




IBC 2026

Demo Showcase





Demos

Click a solution area to explore demos.

- [Content Creation](#)
- [Media Lifecycle Management](#)
- [Multiplatform Distribution](#)
- [Revenue Generation](#)
- [Builder Zone: AWS Elemental Media Services](#)

At IBC 2026, AWS shows how cloud infrastructure and agentic AI services remove traditional barriers between creative vision and global reach. AWS demonstrates how agentic and generative AI accelerate content creation, automate workflows across the media supply chain, and deliver personalized, profitable viewing experiences at scale.

By providing the most purpose-built capabilities across Content Creation, Media Lifecycle Management, Multiplatform Distribution, and Revenue Generation, AWS empowers media organizations to create without limits, transform archives into revenue engines, deliver the highest quality video across every screen, and turn audience data into growth, ensuring every story has the technical foundation to reach and engage audiences at whatever scale the narrative deserves.





Table of Contents

→ [Content Creation](#)

[Automate ingest and accelerate dailies processing](#)

[Build an AI-powered studio on AWS](#)

[Streamlining production workflows with agentic AI](#)

[Studio operations intelligence: from engineer to editor](#)

→ [Media Lifecycle Management](#)

[Agentic AI rights and contracts analysis](#)

[Agentic promo compilation, localization, and reframing](#)

[AI-driven archive migration](#)

[Content compliance with agents](#)

[Media Lake on AWS: One catalog for every content stage](#)

→ [Multiplatform Distribution](#)

[AI-accelerated smart stories for news](#)

[AI-driven archive monetization and channel origination](#)

[Intelligent agentic operations for broadcast and streaming events](#)

[Multi-vendor live production on AWS using MXL](#)

[Near-live highlights with AWS Elemental Inference](#)

[Personalizing the viewing experience at broadcast scale](#)

[Power live fan experiences with a sports data platform](#)

[Reinventing linear contribution and distribution on AWS](#)

→ [Revenue Generation](#)

[AI-powered brand sponsorship measurement for live broadcasts](#)

[Agent-powered ad sales: from manual workflow to conversational commerce](#)

[Agentic yield management for prebid](#)

[Intelligent ad breaks for live sports with AWS Elemental Inference](#)

[Modernize advertising operations with open data architectures](#)

[Monetize AI-crawler traffic with intelligent bot control](#)

→ [Builder Zone](#)

[Dynamic multi-view for the next generation of live video](#)

[Live contextual advertising using AWS Elemental Inference](#)

[Seamless ad monetization for creator-driven video](#)






Content Creation

Content Creation demos show how AI-enabled tools and agentic AI workflows empower creators and studios at every production stage. From AI-powered visual asset generation and cloud-scale rendering with Deadline Cloud to agentic production planning and intelligent studio operations, see how AWS compresses multi-month production cycles into days, removes workflow bottlenecks, and gives every team member data-driven visibility across the production, with the flexibility to scale resources on demand across any format or channel.


- [Automate ingest and accelerate dailies processing](#)
- [Build an AI-powered studio on AWS](#)
- [Streamlining production workflows with agentic AI](#)
- [Studio operations intelligence: from engineer to editor](#)

Automate ingest and accelerate dailies processing



Demo description

Production processing pipelines built with on-premises hardware force a hard choice: over-provision for peak shoot days or accept backlogs on the days that matter most. This demo shows how connecting the workflow to the cloud changes that equation. Once camera originals are ingested to Amazon S3, AWS Deadline Cloud scales transcode and proxy rendering elastically, spinning up hundreds of workers on heavy shoot days and releasing them when the schedule eases. Because cloud compute is billed by the second, one machine running for 1,000 minutes costs the same as 1,000 machines running for one minute. Scaling out to clear a backlog adds no premium over processing sequentially. Finished dailies and proxies flow into the editor's system without manual handoffs, while media stays on Amazon S3 as the single source of truth. The result: productions clear processing backlogs in minutes instead of hours, without paying for hardware that sits idle between shoots.



