. Advertising is the number one innovation opportunity they name in their own services, holding that position for a third consecutive year at 22%, and 28% are already using AI for scene and ad placement detection. Both numbers trace back to the same commercial pressure showing up in two parts of the stack. The opportunity is real, and it is also the hardest part of the stack to get right.

While originally limited to FAST business models, advertising across OTT platforms quickly became a must-have for nearly every streaming service. With consumers willing to sacrifice ads for a lower subscription cost, ad-supported tiers drive the majority of new signups and contribute revenue that grows with usage, as the total number of viewing hours across all streaming services rises each year. Advertising has become an economics-of-scale business, with subscriber growth, retention and profitability of OTT services all depending heavily on ads.

With advertising now more important than ever, streaming services must focus on targeting and delivering the right ad to the right user, on the right device, and in the right format. With no standard ad tech or measurement across the video stack, many are still experimenting with new ad formats like pause ads and squeezebacks, which continue to show traction. Ads within short-form and long-form video vary by device, and vertical video is also changing the way ads are created and delivered.

To increase revenue, many streaming services have doubled the number of ads they show per hour since launching in the market, but they also need to balance what consumers will tolerate in terms of user experience with the results advertisers expect. A lack of standards in ad measurement makes this challenging when advertisers buy across platforms with content that is live and on-demand, as well as one-off large sporting events.

Ad-supported streaming services are crucial to the success of the streaming market, especially as many services have since become profitable, and even though challenges remain in selling, delivering and measuring inventory, the industry is making a lot of progress each year.




Dan Rayburn

Streaming Media Industry Analyst, Writer & Consultant

Where the revenue is

The interesting thing about the monetization mix in 2026 is not that advertising won. It is how fast the free, ad-supported channel closed the distance on the paid one. FAST launched into a market that already had a mature subscription habit, and it did not need to displace that habit to work. It only needed

to reach the people who were not going to subscribe to anything else. The survey puts numbers on it. AVOD leads at 69%. FAST entered the survey in 2024 at 46% and now sits at 54%, effectively tied with SVOD at 54%. Matching the subscription model in two cycles is the most significant monetization shift in this edition.



“The future media company is an intelligent software company. Its advantage will come from how quickly it can compose, govern, monetize, observe, and continuously improve every audience experience globally.”



Scott M Davis

CTO, Media & Entertainment @ GlobalLogic

CTV: The most valuable screen is also the hardest

Connected TV is where the advertising opportunity is largest and the execution challenge is steepest. Watch times on CTV run 2 to 3 times longer than mobile or web. Completion rates are higher and CPMs carry a premium, and ad budgets are migrating across from linear broadcast as audiences do.


CTV devices are not built for advertising. There are more TV operating systems than ever.

Devices are often low-cost with a single video decoder and limited performance, so the heavy lifting of delivering a quality ad experience has shifted from device hardware to video software. Put directly: CTV devices were not designed with advertising in mind. The Player now carries most of that load. Which means ad delivery quality is fundamentally a Player problem, not an ad server problem.

“While CTV’s massive engagement makes it the crown jewel of digital video, publishers will only unlock its full monetization potential when we bridge the historical divide between ad decisioning and client-side execution. In a highly fragmented device landscape, ad delivery quality isn’t just a player challenge or an ad server challenge: it is an integration mandate. True scale requires the major ad servers and dominant video players to act as a singular, unified engine, translating complex programmatic decisioning into flawless, broadcast-quality streams on even legacy hardware-constrained devices.”



Lemu Coker
CTV Ecosystem Partnerships Manager @ Google



New formats, new inventory

There are two levers for sustainable advertising growth. First, new ad formats that create fresh inventory without simply adding more full-screen breaks. Second, data that shows where fill is low by daypart, content, platform, and cohort, so it can be sold and filled. Ad Sales Intelligence is becoming a new product category that ad sales teams need.

New ad formats are not a design exercise. Streaming ran out of places to put linear spots without damaging retention, so the industry went looking for inventory that interrupts nothing. Pause ads sell attention that was already idle. Squeezebacks sell screen space,

not time. Both exist because the obvious inventory was already sold.