Demo partners

Adobe




Suite

Related content


- [Getting started with Deadline Cloud](#)
- [AWS Deadline Cloud code samples](#)
- [AI-powered render troubleshooting with the new Deadline Cloud Assistant](#)
- [Deadline Cloud: Configurable job scheduling](#)
- [Deadline Cloud: Host configuration scripts](#)
- [Suite Studios: S3 File Streaming](#)
- [Origami Platform](#)
- [Frame.io and Avid Media Composer](#)
- [Avid on AWS](#)

Build an AI-powered studio on AWS



Demo description

The gap between a creative idea and a finished visual asset has always been measured in time: time to generate, time to review, and time to start over when the note comes back. This demo shows the latest capabilities of an AI-powered studio on AWS, where content creators generate, augment, and iterate on visual assets at unprecedented speed. Artists can utilize the best generative model for each task through one environment, backed by Amazon Bedrock, without rebuilding their pipeline each time a better model appears. By using an AI-powered studio on AWS, creative teams can reduce overall production timelines while keeping artists in control of creative decision-making.



Demo partners



Related content

- [Guidance for an AI-Powered Studio on AWS](#)
- [AWS Deadline Cloud: FLUX.2 Klein LoRA fine-tuning and image generation](#)
- [Scale LLM fine-tuning with Hugging Face and Amazon SageMaker AI](#)
- [Luma AI in Amazon Bedrock](#)
- [AWS Marketplace: Kling AI](#)
- [Prime Video Orders Three New Animated Series from GenAI Creators' Fund, Backed by Amazon MGM Studios and AWS](#)



aws FOR MEDIA AND ENTERTAINMENT


DEMO WALK-THROUGH

Streamlining production workflows with agentic AI

Jack Le Bon
Assoc. Solutions Architect, Media & Entertainment
AWS

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

From pre-production through distribution and marketing, film and TV teams spend significant hours on coordination work that does not add creatively to the final product. This demo shows how agentic AI gives those hours back. Purpose-built AI tools address the most repetitive bottlenecks across talent management, production planning, social content, and AI tool selection, automating undifferentiated coordination work that distracts from more impactful work across the production. Each capability targets a specific category of manual effort across the production lifecycle. The result: productions move faster, with fewer handoffs, less friction, and more time for the creative decisions that only people can make.

 Social Department


- [Guidance for an AI-Powered Studio on AWS](#)
- [Building an enterprise-ready generative AI platform on AWS](#)

aws FOR MEDIA AND ENTERTAINMENT

DEMO WALK-THROUGH

Coming soon

© 2016, Amazon Web Services, Inc. or its affiliates. All rights reserved.



In fast-moving production environments, engineering, production, and editorial teams often rely on separate tools and disconnected views, making it harder to stay aligned, spot issues early, and keep projects on track. This demonstration shows how AI-powered assistants bring everyone onto the same page through a single conversational interface, applied across three distinct user stories: engineering support querying instance health, scaling events, and cost allocation on-demand; a non-technical producer asking natural language questions about environment status, schedules, and budget; and editors getting contextual guidance on troubleshooting and workflow optimization. Each user gets data-driven insights relevant to their responsibilities, enabling faster and more confident decision-making. The outcome: studios gain production operations visibility across technical and non-technical stakeholders, optimizing costs and timelines while reducing risk.

Related content

- [Get started with Amazon Bedrock AgentCore](#)
- [Build multi-agent site reliability engineering assistants with Amazon Bedrock AgentCore](#)
- [Build a FinOps agent using Amazon Bedrock AgentCore](#)
- [Autodesk Flow Production Tracking](#)






Media Lifecycle Management

Media Lifecycle Management demos reveal how Emmy Award-winning cloud-native services and agentic AI pipelines transform media supply chains and archives from cost centers into revenue engines. From AI-assisted promo assembly and automated metadata creation to content compliance verification and AI-powered localization at scale—see how intelligent orchestration replaces manual processes across the entire content lifecycle.

- [Agentic AI rights and contracts analysis](#)
- [Agentic promo compilation, localization, and reframing](#)
- [AI-driven archive migration](#)
- [Content compliance with agents](#)
- [Media Lake on AWS: One catalog for every content stage](#)

The image shows the AWS logo and the text "FOR MEDIA AND ENTERTAINMENT" in the top left corner. Below this, the text "ZERO WALK-THROUGH" is displayed. The main title "Agentic AI rights and contracts analysis" is prominently featured in the center. A large play button icon is overlaid on the title. In the bottom left corner, the name "Randall Potter" is listed, followed by his titles: "Senior Solutions Architect and Generative AI Subject Matter Expert" and "AWS". At the very bottom, a small copyright notice reads: "© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved."

Every day a title sits in rights limbo is a day of lost revenue. For broadcasters and streamers managing deep catalogues, confirming "clear to air" means manual contract review, cross-functional coordination, and legal bottlenecks that delay distribution. This demo shows an agentic AI architecture powered by Amazon Bedrock that reads complex licensing agreements, extracts territory, window, and exclusivity clauses, and cross-references them against distribution plans in real time. The agent surfaces conflicts and recommends clearance paths while preserving human editorial control for nuanced judgment calls. Rights clearance processes that previously took weeks now resolve in minutes, proven at scale with Condé Nast.



Related content






FAN ENGAGEMENT | DERO WALK-THROUGH

Agentic promo compilation, localization, and reframing

Tobias Nitzsche
Solutions Architect
AWS

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.


Rightsholders of live events and deep archives need highlight reels in every language, for every platform, faster than any manual team can cut them. This agentic workflow ingests content from an archive or a live event still in progress and automatically identifies key moments worth cutting into a reel. A producer can review and approve the selection, then the workflow takes over: intelligent reframing to vertical and platform-native aspect ratios using AWS Elemental Inference, commentary generation in six languages with Amazon Bedrock, voiceover using synthetic or cloned voices, and automated packaging of every version for delivery with AWS Elemental MediaConvert. A second approval gate ensures editorial oversight before publishing to a CMS, Media Lake on AWS, or other distribution endpoint. Drawing on implementations with Sinclair, this demo shows how content owners can reach every audience on every screen, while editors retain full editorial control and turnaround drops from days to minutes.



- [Guidance for a Media Lake on AWS](#)
- TwelveLabs video understanding models are now available in Amazon Bedrock




AI-driven archive migration



Demo description

Broadcasters and content owners hold vast archives locked on aging tape, on-premises MAMs, and regional silos. Buried within those archive shelves are rights-cleared assets that could be earning revenue today if anyone could find them. This demo walks you through the end-to-end fix as a single guided journey. An interactive chatbot builds your business case, estimates what the archive costs today, and shows you both how to optimize cloud storage and what revenue migration could unlock. Next, a governance dashboard takes over, orchestrating the migration of tens of thousands of files, giving you live visibility into throughput, completion forecasts, and emerging risks. Finally, an agentic workflow links your assets against a metadata repository and adds content understanding that opens up new revenue opportunities by making it easy to create shorts, promos, and other derivative content. Walk away with a clear business case: lower storage cost, no more physical media risk, and a measurable path from dormant archive to active monetization.



Related content


- [BBC Archive Case Study](#)
- [Media Supply Chain & Archive on AWS](#)
- [Aligning AWS to MPA security best practices for media archives](#)
- [Modernizing NASCAR's multi-PB media archive at speed with AWS Storage](#)
- [Unlock the value of Media Archives with AWS](#)

Content compliance with agents



Demo description

Broadcasters and streaming platforms struggle with manual content compliance reviews across rights clearance, title management, and regulatory standards—a time-consuming process that doesn't scale efficiently. This demonstration shows how AI agents can automate compliance analysis, integrating directly with Media Asset Management systems like Fonn Group's Mimir. Attendees will see automated workflows that check content against compliance standards, detect profanity, and apply custom compliance profiles—all triggered seamlessly from within existing platforms. The result: faster review cycles, consistent adherence to standards, and a cost-effective compliance process that scales from small creators to large enterprises.



Demo partners



Related content

- [Streamlining content compliance: Automating media analysis with Amazon Nova](#)




DEMO WALK-THROUGH

Media Lake on AWS: One catalog for every content stage

Robert Raver
Principal Solutions Architect
AWS


© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Media teams run disparate systems for each stage of the lifecycle - ingest, versioning, enrichment, search, transformation and distribution. Teams spend more time reconciling these systems than actually working with the content itself. The Guidance for a Media Lake on AWS replaces that fragmentation with a single catalog that spans every stage of the content lifecycle. The media lake registers files where they already live, treating each one as a logical asset with its variants and components. Analysis, content understanding, search, transformation and creation all happen against the catalog. Metadata is versioned, schema-defined, and human-reviewed when confidence is low – promoted or rolled back, never overwritten. A knowledge graph provides full lineage: trace any clip to its source, processing run, and the model behind each value. One catalog. Full lineage. One governed view from ingest to delivery.