Linear ads still lead at 67%, but the momentum sits with formats that did not exist in most streaming workflows three years ago. Pause ads are at 45%, squeezeback L-shape and J-shape at 36%, and overlay ads at 29%, all in production at meaningful scale. The catch is that these formats fail in public. A frozen pause ad or a squeezeback that renders wrong at scale gets covered in national press, not trade press, and the reputational cost lands on the publisher. Technical reliability is a commercial requirement here, not an engineering preference.

The technology that unlocks scale

Streaming devices are not built for advertisers, so the software does the heavy lifting. The shift from SSAI to SGAI is the architectural change that makes addressable advertising work at scale, and it is built on open standards, HLS interstitials and SVTA2053.

Personalized advertising and efficient delivery have been in tension since streaming started carrying ads. The industry's answer for a decade was to give every viewer their own manifest. It works, and it stops working at roughly the moment a service gets large enough to care about what it costs.

That is the bill Server-Side Ad Insertion carries. SSAI has grown to 81% in 2026, its highest ever share, and it personalizes by stitching a unique manifest per viewer. Server-Guided Ad Insertion keeps one shared content manifest and moves the ad decision out to the Player at playback, so each viewer still receives their own pod with targeting applied and nobody pays to generate manifests per person. That is what makes addressability work at scale, and what makes ad replacement possible inside live DVR windows. SGAI entered the survey in 2024 and already sits at 12%. For an architectural change in an industry that does not rewrite its packaging layer casually, that is a fast start.


SGAI on open standards means a regulated public broadcaster can serve compliant, targeted ads into hour-long DVR windows without generating a unique stream per viewer.

"An ad failure is a revenue failure. We measure ad playback with the same rigor as content."



Amit Mathur

Sr. Vice President of Product Engineering @ Sinclair Broadcast Group



"New ad formats are benefitting from industry-wide standardisation for the first time. This is significant because it means there will be consistency when it comes to the way the formats look, how ad calls are made, how they are traded, and how measurement is applied to provide brands with the level of attribution data they require. I expect to see a marked increase in their use for SSAI and SGAI over the next year or two."



Paul Davies

Head of Marketing @ Yospace

Proof follows the same architecture shift, and the stakes are already established earlier in this chapter, when a new ad format fails, it fails in public, with reputational cost landing on the publisher, not the ad server vendor. The same exposure applies to attribution. Verifying that an ad played the way it was bought means reconciling delivery against what was promised, and under a fully personalized SSAI manifest, that means auditing a unique stream per viewer. SGAI's shared manifest model makes that reconciliation simpler in principle, one manifest to check instead of millions. Whether teams have actually built that verification layer yet is the open question for next year's edition to track.

Data and AI turn delivery into commercial advantage. Observability for video playback surfaces fill rates in real time, by daypart, by

content, by platform and cohort. AI puts that data in the hands of non-technical ad sales teams, closing the gap between inventory insight and deal-making. The teams who can act on performance data within the session will fill more inventory than the ones waiting for the next morning's report.

Two shifts look likely over the next 12 to 18 months. AI moves further into the ad operation itself, with ad sales teams asking questions of performance data using AI and acting mid-session. And formats converge with buying, as pause ads, squeezebacks, and shoppable ads move from experiments toward standard sellable inventory, with SGAI on open standards moving from cutting edge toward default for addressable advertising. Both are directional calls, not certainties.

“Ad and viewership monetization lives or dies on client-side beaconing, and beaconing is brutal at scale: every viewer’s device reporting in every few seconds. When a large event hits 15M concurrent devices, at peak this means around 2M beacons per second. If the telemetry holds, the money holds.”



Lionel Bringuier
GM Media & Entertainment @ Memento.

“Performance marketers want more than ever to diversify beyond saturated Search, Social and Retail Media. So it’s up to the industry to work on enabling scaled free ad supply, lean-back friendly SendToPhone / AddToCart feats and accurate data to size the opportunity and make TV Advertising grow again.”



Julien Boyreau
Business Development Director, Advertising @ Broadpeak

WHAT THIS MEANS FOR VIDEO TEAMS

- Treat ad delivery as a Player problem, not an ad server problem. CTV devices cannot carry the load. Your software has to.
- Evaluate SGAI for any live or DVR workflow where addressable advertising matters. The manifest-per-viewer cost of pure SSAI becomes a ceiling at scale.
- Build real-time fill-rate visibility by daypart, platform, and cohort. The teams acting on live data within the session outperform those waiting for next-day reports.

QoE, Playback and Observability

AUTHORS



Reinhard Grandl
CPO @ Bitmovin



Daniel Hoelbling-Inzko
Chief Architect & VP Engineering / Analytics & CA @ Bitmovin

You can hit every backend SLO you set and still find out your stream broke the same way your customers did: from an angry post twenty minutes too late.

By the time that signal arrives, the incident is over and the churn is already booked. A recovery time objective is a promise to fix things inside a window. No such promise is credible if you cannot detect the problem inside that window. And for viewer experience, detection can happen in exactly one place in real time: the Player, on the device, at the end of the line.

This is why the QoE signal has been quietly promoted from back-office report to load-bearing infrastructure. For a decade, quality of experience was something a human read off a dashboard the next morning, so sampled and batched data was perfectly fine because nothing acted on it until a person did. That era is ending. The thing consuming the QoE signal is increasingly a machine acting in the moment: a server-guided ad decision choosing a spot at playback, an ad-sales layer chasing fill, an AI assistant acting on the data directly. The instant something acts in real time, the signal feeding it has to be real time too. And it has to be true.



With almost no standards in the streaming industry, including end-user devices, protocols and the definition of “video quality,” monitoring, measuring and tuning QoE in real time is one of the most difficult pieces of the video ecosystem. Anyone can deliver a stream. But delivering the right video to the right user, for a specific device, with the best QoE possible, at scale, is difficult.

The five key QoE metrics most streamers use today are video start time, buffering/re-buffering, latency, VMAF score and max bitrate, but this can vary based on the business model and type of content. Broadcasters and OTT platforms invest a lot of work in proactively identifying QoE problems before they happen, especially with large-scale live events, but telemetry data can only go so far.

Player data provides the insight needed, but aggregated numbers don’t define or highlight the user experience on their own. Dashboards can look great but provide an outdated picture, since any delay in ingesting data slows detection. Streamers work hard to balance what data to collect, when to collect it, how often to collect it, and what they think matters most to viewers. AI technology is expected to help analyze the data faster and provide better proactive insights, but AI by itself isn’t a quick, simple or easy integration into what’s being used today.

The survey results on video quality provide a unique window into a key aspect of the streamer’s business, giving insight into how they think about QoE, the metrics that matter most and the impact it has on their business.



Dan Rayburn

Streaming Media Industry Analyst, Writer & Consultant


Detection time is the whole game

Sampled and batched telemetry sets the true detection floor to whenever complaints cross a threshold, not whenever the problem actually started. Real-time, per-session QoE from inside the Player is the precondition for any honest recovery objective on experience.


Consistent monitoring at 28% and finding root cause at 27% both rank among the top operational concerns in the challenges question we saw earlier in the State of Video in 2026

section, which is one problem described from two directions. What teams choose to measure inside the session tells you which half of it they have stopped worrying about.


Buffering leads at 57% and error rates at 43%. Video start time has fallen to 28%, down from 64% in 2020, the clearest sign that the industry solved a first-impression problem and has moved on to what happens for the rest of the session.




Which video performance metrics are most important to you?



That shift reflects what the industry has already solved and what remains hard. Getting video to start quickly is now table stakes. Keeping it reliable across the full session, across a



fragmented device landscape that spans browser, mobile, and a dozen CTV platforms, is where the operational complexity now lives.




Video performance metrics most important to you: trend over time (2019-2026)

The player sees what nothing else can

Observability for video playback is instrumented inside the playback engine itself, not bolted on as a page or app beacon. It sees decoder state, buffer levels, bitrate switches, dropped frames, startup phases, and Player error codes, the internal truth of playback that is invisible from anywhere outside the Player and cannot be reconstructed after the fact. This is the honest line between playback observability and the backend monitoring market.