- [Guidance for a media lake on AWS](#)
- [Guidance for a media lake on AWS: Accelerate media operations with agentic AI](#)
- [Using Amazon S3 Vectors \(preview\) to semantically search using a media lake on AWS](#)






Multiplatform Distribution

Multiplatform Distribution demos showcase how AWS enables broadcast and streaming organizations to deliver the highest quality video across every screen. From AI-accelerated newsroom workflows with verified provenance and intelligent live event operations to hyper-personalized streaming, near-live highlights powered by AWS Elemental Inference, multi-vendor cloud production over MXL, and cloud-native linear contribution and distribution, see the unmatched agility, elasticity, scalability, and reliability that enhances discoverability, reaches new audiences, and transforms viewer engagement.

- [AI-accelerated smart stories for news](#)
- [AI-driven archive monetization and channel origination](#)
- [Intelligent agentic operations for broadcast and streaming events](#)
- [Multi-vendor live production on AWS using MXL](#)
- [Near-live highlights with AWS Elemental Inference](#)
- [Personalizing the viewing experience at broadcast scale](#)
- [Power live fan experiences with a sports data platform](#)
- [Reinventing linear contribution and distribution on AWS](#)




AI-accelerated smart stories for news



Demo description

Newsrooms produce stories for more platforms than ever, yet their tools remain disconnected: producers copy metadata between systems and manually verify sources under deadline pressure, risking brand integrity with every shortcut. Industry pressure is dissolving the divide between media and editorial teams — a convergence that becomes an opportunity when workflows share a common thread of data. This demo shows how live feeds land in TAMS, an open-source, frame-addressable media store, where agentic AI orchestrated using Amazon Bedrock AgentCore transcribes, indexes, enriches, and verifies footage in near real time, cross-referencing claims and surfacing confidence scores. From a single verified source, the system assembles multiple outputs — a social clip, a bulletin package, and editorial copy — each with signed provenance. The journalist stays in control throughout, accelerated by the technology rather than replaced by it to create one verified story that can be published in multiple formats.



Demo partners




Related content

- [Time Addressable Media Store - Home of the TAMS community and latest news](#)
- [IBC INCUBATOR 2026: Smart Stories - The New Knowledge Architecture for Content Production](#)
- [Amazon Bedrock AgentCore: The platform for production AI agents. Any framework. Any model. Secure at scale.](#)




AI-driven archive monetization and channel origination



Demo description

Linear channels still run on rigid, hardware-bound playout that demands manual scheduling and constant staffing, and launching a new channel takes months. This demonstration reimagines playout as a cloud-native, AI-driven operation built on TAMS as an open media lake. This enables broadcasters to discover archive content through transparent enrichment, scene detection, and automated localization. They can then create channels, add content, and deliver localized feeds to partners on demand. A conversational agent reads audience demand and proposes channels the operator approves in a click, while AWS Elemental MediaTailor assembles rights-compliant linear feeds automatically. The result: operations teams launch and localize channels in a fraction of the usual time, freeing staff from manual scheduling and reaching new audiences faster.




Demo partners



Related content


- [Live streaming localization and accessibility using AWS Media Services](#)
- [Configure live captions with SyncWords and AWS Elemental Media Services](#)
- [Translate live sports automatically to reach international fans with AWS Media Services and SyncWords](#)
- [Create a scalable workflow using the Intel Library for Video Super Resolution](#)
- [Delivering low-latency captions and voice translation for live sports, news, and OTT platforms with SyncWords and AWS](#)
- [How to Use Channel Assembly with AWS Elemental MediaTailor to Launch Virtual Channels from Existing Sources](#)
- [ITG: How Shoppable CTV Ads are Transforming Viewer Engagement](#)

Intelligent agentic operations for broadcast and streaming events



Demo description

Live events like sports, award shows, and breaking news demand flawless technical performance, but identifying and resolving issues across complex broadcast and streaming workflows is challenging under time pressure. Broadcast and streaming engineers can accelerate troubleshooting and ensure optimal performance throughout events by applying agentic AI across the full signal chain. Using Amazon Bedrock AgentCore and AWS Strands SDK, this demonstration shows how specialized AI agents transform media operations to identify root causes of technical issues end-to-end across broadcast and OTT workflows. Using AI assistants, Amazon Kiro, and Amazon Quick, and integrating observability data from AWS Media Services, CMCD video player data and streaming analytics from AWS Partners, Mux, Hydrolix and TAG Video Systems, teams can achieve full-chain observability from contribution to distribution. This demonstration shows how to achieve faster root cause identification, reduce incident response time, and ensure high-quality experiences that keep subscribers engaged when it matters most.



Demo partners




Related content

- [Agentic Analytics with Amazon Bedrock AgentCore and Hydrolix](#)
- [Agentic Intelligent Media Operations](#)
- [Build an intelligent video streaming monitoring solution with CMCD and MCP](#)
- [Hydrolix MCP Server on GitHub](#)
- [Monitor your Video Infrastructure with Hydrolix and Mux Data](#)
- [MUX MCP Server on GitHub](#)

Demo description

The broadcast industry's move to software-defined production requires interoperability through a shared media exchange layer rather than standardizing applications or creating a universal API. This demo delivers that model. A complete multi-camera sports production control room (PCR) runs live on AWS, with independently developed applications via AWS Partners for vision mixing, replay, audio, and graphics exchanging uncompressed media over the Media eXchange Layer (MXL) v1.0.0, the open-source specification published in March 2026. MXL's shared memory transport delivers sub-millisecond latency with no generation loss. This demonstration showcases how a multi-vendor approach using one open standard can deliver broadcast production quality at low latency.



Demo partners



Related content

- [MXL is the foundation. Here's what it takes to build the house](#)
- [MXL GitHub](#)
- [IBC Accelerator: Multi-Vendor Software Live Media Exchange](#)
- [Dynamic Media Facility](#)
- [AMWA Joint Task Force on Dynamic Media Facilities](#)




Near-live highlights with AWS Elemental Inference



Demo description

Broadcasters now publish live content to more platforms than ever — TikTok, Reels, Shorts, linear — each demanding its own format and editorial angle. Manual workflows simply cannot keep pace. This demo shows how AWS Elemental Inference integrated with AWS Elemental MediaLive runs AI in parallel with encoding to deliver publish-ready content within seconds. Intelligent reframing transforms a single live feed into vertical, captioned, and clipped formats automatically. New contextual metadata capabilities — shot detection, transcription, and content tagging — power a clipping interface where editorial teams navigate hours of content in seconds. An agentic workflow built on Amazon Bedrock recommends highlight selections from natural-language prompts with no ML expertise required. The content is prepped for every format and every platform, while the event is still live.



Demo partners

#tagboard



Related content


- [AWS Elemental Inference](#)
- [AWS Elemental MediaLive](#)
- [AWS Elemental MediaPackage](#)
- [Personalize live viewing, monetize every impression, and reach every device with AWS Elemental at IBC 2026](#)
- [Discover how customers are using AWS Elemental Inference](#)
- [Generate live subtitles with AWS Elemental Inference Smart Subtitles](#)
- [Transform live video for mobile audiences with AWS Elemental Inference](#)
- [Using AWS Elemental Inference with MediaLive](#)
- [How AWS Built a Live AI-Powered Vertical Video Capability for Fox Sports with AWS Elemental Inference](#)

Personalizing the viewing experience at broadcast scale



Demo description

Streaming platforms compete less on catalog size now and more on whether each session feels made for the person watching. This demonstration shows contextually aware streaming that reshapes recommendations, user interface (UI), and experience in real time based on audience interactions with live content. Amazon Bedrock AgentCore orchestrates a multi-agent workflow, including a video highlight agent, while Amazon Personalize drives audience segmentation and recommendations. AWS Elemental Inference powers automated sports highlight detection, surfacing key moments from live content as they happen. Attendees will see how streaming platforms move beyond volume to relevance, lifting engagement and retention by meeting viewers with an experience that adapts to their preferences.