Application performance monitoring is crowded and largely solved. It measures what the system intended to deliver. Viewer truth measures what the human actually saw, and it can only be collected where it happens. The two are complementary, not competing. When a CDN fails over or an ad server times out, the proof of whether viewers were hurt lives in the QoE data, and the QoE pattern points teams at which backend system to investigate.

Observability solution used: trend over time (2020-2026)




The provider landscape says more about how teams buy than about who is winning. Nobody is consolidating. The spread across vendors has tightened year on year, which is what happens when teams stop hunting for one tool that covers everything and start assembling a

stack where each piece answers a question the others cannot. Bitmovin's Observability sits at 27%, alongside NPAW at 36%, Google Analytics at 31%, and Conviva at 30%.


The device landscape and the player codebase

The surface area teams are responsible for shapes the entire QoE challenge. 80% support HTML5 in browser. Native iOS at 49% and Native Android at 47% confirm that mobile is a primary delivery target for nearly half the industry, and HTML5 on Android at 47% and on iOS at 45% mean most teams ship mobile web alongside native. Apple TV at 47%, Android TV

at 41%, Samsung Tizen at 38%, and LG webOS at 38% complete the major CTV platforms. That is not just a support-matrix problem. A native iOS crash, a mobile browser stall, and a CTV decoder error each look different at the Player, and a QoE pipeline that treats them as one signal will miss all three.




Which of the following platforms and devices do you support for streaming video or audio content?



Each platform has its own Player implementation requirements, its own performance characteristics, and its own failure modes. Maintaining a consistent Observability layer across all of them simultaneously is the operational challenge the data reflects.

57% build in-house on an open-source or native foundation, 28% buy commercial, and 12% run a hybrid. Only 3% build from scratch, down from between 24% and 29% in 2020 and 2021. For those 57%, playback instrumentation is something the team configures and maintains itself.



How a team builds its Player decides what that team can see when playback goes wrong, which makes it a QoE question. Open-source and native foundations hand you control over instrumentation along with the maintenance bill. Commercial players invert that trade.


What almost nobody does anymore is start from nothing. The shift away from fully custom

builds has been steady since 2020, and it has not reversed in a single year across the whole period. Writing a Player from scratch stopped being a defensible use of engineering time, and the data shows an industry that made that call once and never looked back.

A dashboard is not data

Most measurement tools hand teams a governed, sampled, aggregated dashboard. Fast detection and automated response depend on something different, the raw event stream itself, per session, at full granularity, in real time, owned by the team and routable wherever they already work. Without that, both stay aspirational.

General-purpose analytics tools sample and aggregate by default, a design suited to marketing funnels rather than sub-second playback. A dashboard is a view of the data. The event stream is the data itself. That distinction matters for correlation against backend systems too, since building a real recovery objective depends on controlling data at the event level, not just viewing a summary of it.



How teams access Observability data: trend over time (2020-2026)

The audience for QoE data is changing, and that changes what the data has to be. A dashboard works when a person can review a number that is an hour old and decide what to do about it. Automated systems do not review anything, they act on whatever data is already in front of them at that moment.

The shift shows up in how teams reach their data. Analysis tool dashboards reached 75% in 2026, but direct API and database queries have grown from 21% in 2020 to 45% today. MCP connectors and AI assistants appear for the first time this year, at 9% and 7%.

Those last two figures are small, and they are the direction of travel. An agent is only as good as the signal it acts on, which means real-time speed, Player-side fidelity, and ownership of the raw event stream. Remove any one of the three and it will act confidently on a second-hand picture it cannot see behind.

Advertising playback QoE and advertising engagement metrics debut as tracked metrics in 2026, at 8% and 6% respectively. The ad experience is now part of the QoE measurement. A missed ad, a slate where a creative should have been, or a buffer event during a break is a revenue event visible only in the Player.


The dashboards were always for us. The next reader of your QoE is a machine, and it will only be as trustworthy as the signal you let it see.



Simon Ouderkirk
VP of Product @ Hydrolix

“Video start time is a resource, not an error. Real per-session data shows viewers barely notice a few extra seconds of startup, but they leave over rebuffering. The mature streamer’s move is pinning VST to actual viewer tolerance per viewing platform, and spending the headroom on session stability.”

Maintenance time vs. root cause identification time (2024-2026)




Detection time and the recovery window

These are two separate clocks. Maintenance time is how much of the week goes to keeping the service running. Root cause time is how long one incident stays unexplained. Across the 2024, 2025, and 2026 groups the maintenance bars shorten while the longest root cause bands grow their share.

The gap between something breaking and somebody knowing why is where viewers leave. Backend telemetry will not close it, because the systems doing the reporting are the ones that believe they are working correctly.

For a lot of teams that gap is still measured in days. 29% take 12 to 24 hours to identify the cause of a streaming issue, a further 20% take 3 to 5 days, and only 9% resolve within the hour. The trend underneath is more interesting than the headline. Median time has improved while

the share of incidents running past 3 days has grown. Common problems are getting easier and edge cases are getting harder, which is what happens when tooling improves faster than the systems it monitors.



On average, how long does your team take to identify the root cause of streaming-related issues?

WHAT THIS MEANS FOR VIDEO TEAMS

- ❑ If your QoE data is sampled or batched, your recovery time objective is aspirational. Real-time, per-session Player-side data is the precondition for an SLO on viewer experience.
- ❑ Start tracking advertising playback QoE alongside content QoE. A buffer event during an ad break is a revenue event. If you cannot measure it, you cannot manage it.
- ❑ Evaluate whether you own the raw event stream or just a dashboard view of it. Any automated or AI-driven response requires event-level data you control.

Low Latency, MOQ and the Future of Video Delivery

AUTHORS



Alex Zambelli

Director of Product / Encoding @ Bitmovin



Martin Smole

VP Engineering / Encoding @ Bitmovin

Video has spent 20 years as a guest on someone else's infrastructure. HTTP and TCP were designed to move documents, reliably and in order, over connections that were expected to be slow and occasionally broken. Video was made to fit that shape through segmentation, manifests, and adaptive bitrate, and the result works well enough that video now accounts for approximately 75% of global mobile data traffic (Ericsson Mobility Report 2026), and roughly 65% of all internet traffic according to Sandvine's 2023 Global Internet Phenomena Report.

Working well enough is not the same as being designed for the job. Every latency technique in production today is a workaround for a transport that was never trying to do this, and the cost of those workarounds shows up as buffer, as bandwidth, and as operational complexity. For most content that trade has been acceptable, and the survey confirms it. The business case for low latency simply is not there for a catalog of on-demand titles.

What has changed is that someone is finally attempting the other approach. MOQ, or Media over QUIC, treats video as a first-class citizen of the transport, not a payload layered onto the web. It is the first serious attempt at that in two decades, which is why it draws attention out of proportion to how much of it is actually running anywhere. The survey reflects exactly that imbalance. MOQ appears with 28% planning intent for distribution in its first year on the survey, against 2% in production today. Intent, not deployment. That distinction is the honest core of the story and this chapter keeps returning to it.

Every annual survey raises two kinds of questions. The first is straightforward: What technologies are organizations using today? The second is more intriguing: Which of today's trends will actually change the way video is delivered over the next several years? The data in this chapter helps answer both.

The distribution format data highlights the balance between proven technologies and emerging alternatives. HLS remains the dominant delivery protocol at 78%, followed by DASH at 66% and CMAF at 47%. At the other end of the spectrum is Media over QUIC (MoQ), with 28% of respondents planning deployments over the next year. Whether that translates into broad adoption remains to be seen, but it's one of the more interesting technologies to watch as the industry looks for more efficient ways to deliver live video at scale.

The latency results may be the chapter's biggest surprise. Five years ago, 60% of respondents considered under five seconds to be low latency. Today that figure has dropped to just 16%. Instead, responses cluster around two targets: 28% expect end-to-end latency under 10 seconds, while another 17% target under 30 seconds. Has the industry become more realistic about latency? Or have organizations concluded that reducing latency below a certain point simply isn't worth the additional cost and complexity? The data doesn't answer that question directly, but it certainly raises it.