Demo partners



Related content

- [AWS GitHub Sample: Speech to Speech Movie Recommendation Solution](#)
- [Deliver hyper-personalized viewer experiences with an agentic AI movie assistant using Amazon Bedrock AgentCore and Amazon Nova Sonic 2.0](#)
- [Amazon Bedrock AgentCore](#)
- [What is Amazon Personalize?](#)



FOR MEDIA AND
ENTERTAINMENT


DEMO WALK-THROUGH

Power live fan experiences with a sports data platform

Rizwan Qureshi
Principal Specialist Solution Architect
AWS

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.


A sports live stream is no longer a single viewing experience. Today's fans engage through multiple dimensions simultaneously — immersive 3D match reconstructions, gamification layers, enhanced statistical overlays, AI-powered companions, and contextual commerce. These experiences are driven by data, and AI has accelerated both their creation and orchestration. This demo shows the platform that powers them. A real-time sports data pipeline ingests low latency player tracking, processes it through ML models to produce match facts, event data, and tactical intelligence, then orchestrates downstream experiences through event-driven architecture and agentic AI. Built on AWS with Amazon SageMaker for inference and Amazon Bedrock for contextual intelligence, one platform serves every screen: broadcast overlays, mobile apps, streaming experiences, and interactive participation.



- [AWS is how fans get closer to the world's game](#)
- [How Bundesliga built Captain, an AI agent for fans, using Amazon Bedrock](#)
- [Bundesliga transforms data into AI-powered fan-favorite stories with AWS](#)
- [Behind the scenes: AWS real-time architecture for Bundesliga Match Facts](#)




Reinventing linear contribution and distribution on AWS



Demo description

As broadcasters increasingly move to IP-based distribution, the need to manage and control live linear feeds across global networks becomes critical. This demonstration shows how AWS Elemental MediaConnect Router receives and manages live and linear feeds from global contributors, through both a custom-built frontend UI and a native Console UI. It also highlights monitoring and automated alerting capabilities powered by TAG Video Systems, alongside virtual headend capabilities with the Zixi Zen Master solution, supporting SCTE-224 and virtual Integrated Receiver Decoder (IRD) features. We will also be showing how live contribution feeds can be ingested into a TAMS store and played out directly using the BCNexxt playout solution. Together, these give customers a clear path to manage their entire terrestrial distribution and linear routing in AWS, with the flexibility to scale across regions and reduce reliance on rigid, high-cost legacy infrastructure. There are many ways customers can implement this workflow. A real-world use case is Thomson Reuters, where the MCR/TV Solutions team distributes live content (N+ Lives) over IP, demonstrating how Reuters live feeds flow in as contribution and are routed out as linear distribution, delivering cloud-native routing at a lower cost and higher flexibility.



Demo partners



Related content

- [Global Cloud-Based Live Video Delivery with AWS and Zixi](#)
- [AWS Elemental MediaConnect Flows Managed & Monitored in ZEN Master](#)
- [TAG Video Systems: Unlocking cloud-based quality of experience \(QoE\) management with TAG and AWS](#)
- [Zixi ZEN Master Control Plane for AWS Elemental MediaConnect and AWS Services](#)




Revenue Generation

Revenue Generation demos show how media organizations and advertising companies combine first-party data with AI to modernize ad tech stacks and diversify monetization strategies. From AI-powered brand sponsorship measurement and agent-powered ad sales to intelligent ad breaks for live sports with AWS Elemental Inference, agentic yield management on the open-source Prebid stack, and a new way to monetize AI-crawler traffic, see how agentic AI-powered automation maximizes revenue across advertising, subscriptions, and content licensing.


- [AI-powered brand sponsorship measurement for live broadcasts](#)
- [Agent-powered ad sales: from manual workflow to conversational commerce](#)
- [Agentic yield management for prebid](#)
- [Intelligent ad breaks for live sports with AWS Elemental Inference](#)
- [Modernize advertising operations with open data architectures](#)
- [Monetize AI-crawler traffic with intelligent bot control](#)

AI-powered brand sponsorship measurement for live broadcasts



Demo description

Sponsorship renewals depend on proof of delivery, yet most rights holders wait weeks for manual audits. This demo processes a live sports broadcast using Amazon Transcribe for audio transcription, Amazon Bedrock Data Automation and Amazon Bedrock foundation models for logo detection, OCR, and brand recognition, storing results in an Amazon S3 Tables-backed data lake. Every sponsor is ranked by Share of Voice with video and audio evidence and on-screen overlays proving where and when each brand appeared, broken down by inventory type. Amazon Quick lets sponsorship teams query results in natural language. Walk into renewals with same-day proof, no manual audits or pre-configured brand lists.



Related content

- [Automate video insights for contextual advertising using Amazon Bedrock Data Automation](#)
- [Analyze media content using AWS AI services](#)




DERO WALK-THROUGH

Agent-powered ad sales: from manual workflow to conversational commerce

Adrien Lebert
Solutions Architect, Media & Entertainment
AWS

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Direct-sold advertising commands premium CPMs, but the sales process is manual, slow, and engineering-dependent. Campaign setup, creative review, and inventory allocation still require human handoffs at every stage. This demo shows how AI agents built on Amazon Bedrock automate the direct sales workflow for publishers running Prebid Server without replacing your team or locking you into proprietary platforms. A Prebid Sales Agent handles campaign management, product ranking, and creative policy review through natural language, giving sales teams an agent-accessible interface to premium inventory. The agents communicate over open protocols, connecting demand- and supply-side systems without vendor lock-in. The result: faster deal execution, fewer manual touchpoints, and direct-sold revenue captured at the speed of programmatic.



- [Guidance for Advertising Agents on AWS](#)
- [Ad Context Protocol \(AdCP\)](#)
- [Amazon Bedrock AgentCore](#)




Agentic yield management for prebid



Demo description


Publisher ad-ops teams leave yield on the table because bidder configurations and price floors are locked in code, requiring engineering tickets for every change. This demo puts control back in business hands. AI agents built on Amazon Bedrock let ad-ops users manage Prebid Server bidder settings and price floors through natural language, with every change governed and auditable. AWS RTB Fabric analyzes bid-stream data in real time, surfacing optimization signals and dynamically shaping traffic to maximize fill and CPM. Built on the open-source Prebid stack with no vendor lock-in, the solution is designed for broadcasters, digital publishers, and streaming platforms seeking incremental programmatic revenue without incremental engineering cost.



Related content


- [Guidance for Deploying a Prebid Server on AWS](#)
- [Prebid Server \(open source\)](#)
- [Amazon Bedrock AgentCore](#)
- [Deploy MCP servers on AgentCore Runtime](#)
- [AgentCore Gateway — MCP aggregation](#)
- [Amazon Quick — MCP integration](#)

Intelligent ad breaks for live sports with AWS Elemental Inference



Demo description

Ad placement in live sports is still a manual job. This demo makes it fully automated. AWS Elemental Inference generates real-time contextual metadata from the live feed — scenes, actions, and speech — as it happens. The operator writes the ad strategy once, in plain language, and Amazon Bedrock turns it into rules. From that point, no human is needed: the engine matches the live metadata against the rules and triggers each ad automatically, at the right moment and in the right format.



Related content

- [Generate live subtitles with AWS Elemental Inference Smart Subtitles](#)
- [Using AWS Elemental Inference with MediaLive](#)
- [Sample Elemental Inference Video Clipping Portal](#)

Modernize advertising operations with open data architectures



Demo description


Broadcasters and streaming platforms sit on mountains of first-party audience data but struggle to activate it for revenue. Fragmented profiles, siloed viewing signals, and opaque churn patterns leave CPMs and fill rates on the table. This demo shows how an open data architecture built on Amazon S3 and Amazon Bedrock unifies customer profiles, segments audiences by behavior, and models retention risk. Amazon Quick serves as the conversational analytics layer, letting content and revenue strategists ask natural-language questions about their audiences. A CDP or clean-room partner owns the customer-facing activation, while AWS provides the scalable, open foundation. The result: higher yield, addressable ad revenue, and data-driven programming decisions.



Related content

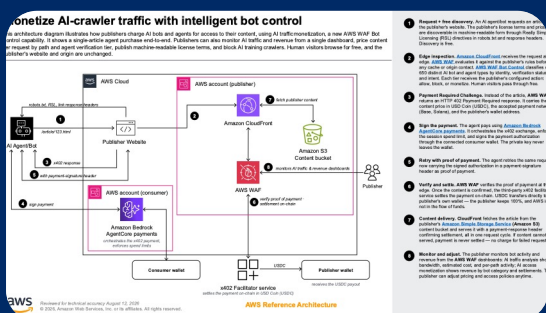
- [Guidance for Customer Data Platform on AWS](#)