Contribution workflows continue to evolve as well. SRT now leads all contribution protocols at 29%, ahead of HLS at 25% and RTMP at 19%. That's a remarkable increase from the previous report, where SRT measured just 7%. SRT and MoQ solve different problems, but SRT's continued growth suggests the industry is increasingly comfortable moving away from traditional TCP-based workflows when the operational benefits are clear.

The codec data raises another interesting question. H.264 continues to dominate production deployments at 85%, followed by HEVC at 58% and AV1 at 19%. Yet respondents are

more optimistic about deploying AV1 over the next twelve months, with 35% planning implementation compared to 26% for HEVC. Is AV1 finally approaching a broader deployment phase, or will this prove to be another year where plans outpace production use? VVC's projected 20% adoption over the next year was perhaps the biggest surprise in the codec data and will be worth revisiting in the next survey.

Finally, the CDN data suggests that some parts of the streaming workflow have largely stabilized. 80% of respondents rely on one or more third-party CDN providers, while just 10% operate their own CDN and 8% use a hybrid approach. For most organizations, the decision appears to have been made: building a CDN simply doesn't offer enough operational or economic advantage to justify the investment.

Like the best survey data, these results don't provide every answer. They do identify the technologies gaining momentum, those that continue to dominate despite predictions to the contrary, and, perhaps most importantly, the questions the industry will be debating over the next year.




Jan Ozer

Founder @ Streaming Learning Center; Contributing Editor @ Streaming Media

Twenty years of borrowed infrastructure

When online video delivery became viable in the early 2000s thanks to the spread of broadband internet, state-of-the-art streaming protocols were still based on designs from the previous decade. Proprietary protocols from RealNetworks and Microsoft, along with standardized RTSP, were built while the internet was still in its infancy, borrowing ideas from telecommunications networks that were closed and managed, not open and best effort. Typical connection speeds at the time ranged from 56 kbps to 1.5 Mbps, and reliable video delivery over the public internet at that speed was a real challenge.


That changed in the mid-2000s when Move Networks introduced a different approach, delivering media over standard HTTP in short, cacheable segments. Microsoft, Apple, Adobe, and eventually MPEG followed, and by the 2010s the industry had shifted almost entirely to HTTP-based delivery. That is the two decades this chapter keeps returning to.



Which streaming protocols are you using in production for distribution?

Every streaming stack in production today is a workaround for a transport that was built to move web pages, not video. HLS/TS leads at 78%, DASH at 66%, HLS/CMAF at 47%. The HTTP-based stack works. But TCP's reliability model and head-of-line blocking were never built for real-time media. Video now dominates internet traffic, mobile and fixed alike, and still runs on transport designed for web pages. That is the gap MOQ exists to close.

Encoder choice shapes the same economics MOQ is arguing about. Commercial encoders anchor live delivery for reliability under SLA, used by 73% of respondents for live workflows, more than double the 32% running in-house on open-source and the 19% on open-source directly. That preference for a supported vendor over saving license cost is itself a lever on the delivery cost equation MOQ is trying to change.




Encoder type used: Live

MOQ: What it is, and the honest state of it

MOQ starts from a different premise, which is to build the transport for video rather than around it. It runs on QUIC, the transport under most modern web traffic, and it collapses problems the current stack solves separately. One protocol can carry contribution and distribution, live and on-demand, at sub-second latency, over commodity networks. Before weighing what MOQ could change, it's worth seeing what the industry has done without it, using low latency workarounds instead of a new transport.


Low-latency HTTP did the job it was asked to do, which was to make live streaming tolerable without replacing the delivery stack. It succeeded well enough that the industry stopped treating latency as a crisis, and that success is exactly what defines the ceiling MOQ is arguing about.



Low-latency technology used over time (2021-2026)

Adoption looks like a mature workaround that is still growing. LL-HLS has grown from 33% in 2021 to 42% today, LL-DASH tracks at 37%, and the share using no low-latency approach at all has fallen from 37% in 2021 to 32% today. These are the layers MOQ would supersede, and all three are moving in the direction that makes the argument possible.

The gap between the two bars is the point: teams are already planning for MOQ, but almost none are running it in production yet. That trajectory points to something bigger than transport alone, a broader shift in how the underlying infrastructure needs to evolve to support it.



“The future of video delivery isn’t defined by latency alone, it depends on transport and edge architectures working together at global scale. MOQ is a key step, and distributed edge infrastructure improves resilience, simplifies operations, and supports emerging video workloads.”



Kip Schauer
Senior Product Manager, Media Partnerships @ Oracle Cloud Infrastructure


Low Latency is the catalyst, not the point

It is worth being direct about where low latency does and does not carry value. For most content there is no measurable business case for it, and forcing it adds cost and complexity that can cut into real-time ad-auction time. Low latency was the number one innovation opportunity in 2021 at 43%. In 2026 it sits at 8%, consistent with the State of Video in 2026 section earlier. That's the industry settling, not solving. The gap between the two numbers tells the real story: live latency still leads the challenge list at 36%, because an expectation that applies everywhere is harder to meet than a frontier a few teams are chasing. The innovation number fell because teams stopped looking for a breakthrough.

The challenge number stayed high because they still have to deliver it, at scale, on budget.


Expectations came down as understanding went up. In 2020 and 2021, a quarter of respondents said they needed sub-second delivery, largely because the industry had not yet worked out what sub-second actually costs. Now that it has, targets have moved to what the content justifies.

Sub-1-second expectations have fallen from 25% and 26% in 2020 and 2021 to 8% in 2026, and the under 10 seconds bracket leads at 28%. That is calibration, not lowered ambition.



Where sub-second still matters, the reason is commercial. Live sports, betting, auctions, and interactive formats are the cases that justify leaving the HTTP workaround behind, and they


are the same cases where the ad architecture from Chapter II starts to depend on the delivery architecture here.



The Contribution Layer: Proof That UDP Works in Production

SRT already leads live contribution at 29%, ahead of HLS at 25% and RTMP at 19%, not because TCP-based options are disappearing, but because UDP-based transport is proven and trusted in production. The distribution layer is what MOQ is positioned to address next.

SRT's lead is a proof point, not a precedent for replacement. Distribution is a bigger, more varied problem, so MOQ's case still depends on ratification completing and a production use case at scale being demonstrated.



The business case: Cost and efficiency

Speed is what gets MOQ attention. Cost and efficiency at scale is what would get it deployed. Video's growing share of mobile data traffic, as we saw earlier, makes delivery efficiency a line item on the P&L, and a video-native transport aims to move the same content using less capacity. Codec efficiency pulls the same direction from a different place, and the two should not be confused. MOQ is transport. Codecs are compression.

Codec transitions are the slowest decisions in video and the most consistently over-predicted. Every generation arrives with a compression advantage and a forecast attached, then spends years waiting on device reach, licensing clarity,


and a workload large enough to pay for re-encoding the library.

Which is why production and intent tell two different stories. H.264 leads production at 85%, with HEVC at 58%, because switching costs keep legacy formats in place long after a newer codec is technically ready. Planning intent puts AV1 first at 35%, ahead of HEVC at 26% and VVC at 20%, which is where the momentum is headed next, layered on top of the production base rather than replacing it.

Codec adoption depends on three things. Playability across the devices you ship to, a workload where the bandwidth saving

clears the encoding cost, and licensing terms someone will sign. AV1 leading intent at 35% says teams expect those conditions to line up. Production share still sits with H.264 at 85%. A methodology note on these figures: this survey did not capture respondent type as publisher


versus technology or service provider. Codec figures like AV1 and VVC adoption may skew toward what service providers are ready to support rather than what publishers are actively deploying. A future edition could ask this directly and report the two populations separately.