Monetize AI-crawler traffic with intelligent bot control



Demo description

AI crawlers now generate more traffic than human visitors on many publisher sites, yet that traffic produces zero revenue and inflates infrastructure costs. This demo shows how publishers can take back control. AWS WAF Bot Control identifies and classifies AI-crawler traffic in real time, giving publishers granular control over which bots access their content and under what terms. The demo layers a licensing and pay-per-crawl monetization framework on top of WAF signals, enabling publishers to charge for authorized AI training access while blocking unmonetized scraping. The result: a net-new revenue stream from web traffic that previously hasn't been monetizable, with content protection enforced automatically.



Related content

- [AWS WAF announces AI traffic monetization](#)
- [AWS WAF adds AI traffic monetization capability to help content owners charge AI bots for content access](#)
- [AI traffic monetization](#)
- [Agents that transact: Introducing Amazon Bedrock AgentCore payments, built with Coinbase and Stripe](#)



Builder Zone

Builder Zone demos showcase the latest AWS Elemental Media Services innovations that push the boundaries of live video technology. From dynamic multi-view experiences assembled on the fly with AWS Elemental MediaPackage to live contextual advertising powered by AWS Elemental Inference and seamless ad monetization for creator-driven video with AWS Elemental MediaTailor, explore next-generation solutions that give operators interactive viewer control, cookie-free ad targeting tied to what's happening on screen, and scalable monetization for the fastest-growing content formats.


- [Dynamic multi-view for the next generation of live video](#)
- [Live contextual advertising using AWS Elemental Inference](#)
- [Seamless ad monetization for creator-driven video](#)

Dynamic multi-view for the next generation of live video



Demo description

Audiences today expect to choose their own experience, picking camera angles, isolating player feeds, or switching between commentary tracks in real time. But delivering multi-view experiences has traditionally meant duplicating packaging pipelines for every permutation, resulting in increased costs and operational complexity. In this demo, see how AWS Elemental MediaPackage enables Dynamic Multiview experiences from a single packaging endpoint. Passing view-selection parameters directly in the playback URL enables MediaPackage to dynamically assemble the right combination of video, audio, and metadata tracks at request time — no pre-packaged variants, no player-side SDKs, no proprietary extensions.




Live contextual advertising using AWS Elemental Inference



Demo description

Third-party cookies are disappearing and audience-level targeting is losing signal, but advertisers still need to reach the right viewer at the right moment – and the most powerful signal was in the content all along. Elemental Inference turns live video streams into actionable, monetizable moments in near-real-time with contextual ad decisioning tied to what's happening on screen. Contextual targeting is achieved in near-real-time: no cookies, no PII, no post-production delay. For live sports, serve a sponsor creative the instant a star player scores — matched by moment type, player, and excitement level. For live news, identify and monetize brand-safe ad breaks that won't conflict with sensitive on-screen content.





FOR MEDIA AND
ENTERTAINMENT

DEMO WALK-THROUGH

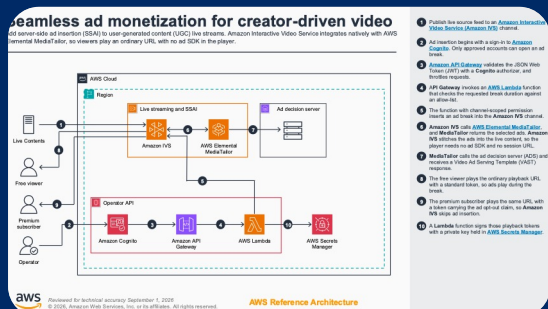
Seamless ad monetization for creator-driven video

▶

Yohan Choi
Senior Specialist Solutions Architect, Media Edge
AWS

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Creator-driven video is one of the fastest-growing segments in media, but monetizing hours of content into advertising revenue remains a challenge. Unpredictable stream lengths, variable video quality, and the sheer volume of concurrent creators make traditional ad insertion approaches difficult to scale. In this demo, see how to build a complete monetization pipeline using Amazon Interactive Video Service (IVS) and AWS Elemental MediaTailor that lets creators go live with minimal setup while delivering broadcast-quality, personalized ad experiences with no buffering, client-side complexity or disruption to the viewing experience.





Create. Connect. Captivate.

Contact us

Learn more