CDN: The delivery infrastructure layer

Most teams no longer bet on one provider. The split below, across multi-CDN, single-vendor, and in-house delivery, shows how that consolidation has already played out. Multi-CDN is the default. 52% of respondents run multiple

third-party CDNs, and a further 28% rely on a single third-party provider. In-house (DIY) CDN sits at 10%, hybrid third-party-plus-DIY at 8%, and 2% deliver with no CDN at all.



CDN strategy reflects the delivery maturity that already exists at the layer above the transport. Third-party CDN usage has stayed the dominant approach every year since 2020. In-house DIY CDN has declined from 31% in 2021 to 10% today, as organizations recognize that the complexity of running your own CDN rarely justifies the cost relative to multi-provider resilience. The CDN layer already works well. MOQ proposes a video-native transport foundation underneath a delivery infrastructure that currently runs on borrowed protocol.

The stake in the ground is this. If the efficiency case holds and standardization completes, the first serious MOQ production deployments could arrive as early as Q1 2027, though that timeline depends on two things that have not happened yet. The draft is on the Standards Track and moving toward Working Group Last Call, with no RFC published yet, and a compelling production use case at scale has not been demonstrated. But the direction is set. Video outgrew the transport it was built on. MOQ is the first credible attempt to fix that at the transport layer instead of patching around it, and the survey shows the industry is already looking that way.

“SRT’s rise at contribution confirms that open, interoperable workflows are earning trust at scale. As MOQ extends that shift to distribution, the winners are the teams that stay open across protocols and partners, not locked inside walled gardens.”



Emeka Okoli

SVP Strategic Business Development @ Zixi

WHAT THIS MEANS FOR VIDEO TEAMS

- Track MOQ inside your organization now. The ecosystem is forming and teams that have done the evaluation before it ships will move faster when it arrives.
- The business case to make internally for MOQ is cost and efficiency, not latency. Frame it as a delivery infrastructure investment.
- If you run live sports, betting, or auction workflows, the contribution-to-distribution path is where MOQ’s first production value will appear. SRT’s lead in contribution shows the upstream shift is real.


Methodology

This is our 10th year of running our Bitmovin Video Developer Report. The survey was conducted in July 2026. The responses provided strong representation of the global market, coming from a diverse range of video developers and industry experts working at companies located in 94 different countries. Participation by region was led by North America with 45% of our participants, followed by EMEA at 38%, the Asia Pacific region at 14% and Latin America with 2% of total responses.


As a number of the questions are multiple choice, you will notice that a number of these multiple-choice/answer questions have an aggregate that will not add up to 100%. In these cases, the number indicates the percentage of respondents that chose each option. Like last year, we have used a ranking option for a number of the questions and have adapted some of the questions to align with the recent advancements in technology development.

Please keep in mind that the survey was open for everyone to participate, but the results may be somewhat skewed toward the Bitmovin customer base.


COMPANY SIZE




REGION



JOB TITLE



INDUSTRY



About Bitmovin

Bitmovin is a global leader in video streaming technology, powering the encoding, playback, observability, and monetization of video for the world's largest broadcasters, streaming services, and enterprises. Trusted by more than a third of the world's top streaming platforms, its cloud-native platform processes over 14 million minutes of video encoding, 6,000+ live events, and 4.4 billion plays every month, backed by 78 billion observability data points that help streaming teams catch playback issues before viewers ever notice them.

Founded in 2013 out of academic research and a Y Combinator cohort, Bitmovin co-created the MPEG-DASH standard and built the first commercial adaptive streaming player, a developer-first heritage

that continues today through ATHENA, its dedicated research lab, and a portfolio of 30+ patents in video technology. That same spirit of innovation is now bringing AI directly into the video workflow, from AI-powered scene analysis to agentic tools that let teams test, monitor, and optimize streaming operations in real time.

Bitmovin serves 400+ customers worldwide, supported by a global team and one of the industry's largest partner ecosystems, from offices in Vienna, London, Berlin, Denver, and Klagenfurt.

For more information:

www.bitmovin.com

FURTHER AUTHORIZING CONTRIBUTIONS



Ian Baglow

Co-CEO @ Bitmovin



Stefan Lederer

Co-CEO and Co-Founder @ Bitmovin

